

INTEGRATION OF AIRBORNE LASER SCANNER AND OPTICAL DATA
USING IMAGE TIE POINTS AND THE QUASI-LOG-POLAR
ANALYTICAL FOURIER MELLIN TRANSFORM

MICHAEL S. P. BAUMGAERTNER

A DISSERTATION SUBMITTED TO THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILMENT OF THE REQUIREMENTS FOR THE DEGREE OF
DOCTOR OF PHILOSOPHY

GRADUATE PROGRAM IN EARTH AND SPACE SCIENCE
YORK UNIVERSITY
TORONTO, ONTARIO

November 2023

Abstract

LiDAR and camera have complementary data characteristics. Their combined usage allows for achieving better accuracy and enhanced inference about the environment than through the isolated use of the data of one sensor only. Co-registration is the process of spatial alignment of data from different types of sensors. It enables cross-aided sensor- and data calibration and the combined usage of the data. Automatic co-registration of airborne optical imagery and LiDAR data from large-area surveys has not been addressed, unless the datasets come with metadata or some prior knowledge about the scenery or have been pre-processed and converted into data with similar density and structure.

In this dissertation, an innovative co-registration concept that uses image tie points, projected into object space, and Lidar points, is proposed and researched. Since 3D tie points are generated automatically with centimeter-level precision within the photogrammetric block adjustment, this approach has the potential of providing high accuracy and full automation. However, it must address that large area photogrammetric tie points provide extremely low point density compared to laser point clouds, and they describe different aspects of the scanned surface, since tie points coincide with radiometric corner locations on the visual surface whereas laser points locate anywhere on surface objects and on the ground.

This research provides the following contributions: Firstly, the use of 3D photogrammetric tie points for co-registration of large-format sensor data is proposed and analyzed. Secondly, new methods for coarse- and fine registration in space domain are presented and evaluated. Thirdly, five methods for co-registration in frequency domain are researched. Lastly, the novel Quasi Log-Polar Fourier Mellin Transform (QAFMT) method is suggested as tool for automatic coarse co-registration of large area datasets. The learnings and potential use for other applications are discussed in the end. The performance of the QAFMT and its competing methods is analyzed using one simulated and two real-world datasets. While existing algorithms estimated 45 of 144 co-registration parameter pairs correctly (31% success), the QAFMT yielded a 7% improvement (50/144 accurate pairs, 38% success).

Dedication

*In memory of Herrn Georg Strobel
and with heartfelt love for Patricia Carol.*

Acknowledgements

I would like to express my deep gratitude to Prof. Costas Armenakis and Prof. Jianguo Wang for their constructive and devoted support. Without Costas and Jianguo, a successful completion of my dissertation would have been impossible. A big thank you to my friend Prof. Avesta Goodarzi for guiding me in restructuring and condensing this thesis and the many consultations towards the end of this work and guiding me successfully to the finish line. A heart-felt thank you also to my friend Dr. Hamid Goldman, who supported me throughout the years with fruitful consultations related to my PhD studies and life in general. A warm thank you to S. D. Sharma for sharing his wisdom, reminding me to seek balance, and enjoy the ride. Thank you to Prof. Sunil Bisnath for making positive key decisions, providing lab space, software licenses, a laptop, serving on my committee, and for pushing the advancement of my studies.

A very special thanks to Dr. Peter Friess, freelancing scientist, and consultant at Teledyne Optech, for providing the topic for this dissertation, for helping with the mathematical formulation of the fine registration algorithms, and for supporting my LiDAR calibration-related publications. A warm thank you to my former boss Dr. Rachel Brooks, formerly Manager of the Applications group at Teledyne Optech, for supporting my part-time studies and my first LiDAR calibration-related publication. Thank you, Ken Andrews, Six Sigma Master Black Belt instructor at General Motors, for sharing tools and helping me design the structure of FMT experiments in this thesis. Thank you also to Prof. Shabnam Pejhan for her input.

Thank you also to Scott Andrews, formerly Technical Writer at Trimble Applanix, and to my former boss Peter Canter, formerly Director of the Marine Navigation division at Trimble Applanix, for supporting my PhD studies and advancing my GPS-related research proposal and my internal publication about the moving GNSS base station concept.

A special thanks to Prof. Martin Kada who was so kind to be my external examiner. Thank you to Sarah Howe and Prof. Cheryl Giblon from Innovation York and to Don Carswell, formerly President of Teledyne Optech, for providing and supporting the data

sharing agreement. Thanks also to Prof. Michael Cramer and Prof. Nobert Haala for recommending me to YorkU's PhD program.

A big thank you to my former colleagues at Teledyne Optech, particularly Yu Gao, Livia Theriault, Shahrzad Pandhe, and Martin Pokorny for their help with my LiDAR calibration-related publications and preparation of the LiDAR and camera datasets.

Thank you also to Prof. James Elder for pointing me towards the frequency domain. Thanks to Dr. Bruno Scherzinger, formerly head of R&D at Trimble Applanix, and Prof. Anthony Szeto for serving as committee members during the time when I was researching the GNSS moving reference station PhD topic. Thank you to Marcia Gaynor, Graduate Assistant of the Earth- and Space Science program, for her outstanding assistance throughout the years. Thank you to Prof. Spiros Pagiatakis for serving as chair, and Prof. Pavlos Motakis for serving as examiner.

Lastly, a heart-felt thank you to Dr. Michael Gruber, head of R&D and Applications at Vexcel Imaging, who allowed me to work part-time temporarily so that I could dedicate more time to my studies, and for keeping it private. Thank you to my former fellow PhD students at YorkU (Garrett Seepersad, Paul Collins, Surabi Gautabi, and Guy Bernardi) for providing company, for sharing tips on how to write this thesis, and for welcoming me in the lab. Lastly, a big thank you to my parents who supported my studies in Geodesy and Geomatics, and my studies in Technical Pedagogics in Automotive Engineering at University of Stuttgart, Germany, which enabled me to start my PhD studies in Canada.

Table of Contents

Abstract.....	ii
Dedication.....	iii
Acknowledgements.....	iv
Table of Contents.....	vi
List of Tables	xi
List of Figures.....	xiii
List of Symbols	xix
List of Acronyms	xxi
1.0 Introduction.....	1
1.1 Motivation.....	1
1.2 Problem Definition.....	1
1.3 Thesis Objectives	5
1.4 Research Contributions	6
1.5 Thesis Outline	6
2.0 Literature Review	9
2.1 Introduction.....	9
2.2 Registration Paradigm and Categorization of Registration Algorithms	9
2.3 Registration Primitives and Matching Strategies.....	11
2.3.1 Existing Co-Registration Methods.....	12
2.3.2 Co-Registration Methods of Special Interest	15
2.4 Summary	26
3.0 Test Data and Assumptions.....	28
3.1 Introduction.....	28
3.2 Reference Dataset	28
3.3 Real-World Datasets	30
3.3.1 Oshawa Dataset.....	30
3.3.2 Keyworth-Nottingham Dataset	33
3.4 Simulated Dataset	37

3.5	Assumptions.....	39
4.0	Coregistration in Space Domain.....	44
4.1	Introduction.....	44
4.2	Overlap Check in Space Domain.....	45
4.2.1	Global Cost Functions.....	45
4.2.2	Surface Analysis using Differential Geometry.....	47
4.2.2.1	Definition and Characteristics.....	47
4.2.2.2	Discussion.....	54
4.2.3	Fractal Analysis.....	55
4.2.4	Surface Analysis Using Entropy.....	57
4.2.4.1	Definition and Characteristics.....	57
4.2.4.2	Discussion.....	63
4.2.5	Probabilistic Analysis Using Graphical Models.....	65
4.3	Coarse Registration in Space Domain.....	66
4.3.1	Kalman Filtering.....	67
4.3.2	Accelerated Key Point Matching.....	67
4.3.3	Discussion.....	69
4.4	Fine Registration in Space Domain.....	69
4.4.1	Introduction.....	69
4.4.2	Fine Registration and cross-aided Sensor Calibration using unclassified Tie Points and LiDAR Roof Planes.....	72
4.4.2.1	Methodology and Implementation.....	72
4.4.2.2	Results.....	78
4.4.2.3	Discussion.....	89
4.4.3	Fine Registration using unclassified Tie Points and classified LiDAR Planes	91
4.4.3.1	Methodology.....	91
4.4.3.2	Results and Discussion.....	93
4.4.4	Fine Registration using unclassified Tie Points and locally fitted Surface Patches.....	93
4.4.4.1	Methodology.....	93
4.4.4.2	Results.....	95
4.4.4.3	Results and Discussion.....	97
4.4.5	Fine Registration using classified Tie Points and classified LiDAR Patches	98
4.4.5.1	Methodology.....	98
4.4.5.2	Results and Discussion.....	101

4.5	Conclusions.....	106
5.0	Coregistration in Frequency Domain: Performance Analysis of Existing Methods.....	108
5.1	Coarse Registration in Frequency Domain	109
5.1.1	The Phase Only Matched Filter	109
5.1.1.1	Definition.....	109
5.1.1.2	Properties.....	111
5.1.1.3	Co-Registration using POMF	112
5.1.1.4	Recovery of the Cartesian Translation using POMF.....	116
5.1.1.4.1	<i>Oshawa Data</i>	117
5.1.1.4.2	<i>Keyworth-Nottingham Data</i>	128
5.1.1.4.3	<i>Simulated Data</i>	132
5.1.1.5	Summary.....	136
5.1.2	The Fourier Mellin Transform	138
5.1.2.1	Definition.....	138
5.1.2.2	Co-registration using FMT	140
5.1.2.2.1	<i>The FMT Algorithm</i>	140
5.1.2.2.2	<i>Properties</i>	151
5.1.2.2.3	<i>Considerations about the Design Space and Angular- and Scale Offsets</i> 160	
5.1.2.2.4	<i>Numerical Results</i>	162
5.1.2.3	Summary.....	169
5.2	Conclusions.....	172
6.0	Coregistration in Frequency Domain: Expanded Existing and Novel Methods	174
6.1	The Analytical Fourier Mellin Transform	174
6.1.1	Introduction.....	174
6.1.2	Definition	175
6.1.3	Co-Registration using Log-Polar AFMT	177
6.1.3.1	The LAFMT Algorithm.....	177
6.1.3.2	Properties.....	185
6.1.3.3	Numerical Results	189
6.1.3.4	Discussion.....	195
6.1.4	Co-Registration using Polar AFMT	196
6.1.4.1	The PAFMT Algorithm	196
6.1.4.2	Properties.....	200

6.1.4.3	Numerical Results	203
6.1.4.4	Comparison of LAFMT and PAFMT Angle Estimates	205
6.1.4.5	Discussion.....	210
6.1.5	Co-Registration using Extended AFMT	212
6.1.5.1	The EAFMT Algorithm.....	212
6.1.5.2	Discussion.....	214
6.1.6	Conclusions.....	215
6.2	The Quasi Log-Polar Analytical Fourier Mellin Transform.....	216
6.2.1	Co-Registration using Tightly-Coupled QAFMT.....	217
6.2.1.1	The QAFMT_TC Algorithm	217
6.2.1.2	Properties	221
6.2.1.3	Numerical Results	226
6.2.1.4	Discussion.....	240
6.2.2	Co-Registration using Loosely-Coupled QAFMT.....	241
6.2.2.1	The QAFMT_LC Algorithm	241
6.2.2.2	Properties	243
6.2.2.3	Numerical Results	247
6.2.2.4	Discussion.....	253
6.3	Summary	254
7.0	Conclusions and Recommendations.....	265
7.1	Conclusions.....	265
7.2	Recommendations for future Work	268
8.0	References.....	270
APPENDIX A: Related AI Based Co-Registration Methods.....		293
APPENDIX B: Characteristics of Image-based and Keypoint-based 3D Tie Points and Laser Points.....		298
APPENDIX C: Test Data		311
APPENDIX D: Fine Registration using unclassified Tie Points and unclassified Laser Points using ICP		357
APPENDIX E: Design of Experiments, Reduction of Design Space, and Refined Experiments with Simulated Data.....		365
APPENDIX F: Interpretation of FMT Experiments with Real-World Data.....		399

APPENDIX G: Precise Alignment of Point Clouds to build a Reference Data	403
APPENDIX H: Change of Data Footprint and Entropy when Estimating Data Shifts using POMF	413

List of Tables

Table 5.1: Thresholds used to evaluate FMT-estimated co-registration parameters of real-world height maps.....	163
Table 5.2: FMT self-registration results of Oshawa laser point cloud at scale 0.3048...	164
Table 5.3: Recovery of scale, rotation, and translation using Oshawa laser point cloud at scale 1.0.....	165
Table 5.4: Recovery of scale, rotation, and translation using Oshawa data at scale 3.2808.....	166
Table 5.5: Recovery of scale, rotation, and translation using KN data at scale 0.3048..	167
Table 5.6: Recovery of scale, rotation, and translation using KN data at scale 1.0.....	167
Table 5.7: Recovery of scale, rotation, and translation using KN data at 3.2808.....	168
Table 6.1: Registration of Oshawa height maps using LAFMT.....	190
Table 6.2: Registration of KN height maps using LAFMT	191
Table 6.3: Registration of simulated height maps using LAFMT	194
Table 6.4: Registration of real-world and simulated elevation profiles using PAFMT..	203
Table 6.5: Registration of real-world and simulated elevation profiles using PAFMT to scale 3.2808.....	204
Table 6.6: Comparison of LAFMT and PAFMT results for Oshawa height maps	206
Table 6.7: Comparison of LAFMT and PAFMT results for KN height maps	208
Table 6.8: LAFMT results, in comparison to the PAFMT results for simulated height maps	209
Table 6.9: Registration of Oshawa height maps using QAFMT_TC with the confidence interval of 20%.....	228
Table 6.10: Registration of Oshawa height maps, with scale 3.2808, using QAFMT with the confidence interval of 50%	230
Table 6.11: Registration of KN height maps using QAFMT_TC with the confidence interval of 20%.....	232
Table 6.12: Registration of KN height maps using QAFMT with the confidence interval of 50%.....	234

Table 6.13: Registration of simulated height maps using QAFMT_TC with the confidence interval of 20%	236
Table 6.14: Registration of simulated height maps using QAFMT_TC with the confidence interval of 50%	238
Table 6.15: Registration of Oshawa height maps using QAFMT_LC	248
Table 6.16: Registration of KN height maps using QAFMT_LC	250
Table 6.17: Registration of simulated height maps using QAFMT_LC.....	252
Table 6.18: Number of correct co-registration parameter pairs (angle- and respective scale estimate; min=0, max=4)	256
Table 6.19: Number of correct parameter pairs (min=0, max=4), self-registration of laser point clouds	258
Table 6.20: Number of correct parameters pairs (min=0, max=4), self-registration of tie point clouds	259
Table 6.21: Number of correctly estimated co-registration parameter pairs of laser- and tie point cloud	261
Table 6.22: Number of correctly estimated co-registration parameter pairs (tie- and laser point cloud)	262

List of Figures

Figure 1.1: Challenges of co-registration methods. Drawing after Huang <i>et al.</i> (2021)	3
Figure 1.2: Spatial forward intersection of 2D tie points into object space.....	4
Figure 2.1: Dense tie point cloud from an image pair of a stone statue. Source: Zheng <i>et al.</i> (2013).....	15
Figure 2.2: Extraction of roof peak lines. Source: Choi <i>et al.</i> (2013).....	17
Figure 2.3: Registration of a “Happy Buddha” statue. Source: Makadia <i>et al.</i> (2006)	18
Figure 2.4: Registered surface models and respective height difference map of the overlap area. Source: Akca (2007).....	19
Figure 2.5: Main logical components of the co-registration and cross-aided sensor calibration method suggested by Gneeniss <i>et al.</i> (2015)	22
Figure 2.6: Surface matching residuals of Gneeniss’ British data set using 50 LCPs. Source: Gneeniss <i>et al.</i> (2015)	24
Figure 3.1: Aerial view of the OEEPE test site in Vaihingen an der Enz, Germany. The dark green areas depict forest and water bodies. Source: IFP (2017).....	29
Figure 3.2: Digital terrain model of the Vaihingen and der Enz test site. The terrain variations are exaggerated by factor two. Source: IFP (2017).....	29
Figure 3.3: The Oshawa tie point cloud overlaid with the MATLAB-generated surface mesh	30
Figure 3.4: The Oshawa tie point grid with 81x115 40m cells.....	31
Figure 3.5: The Oshawa laser point cloud overlaid with the surface mesh	32
Figure 3.6: The Oshawa laser point grid with 117x138 40m cells	32
Figure 3.7: Overlay of the Oshawa laser- and tie point cloud and the surface mesh.....	33
Figure 3.8: The KN tie point cloud overlaid with the MATLAB-generated surface mesh	34
Figure 3.9: The KN tie point elevation grid (m) with 36x35 40m cells	35
Figure 3.10: The KN laser point cloud overlaid with the MATLAB-generated surface mesh	36
Figure 3.11: The KN laser point elevation grid (m) with 59x71 40m cells.....	36

Figure 3.12: Overlay of the KN laser points, 10m elevated tie points, and the surface mesh	37
Figure 3.13: The simulated X5d4 DTM (m) and the X5d5 DSM (m) with further modified second peak	38
Figure 3.14: The padded X5d4 DTM (m) image used for co-registration testing.....	38
Figure 3.15: Laser points and photogrammetric tie points of one flight line acquired in the Oshawa area	40
Figure 3.16: The geometric mapping between point cloud-rendered images in terms of translation, rotation- and scale parameters. Illustration after (Bargoti <i>et al.</i> , 2012)	42
Figure 4.1: The functional steps of the proposed co-registration workflow	44
Figure 4.2: Delaunay triangulation of a small set of points in MATLAB	46
Figure 4.4: Parameters used for geometric characterization of a topographic surface. Source: Peckham and Jordan (2007) and (*) Li <i>et al.</i> (2004).....	47
Figure 4.5: Oshawa elevation maps in 40m grid: height (m), slope, aspect, and curvature profiles	49
Figure 4.6: KN elevation maps in 40m grid: height (m), slope, aspect, and curvature profiles	51
Figure 4.7: X5d5 and X5d4 simulated elevation maps (pix): height, slope, aspect, and curvature profiles	52
Figure 4.8: Standard deviation (A1, A2) and range filter (B1, B2) profiles of Oshawa elevation maps (m).....	54
Figure 4.9: Schematic of a fractal snow crystal. Source: Sakuraba (2013)	56
Figure 4.10: Memory-free characteristic of Shannon’s entropy. Source: Sabuncu (2006)	58
Figure 4.11: Characteristics of Shannon’s entropy – effects of mapped image intensity. Source: Sabuncu (2006).....	59
Figure 4.12: Entropies of the Oshawa elevation maps	61
Figure 4.13: Entropies of the KN laser- and tie point elevation maps.....	62
Figure 4.14: Entropies of the X5d5 and X5d4 simulated elevation maps	63
Figure 4.15: Image used as a graph. Source: Förstner (2011) referring to Roscher (2010)	65

Figure 4.16: Graph structures. Source: Förstner (2011) referring to Perez (1998).....	66
Figure 4.17: Co-registration workflow suggested by Persad and Armenakis (2017).....	68
Figure 4.18: Integration of laser points and optical data ⁰	71
Figure 4.19: Orthogonal minimization of the tie-point-to-roof-plane distance	72
Figure 4.20: Main components of the proposed co-registration and cross-aided sensor calibration method	73
Figure 4.21: Sensor driven calibration and co-registration of LiDAR and camera using tie points and roof patches as registration elements.....	74
Figure 4.22: Software components and data flow of the implemented co-registration method.....	76
Figure 4.23: Information of a tie point and proximate planar patch in MATLAB.....	77
Figure 4.24: Automatically extracted tie planes in LMS software (Friess, 2006). Planes in red were included in the LSA	79
Figure 4.25: Examples of very irregularly shaped LiDAR roof polygons (m).....	80
Figure 4.26: Tie points located in corners of surface patches.....	81
Figure 4.27: Examples of image-based tie points located at elevated corner locations and flat on the ground	82
Figure 4.28: Key-point generated tie points situated at elevated corner locations and flat on the ground	84
Figure 4.29: Roof planes and roof peak lines before and after rectification of the laser point cloud	85
Figure 4.30: Visual overlay of LMS roof polygons and imagery of the Oshawa dataset.	86
Figure 4.31: Tie points that fell into the 2D boundaries and close to roof polygons due to spatial misalignment of the point clouds (m).....	87
Figure 4.32: Tie points that fall outside the roof polygon due to spatial misalignment of the point clouds (m)	89
Figure 4.33: Examples of tie points that situate on bare earth surface	92
Figure 4.34: Data integration using both sloped and coplanate polygons as registration elements	92

Figure 4.35: Data integration using locally-optimized surface patches as LiDAR registration elements	94
Figure 4.36: Locally-fit surface patch of (a) first and (b) third order wrap around ground control point 27 of the Oshawa data (m)	95
Figure 4.37: Locally-fit surface patch of (a) first- and (b) third order around ground control point 58 of the Oshawa test site (m)	96
Figure 4.38: Integration of the classified tie points and the classified LiDAR planes for fine registration and cross-aided sensor calibration excluding the tie points in transitional zones	98
Figure 4.39: Integration of tie points and classified LiDAR patches for fine registration	99
Figure 4.40: Fine registration using the classified tie points and the classified laser point patches.....	100
Figure 4.41: Horizontal (a) and vertical (b) differences between 30m resolution SRTM data and 1m resolution reference DEM. Source: Carpenter and Hogarty (2007).....	102
Figure 4.42: Thresholds of two slope based DTM filters for the maximum point difference between a terrain point and its nearest neighboring point. Source: Vosselman (2000).....	104
Figure 4.43: Close-up view of tie- (turquoise) and control points (red) in KN imagery	105
Figure 5.1: The POMF algorithm. Drawing after Bargoti <i>et al.</i> (2012)	113
Figure 5.2: Graphical illustration of the POMF using translated X5d4 DTM.....	116
Figure 5.3: Autocorrelation of Oshawa elevation maps (pix), translated South.....	119
Figure 5.4: Autocorrelation of Oshawa elevation maps (pix), translated East (left)	121
Figure 5.5: Autocorrelation of Oshawa elevation maps (pix), translated South and East	123
Figure 5.6: Cross-correlation of Oshawa elevation maps (pix), translated South (down)	124
Figure 5.7: Cross-correlation of Oshawa elevation maps (pix), translated East.....	126
Figure 5.8: Cross-correlation of Oshawa elevation maps (pix), translated South-East..	127
Figure 5.9: Autocorrelation of KN elevation maps (pix) - expected vs recovered translations	129

Figure 5.10: Cross-correlation of KN elevation maps – expected vs recovered translations	131
Figure 5.11: Autocorrelation of X5d5 and X5d4 (pix) – expected vs recovered translations	133
Figure 5.12: Cross-correlation of X5d4 and X5d5 (pix) – expected vs recovered translations	135
Figure 5.13: POMF that considers N CPS peaks in the search for the correct translation offset	137
Figure 5.14: Computational steps of the standard FMT of image spectra. Illustration after Bargoti <i>et al.</i> (2012).....	141
Figure 5.15: Graphical illustration of the standard FMT workflow using rotated X5d4 DTM (pix).....	150
Figure 5.16: Periodicity of 2D DFT transformation and illustration of discontinuities where the repeated image meets. After Le Maistre Pratt (2000)	154
Figure 5.17: Decrease of discontinuities at image borders through reduction of minimum value and best fitting plane	155
Figure 5.18: Zero-padded image border (pix) creating a “+” artifact in the Fourier spectrum.....	157
Figure 5.19: Multi- “+” artifact in FMT phase spectra and corresponding CPS of a 30deg and 80deg rotated X5d4 DTM (cycles/pix)	158
Figure 5.20: Pugh diagram of the standard FMT.....	160
Figure 6.1: Modification of the FMT into the LAFMT algorithm	178
Figure 6.2: The LAFMT algorithm.....	179
Figure 6.3: Self-registration of a 30deg-offset X5d4 DTM (pix) using LAFMT.....	182
Figure 6.4: Effects of rotation and scale for the X5d4 DTM (pix) in Cartesian (a,b) and log-polar (c,d) base	183
Figure 6.5: CPS of self-registered KN tie point cloud (pix) in (a) original orientation and (b) with flipped axis	185
Figure 6.6: Effects of horizontal translation of the X5d4 DTM (pix) in cartesian (a,b) and log-polar (c,d) base	187
Figure 6.7: Modification of the LAFMT into the PAFMT algorithm	197

Figure 6.8: The PAFMT algorithm.....	197
Figure 6.9: Self-registration of a 30deg offset X5d4 DTM (pix) using PAFMT	200
Figure 6.10: Co-registration of Oshawa and KN point clouds (pix), native scale and orientation, using PAFMT (left) and LAFMT (right).....	202
Figure 6.11: The EAFMT algorithm.....	213
Figure 6.12: The QAFMT_TC algorithm	218
Figure 6.13: Search of CPS peaks for co-registration of Oshawa laser- and tie point clouds (pix) in QAFMT_TC	220
Figure 6.14: CPS of log-polar (left side) and polar (right side) sampled KN point clouds (pix), rotated by 90deg and scaled by 3.2808	223
Figure 6.15: CPS of log-polar (left side) and polar (right side) sampled Oshawa and simulated point clouds (pix), rotated by 90deg and scaled by 3.2808.....	225
Figure 6.16: The QAFMT_LC algorithm	242
Figure 6.17: Self-registration of 90deg rotated X5d4 point cloud (pix) using log-polar (left) and polar (right) AFMT	244
Figure 6.18: Self-registration of scaled X5d4 point cloud (pix) in native orientation using log-polar (left) and polar (right) AFMT.....	245
Figure 6.19: Self-registration of 90deg rotated and scaled X5d4 point clouds (pix) using log-polar (left) and polar (right) AFMT.....	246

List of Symbols

Symbol	Definition
$\mathbb{C}$	complex number domain
$\mathbb{R}$	real number domain
x_p	x coordinate of point cloud image p
y_p	y coordinate of point cloud image p
x_q	x coordinate of point cloud image q
y_q	y coordinate of point cloud image q
Δx	translational offset in image space in x direction
Δy	translational offset in image space in y direction
$s, \Delta s$	scale and scale correction
$\Delta \varphi$	rotational offset in image space
f_p	image of point cloud p when used in Fourier related equations
f_q	image of point cloud q when used in Fourier related equations
$P(x, y)$	image of point cloud p in 2D space domain
$Q(x, y)$	image of point cloud q in 2D space domain
F_p	Fourier transform of point cloud image p
F_q	Fourier transform of point cloud image q
ω_x	frequency in x direction
ω_y	frequency in y direction
ρ	radial coordinate in a polar coordinate system
φ	angular coordinate in a polar coordinate system in radians

r_{xy}	location of maximum translational correlation peak
$r_{s\Delta\varphi}$	location of maximum rotational and scale correlation peak
X, Y, Z	world coordinates of a point
$x, y, 1$	image coordinates of a point
x_0, y_0	camera principal point
f	calibrated focal length
c	nominal focal length

List of Acronyms

Acronym	Definition
AFMT	Analytical Fourier Mellin Transform
AGL	Above Ground Level
AI	Artificial Intelligence
ASL	Above Sea Level
CNN	Convolutional Neural Networks
CT	Computed Tomography
CFT	Continuous Fourier Transform
CPS	Cross Power Spectrum
DEM	Digital Elevation Model
DFT	Discrete Fourier Transform
DOE	Design Of Experiments
DG	Direct Georeferencing
DSM	Digital Surface Model
DTM	Digital Terrain Model
ECEF	Earth Centred Earth Fixed
EO	Exterior Orientation
EOP	Exterior Orientation Parameters
FFT	Fast Fourier Transform
FT	Fourier Transform
FMT	Fourier Mellin Transform
GCP	Ground Control Point
GNSS	Global Navigation Satellite System
GPS	Global Positioning System
GSD	Ground Sampling Distance
HF	High-pass Filter
IFT	Inverse Fourier Transformation
IFFT	Inverse Fast Fourier Transformation
IL	Incremental Learning

IMU	Inertial Measurement Unit
INS	Inertial Navigation System
IO	Inner Orientation
ISO	Integrated Sensor Orientation
LAFMT	Log-polar Analytical Fourier Mellin Transformation
LCP	LiDAR Control Point
LiDAR	Light Detection And Ranging
LSA	Least Squares Adjustment
MI	Mutual Information
MT	Mellin Transformation
MRI	Magnetic Resonance Imaging
NFFT	Non-equispaced Fast Fourier Transform
PAFMT	Polar Analytical Fourier Mellin Transform
Pixel	Picture element
PCL	Point Cloud
POS	Positioning and Orientation System
POMF	Phase Only Matched Filter
PPFMT	Pseudo-Polar Fourier Mellin Transform
QAFMT	Quasi-log-polar Analytical Fourier Mellin Transformation
QAFMT_TC	Tightly-Coupled Quasi-log-polar Analytical Fourier Mellin Transformation
QAFMT_LC	Loosely-Coupled Quasi-log-polar Analytical Fourier Mellin Transformation
RANSAC	RANdom SAMpling Consensus
RMSE	Root Mean Squared Error
SID	Salient Image Disc
SIFT	Scale Invariant Feature Transform
Voxel	Volumetric pixel
2D	Two Dimensional
2.5D	Two- and /a half Dimensional
3D	Three Dimensional

1.0 Introduction

1.1 Motivation

Accurate co-registration of optical images and LiDAR data is important in many remote sensing applications such as (Al-Ruzouq, 2004; Huang *et al.*, 2021; Mishra and Zhang, 2012): Building extraction, image classification, 3D City modelling, canopy modelling, and object reconstruction in street side-, indoor mapping-, and UAV applications. Registration of imagery and surface scans is also important in automotive vision applications.

Over the past decade, new sensors emerged in the geo-spatial market that made data acquisition more user friendly and cheaper than ever before. The market of high-end airborne and street-side sensors has been augmented by a segment of low-cost UAV-, backpack-, and cell phone devices. At the same time, the market segment of fast response applications and change detection, e.g., impact of flooding, hurricanes, and forest fires has become stronger than ever before (Armenakis and Nirupama, 2014; Dilley, 2005). There is also a growing tendency to overlay and compare data from different sources and epochs for map revision (Armenakis, 2000) and to monitor and interpret effects of climate change and natural disasters (Bszukalski, 2017; Ip *et al.*, 2004). These are just a few examples as to why there is a strong need for methods that allow for fast, robust, and fully automated coarse co-registration, so that un-trained operators can quickly conduct overlap checks and overlay and evaluate large amounts of data from different sources.

1.2 Problem Definition

Airborne imagery offers high-density and visually clear information content whilst LiDAR data provides irregular, accurate 3-D geometric surface information which helps the interpretation and classification of features on the ground. Combining LiDAR and camera data is desirable since the combined information allows for better quality surface information than each data source can provide by itself. An example is the color coding of LiDAR data which makes classification for the human operator easier. Another example is the high-density surface information in

homogenous areas and low-density surface information elsewhere provided by LiDAR whilst detailed information is provided along object boundaries everywhere by optical imagery. The prerequisite of using both data sources together is the precise alignment in a common reference frame.

The process of geometrically aligning two or more datasets from different sensors with respect to each other is referred to as co-registration. Often, the term “co-registration” or “cross-source registration” (Huang *et al.*, 2021) is used to emphasize that datasets from different sources are superimposed whilst the term “registration” is used to stress that data from the same type of sensor is spatially aligned. In most cases, it is desirable to align the datasets with high precision. This process is referred to as fine co-registration. To work successfully, most fine co-registration algorithms require an initial estimate of relative pose, referred to as global- or coarse co-registration.

According to Al-Durgham (2014) and Al-Ruzouq (2004), the process of co-registration needs to address four matters:

- (1) Selection of the registration elements
- (2) The matching strategy
- (3) The transformation parameters for relating the reference frames
- (4) The similarity measure describing the coincidence of conjugate elements after registration

RE (1): Typical registration elements are points, lines, planes, higher order primitives, certain objects of interest within the scene, point clusters, or a mixture of these. RE (2): In terrestrial applications, typically a 7-parameter Helmert conformal transformation (three rotations, three translations, and one scale) or a rigid body transformation (three rotations and three translations) are used to relate the coordinate systems. When registering elevation imagery and 2.5D digital surface models, 4 parameters (two translations, one rotation, and one scale) are a popular choice for the transformation model. RE (3): The similarity measure is the mathematical constraint that describes the coincidence of conjugate elements after registration. Typical measures are the spatial discrepancy between tie elements, the differential volume between registered surfaces, the magnitude of a correlation peak, or the mean of the square sum of the residuals of all tie elements.

RE (2): The matching strategy represents the framework (Al-Durgham, 2014) for the automatic co-registration process.

The large volume of research activities in the field of co-registration demonstrates how relevant this topic is. The challenges of creating a matching strategy capable of establishing correspondence between datasets with a high reliability and without the input of an operator, yet also without inferring restrictions or pre-conditions for the datasets, can be summarized in two categories:

- The scenery is acquired at different times, from different viewpoints, and/or with different sensor equipment. Thus, the data sets can differ in alignment/overlap and/or scale but come with different geometric and/or radiometric resolution and/or with different levels of residuals.
- Variances in the data may also be caused by occlusion, moving objects, or partial coverage – some scenery may be visible in one dataset yet not in the other because the view was obstructed, or because one sensor did not collect data over the exact same areas, and the areas of interest may have changed due to urban development or deformations over time.




Figure 1.1: Challenges of co-registration methods. Drawing after Huang *et al.* (2021)

Lastly, overcoming these limitations will increase the computation cost, however a matching algorithm must be efficient numerically (Kim and De Araújo, 2007; Huang *et al.*, 2021).

Although many methodologies have been presented that deal with the problem of automated coarse and fine registration, “automatic and accurate registration of airborne LiDAR and camera data remains an open problem. No techniques have been found that can register optical images and LiDAR data automatically with the required accuracy and efficiency for practical applications” (Mishra and Zhang, 2012; Zitova and Flusser, 2003; Wodzinski and Skalski, 2021). While the class of ICP algorithms can be considered a good solution for automatic fine registration in many

applications, there are situations – as for example cross-aided sensor calibration – where ICP is not the optimal tool (Akca, 2007; Gruen and Akca, 2005; Makadia *et al.*, 2006). In case of coarse registration, there is still no efficient method available that can automatically coarse-align datasets or deliver a quick overlap check for large amounts of airborne survey data: Current algorithms require either input from an operator, application of metadata, non-generic assumptions, or they involve computationally intense processes. Furthermore, most existing registration algorithms consider datasets of similar point density only (Huang *et al.*, 2021; Maiseli *et al.*, 2017; Zhu *et al.*, 2019). These circumstances call for more research of co-registration of airborne survey data with the focus on full automation.

Addressing these challenges, this research will propose, systematically research, and develop a fully automatic and numerically efficient registration paradigm. The topic of this research lunged out from discussions with Dr. Peter Friess, chief of the LMS LiDAR software team at Teledyne Optech (Friess, 2006; Lindenthal *et al.*, 2011; Lindenthal *et al.*, 2015). While looking at a plot of 3D tie points, Dr. Friess realized that they could potentially be used for co-registration, and the idea for the topic of this dissertation was born.

While it is common knowledge that photogrammetric software generates 2D tie points to stitch a collection of images together to one coherent image block, the 2D tie points are also projected from image space on the ground, as illustrated conceptually in Figure 1.2, and the respective 3D coordinates are calculated automatically during the block adjustment (Grafarend and Krumm, 2001). These 3D points – often many thousands of points – are mostly sub-pixel accurate and are traditionally only used as temporary figure and have historically not been used for anything other than blunder detection and generation of a coarse surface model.




Figure 1.2: Spatial forward intersection of 2D tie points into object space

If 3D photogrammetric tie points could be used for automatic co-registration of camera and LiDAR data, it would imply a drastic reduction in computational efforts that are typically associated with existing co-registration methods. For example, existing methods either convert the camera data into a dense point cloud, so that they can match two point clouds with similar characteristics, or they convert the LiDAR data into intensity- or elevation imagery, so that they can conduct image matching. Contrary, 3D photogrammetric tie points are created automatically per default.

For these two point-data sets, however, we cannot define the correspondence of anchor points for performing the co-registration transformation. If we use each point cloud separately most likely we end up with two different surface representations of the same topographic surface: One surface representation, the one derived from lidar based on the high density and quality approximates the topographic surface very well. The other one, the one based on the photogrammetric tie points due to their sparsity is a more generalised (down sampled) representation of the same topographic surface. Further the photogrammetric tie points give a digital surface representation while the lidar points give a mix of surface and terrain (ground) surface representation. If for example, we had bold earth the problem would be simpler as both type of points would lie on the same surface (both surfaces would have the same trend).

1.3 Thesis Objectives

The objective of this study is to analyze how photogrammetric tie points can be used for automated co-registration of airborne imagery and LiDAR data from traditional large-area mapping. The initial phase of this research proposes new matching strategies in space domain, first for coarse registration of optical and laser point clouds, and then for fine registration and respective cross-aided sensor calibration using photogrammetric tie points and LiDAR extracted planes as tie primitives. The second phase of this thesis analyzes what frequency domain-based algorithms can be used for robust automatic coregistration. While exploring these avenues, this study provides a detailed description of the mathematical components of the proposed photo-LiDAR registration methodologies, including generation and selection of the tie elements, transformation function, similarity measure, and matching strategy.

1.4 Research Contributions

The main contribution of this dissertation is the proposal of a new co-registration paradigm where photogrammetric tie points are used as registration features with Lidar data, and the systematic research how these elements can be used best for co-registration. Through analyzing existing methodologies and through deploying new algorithms, this research generates a better understanding of how the different characteristics of photogrammetric tie points and airborne LiDAR data come to play in the context of co-registration.

Summarized, this dissertation provides the following major contributions: (A) Proposal and systematic research of a new and powerful concept for co-registration of airborne LiDAR and camera data – to the best knowledge of the author this is the first time ever someone has considered extremely sparse tie points as tie features for co-registration, (B) Systematic research of space-domain based concepts typically not used for coarse registration, (C) Proposal and study of four novel space-domain based fine registration and integration algorithms, (D) Isolation of fundamental criteria which limit the use of tie points for fine registration and cross-aided sensor calibration, providing a new and deeper understanding of tie points in the context of coregistration, (E) Application of the Phase Only Matched Filter (POMF) and Fourier Mellin Transform (FMT) to unclassified airborne point clouds, (F) Detailed description of the Analytical Fourier Mellin Transform (AFMT) and its application for co-registration of airborne LiDAR and camera data, and (G) introduction of the new Quasi Log-Polar Analytical Fourier Mellin Transform (QAFMT) as optimal method to coarsely co-register large area airborne photogrammetric tie point clouds and LiDAR data.

1.5 Thesis Outline

This dissertation consists of six chapters, which are briefly explained below:

Chapter Two reviews previous works relevant to co-registration of optical and laser data, primarily for the case of traditional large-format airborne applications but also for a few selected close-range

terrestrial- and UAV cases. The methodologies are grouped primarily by the components of their registration model.

Chapter Three describes the datasets, how they compare against data from the OEEPE test site, and what mathematical models were chosen to support coarse- and fine registration.

Chapter Four presents seven space domain-based methodologies that did not surface during literature review but could be used for coarse co-registration and overlap check. It also proposed novel methods for fine registration and cross-aided sensor calibration in space domain, using photogrammetric tie points as tie elements.

Chapter Five presents seven methods of coarse registration in frequency domain and yields the QAMFT as a novel method for co-registration of extremely sparse tie points with large format airborne laser data. It also discusses why fine registration cannot occur in frequency domain.

Chapter Six provides the conclusions and recommendations for future work.

The following appendices include tables, graphs, and more information, supporting the claims made in the main body: *Appendix A* reviews the existing AI-related co-registration methods and their relation and contrast to the proposed registration primitives; *Appendix B* covers the use, generation, and characteristics of photogrammetric tie points and airborne Lidar data, including their respective advantages and disadvantages; *Appendix C* provides a detailed description of the test data; *Appendix D* presents the results of the fine registration using unclassified tie points and unclassified laser points using ICP; *Appendix E* describes the design of experiments, reduction of design space, and refined experiments with simulated data; *Appendix F* describes the interpretation of FMT experiments with real-world data; *Appendix G* describes the precise alignment of point clouds to establish the reference data; and *Appendix H* explains the change of data footprint and entropy when estimating data shifts using POMF.

This dissertation research has the following publications:

Granted:

Lindenthal, S.M., Wang, J.G., Bisnath, S., and Canter, P., 2009. Performance analysis of GNSS network extrapolation supporting the moving reference station concept. *York University internal report*, generated in cooperation with Trimble Applanix, 2009

Lindenthal, S., Ussyshkin, V., Wang, J. and Pokorny, M., 2011. Airborne LIDAR: A fully automated self-calibration procedure. *ISPRS Achieves, ISPRS Calgary 2011 Workshop*, pp.29-31.

Lindenthal, S.M., Armenakis, C., Wang, J.G., Bisnath, S. and Goldman, H., 2015. Sensor-Driven Calibration and Co-Registration of Airborne Multisensor LiDAR Data. *Canadian Journal of Remote Sensing*, 41(3), pp.145-158.

Prospected:

Covering the Co-Registration of Airborne Laser Scanner and Optical Data Using Image Tie Points And The Quasi-Log-Polar Analytical Fourier Mellin Transform.

2.0 Literature Review

2.1 Introduction

This chapter reviews the literature on the co-registration of laser and camera data. First, the elements of a registration paradigm will be described. Then, different co-registration methods are reviewed. Finally, an analysis of a few studies that showed similarity to the proposed co-registration approach will be provided at the end of this chapter.

2.2 Registration Paradigm and Categorization of Registration Algorithms

The efficient registration of geospatial data has been an active field of research over the past few decades and many methodologies have been presented. These methods can be categorized based on what tie elements they use, what optimization strategy they implement, the domain in which they work – as for example computer vision and pattern recognition, medical image analysis, or remotely sensed data processing (Brown, 1992) –, the type of data which they register – imagery and/or point clouds –, the geometric transformation function which they implement, whether they consider statistical data or meta data or only geometric elements during the registration process, or whether they align datasets coarsely or with high accuracy.

How to categorize registration approaches also depends on the scientific domain, the type of applications, and the academic surroundings in which the authors dwell. Prof. Ayman Habib and his team for example showed a strong focus on the registration and segmentation of static and mobile LiDAR and camera data over the past decade (Al-Durgham *et al.*, 2011; Al-Durgham, 2014; Al-Ruzouq, 2004; Bang, 2010; Canaz, 2012; Habib *et al.*, 2006; Habib *et al.*, 2007a; Habib *et al.*, 2010a; Habib, 2013; Kwak, 2013; Kwak and Habib, 2014, Renaudin *et al.*, 2011). Their work postulated the following four-problem-oriented (co-)registration paradigm – or “steps” as referred to by Fonseca and Manjunath (1996) and Wyawahare *et al.* (2009), or “choices for

components” as referred to by Brown (1992), or “matching strategy” (Al-Ruzouq, 2004; Canaz, 2012; Al-Durgham, 2014) – which the process of co-registration must address:

- (1) *Selection of co-registration primitives*: After data acquisition and data preprocessing, it is necessary to decide which types of primitives will be used for the alignment. This can be identified manually or automatically with high certainty. With airborne data, common tie features are salient points or distinctive geometric features such as photogrammetric tie points, lines, line intersection points, planar patches, homogenous regions, curved edges, contours, object boundaries, etc. which are manually or automatically detected from optical images and LiDAR data.
- (2) *Feature matching* (Zitova and Flusser, 2003; Mishra and Zhang, 2012), also referred as *matching strategy* (Habib *et al.*, 2006): Establishing the correspondence between conjugate tie features, as for example detected in optical imagery and laser point clouds (Al-Ruzouq, 2004). Various feature descriptors and similarity measures along with spatial relationships are used for matching the features (Mishra and Zhang, 2012).
- (3) *A transformation function*: Once the correspondence between the matched features has been established, the transformation parameters can be computed. The mathematical transformation model must accurately reflect the geometric relationship of reference and moving data. Examples of frequently used mapping functions are the 7-parameter Helmert transformation (3 rotations, 3 translations, 1 scale) for overlay of 3D point clouds and the 2D rigid body transformation (1 rotation, 2 translations, 1 scale) for registration of airborne imagery or elevation maps.
- (4) *A similarity measure*: After a decision is made on the type of registration primitives and the transformation function, a mathematical constraint, which is necessary to ensure the correspondence of the conjugate primitives (Al-Ruzouq, 2004) must be selected. The formulation of the similarity merit depends on the type of registration elements and the mathematical representation of the respective optimization function (Canaz, 2012;

Renaudin *et al.*, 2011). In some cases, this component could be expressed through a certain expectancy level of the magnitude of correlation peaks.

There are two more components that may appear self-evident but are worthwhile defining:

- *Transforming and resampling*: This component is mentioned as an explicit step only by Fonseca and Manjunath (1996) and Zitova and Flusser (2003). These authors simply assumed that source data were transformed onto the target data after the transformation parameters had been determined. For example, to co-register LiDAR and camera data, the imagery is projected onto the LiDAR elevation map by means of the mapping function, and positions in non-integer coordinates are computed by means of interpolation.
- *The search space*: This component was highlighted by (Brown, 1992): A decision on the spatial extent of the comparative data analysis must be made. The assumptions of the underlying transformation model determine whether all data are considered (global registration approach) or if the search for correspondence should occur only within a local proximity (fine registration approach).

The next chapter will discuss how these six components have been addressed by the current co-registration techniques.

2.3 Registration Primitives and Matching Strategies

The matching strategy comprises the overall scheme of the solution of the matching problem (Al-Durgham, 2014). It comprises the selected primitives, transformation functions and similarity measures for automatically solving the registration problem (Al-Ruzoq, 2004; Schenk, 1999). Matching strategies for coarse co-registration do not require approximate values for the transformation parameters and fully automated approaches can derive an estimate of initial pose without any geometric pre-conditions. However, fine co-registration approaches require an initial

estimate of the transformation parameters to successfully superimpose datasets with high precision.

2.3.1 Existing Co-Registration Methods

As mentioned in the previous chapter, there are many attributes that can be used to classify registration algorithms. The classification schemes can be grouped as follows (Al-Durgham, M., 2014; Fonseca and Manjunath, 1996; and (Mishra and Zhang, 2012): (1) direct geo-referencing (DG), (2) the target-based, (3) the point-based, (4) the feature-based – which can be categorized further into methods that use (4a) lines, planes, and intersection points, and methods that use (4b) freeform primitives – , (5) the mutual information (MI) -based, (6) salient point -based, (7) the frequency-based, and the Artificial Intelligence (AI) -based methods. The last category of co-registering overlapping multi-sensor data is described in detail in APPENDIX A; the following sections summarize their limitations:

1) Direct Georeferencing: DG is a well-established method to coarsely position the raw camera and LiDAR data, as it infers geodetic projection to the data. However, as outlined in the thesis statement, coarse co-registration should be able to register datasets without knowledge about the map projection and units in which the data have been referenced. Furthermore, most UAV and low- cost airborne mobile mapping systems operate with a lower grade navigation device. These factors create the need for algorithms that can achieve coarse level accuracies without specific DG.

2) Target-based methods: They rely on measuring targets in the survey site. Whilst ground targets provide a means to achieve high registration accuracy, they may not always be available in every data set. Since identifying and measuring ground targets typically requires a human operator, target-based approaches do not lend themselves for automatic co-registration.

3) Point-based methods: They have been proven to be accurate in many applications. However, information about the initial position of the datasets must be available. Furthermore, for the fine registration and optional cross-aided sensor calibration to occur with high precision, point-to-point

and point-to-plane algorithms require point clouds with strong elevation contrast and similar point density, or some sort of classification and pre-processing, which goes against the thesis statement.

4) Feature-based methods: In this dissertation, co-registration algorithms that use lines, planes, line intersection points, 3D shapes, and free forms as tie primitives are considered “feature based” algorithms. The feature-based methods that use lines, planes, and line intersection points rely on the existence of human made structures in the scenery. They also require ample depth in object space and similar point densities. In the proposed airborne application, these conditions are not met. The feature-based methods that involve that use 3D shapes, and free forms are suited to function with both natural scenes and built areas, however they require feature extraction and/or classification with visual interpretation of an operator and do therefore not fall within the category of fully automated, robust registration approaches.

5) Mutual Information-based methods: MI methodologies are area-based and tend to perform better with low resolution geodetic elevation data than feature-based methods, since they consider not only one type of geometric element during the co-registration process, but statistical information of all common surface data. Unfortunately, when the images have significant rotation and/or scale differences, MI methods either fail or become extremely time expensive. Also, they are ineffective for non-urban scenery since information-theoretic methods require well-pronounced structures, and some sort of prior knowledge about the datasets and their relation and cannot be applied to natural scenery.

6) Salient point-based methods: Salient point algorithms that use a combination of multiscale image analysis can be helpful to improve overall registration accuracy, especially in suburban areas and natural scenery. Here, salient point matching will fail since it cannot find the feature footprints of human made structures implemented in its feature map data base. However, “salient point-based methods do not address the problem of relief displacement”; also, “its complexity and high computational requirements make it problematic for real-world applications” (Mishra and Zhang, 2012).

7) Frequency analysis-based methods: Frequency methods promise to be the most effective approach for coarse registration since they consider all information of the data input at once. Existing frequency analysis methods proved unable to produce accurate registration results. Thus, in this dissertation, co-registration algorithms will be applied to elevation maps only, and not to intensity images. As will be shown in Chapter 5.0, the FMT and AFMT and thereof derived frequency algorithms will be used to correlate the elevation maps. To date, the performance of FMT and AFMT with low-resolution airborne point clouds is still unknown since the numerical examples in existing publications used image- and laser data of same resolution.

8) Artificial Intelligence (AI)-based methods: APPENDIX A provides an overview of the most recent AI-based methods for registration and co-registration of point clouds and shows how these methods relate to the use of traditional airborne laser- and photogrammetric tie points as tie features.

For the point clouds used in the optimization-based and deep-learning-based registration methods, it is typical to have a lot of information about the depth, and to sample the same surface elements than their counter part. Also, they differ in point density only by approx.1:100. So, often, the research community refers to these point clouds with the lower point density as “sparse point clouds”. Contrarily, tie point clouds from traditional large- and medium-format airborne applications result point densities approx. 5,000-10,000 times lower than their counterpart.

As deep learning is a technique to implement machine learning and machine learning is an approach to achieve artificial intelligence, the “sparse” point clouds used in the AI community allow for feature extraction and learning of local point cloud characteristics. From the experiments given in Chapter 5, the “extremely sparse” tie point clouds targeted in this framework pick up different surface characteristics and describe the surface with insufficient detail to extract features that would be compatible to those extracted from the respective Lidar data in resolution and information content. Consequently, the current AI methods cannot be used with the proposed intake of point clouds. Likewise, the application of AI methods to imagery, for example, the application of Convolutional Neural Networks (CNN) to conducting image- or patch matching, will not work with the elevation maps targeted in this dissertation, either.

An alternative way to categorize the existing co-registration methods is to group them into information theoretic-based methods and probability-based methods (Zhu *et al.*, 2019).

2.3.2 Co-Registration Methods of Special Interest

From literature review, five publications surfaced which followed similar strategies as proposed in this dissertation. This section pays special attention to these five co-registration methodologies.

- (1) As mentioned in Chapter 1.0, the research topic of this dissertation surfaced in 2010. In 2013, photogrammetric tie points and laser points were used as registration primitives in a close-range application (Zheng *et al.*, 2013), in which data from a static terrestrial laser scanner and a camera at different exposure locations were used to map a stone statue (Figure 2.1).



Figure 2.1: Dense tie point cloud from an image pair of a stone statue. Source: Zheng *et al.* (2013)

This multi-modal arrangement appeared similar at first sight; however, the closer examination revealed crucial differences:

- The sensors sampled the surface approximately from the same location with the similar field of view. This a-priori knowledge about the location, scenery, coordinate system, and units made the coarse registration unnecessary
- The point densities of the laser and tie point clouds were similar, i.e., their point ratio was approx. 1:1-to-1:100, yet specifically approx. 1:10,000 in this dissertation
- The scanned object had a convex surface with large changes in depth resulting in point clouds without occlusions
- The statue provided homogenous surface characteristics. Thus, there were only single returns in the laser data
- The tie- and laser point clouds were of equal accuracy
- Their co-registration algorithm was based on a least squares adjustment in space domain subject to minimizing the sum of the squared distances between the tie points and the LiDAR surface. The statue filled most of the image space, which allowed to additionally estimate the effects of lens distortion using the laser points as the ground truth
- The ratio of the surface elevation of the scanned statue and the distance between the sensors and the statue was much bigger than in airborne applications, allowing for an exceptional depth resolution. The result was a truly 3D point cloud with ample depth and strong changes in direction of surface normals rather than a fairly flat 2.5D surface point cloud with small changes in elevation

(2) Another similar co-registration strategy was proposed by Choi *et al.* (2013) which suggested using 3D photogrammetric points and automatically LiDAR-extracted planar patches as its registration primitives. However, the roof peak points in the images were manually measured to, and the roof peak lines were derived from these points. The roof peak lines were also extracted from the laser data by intersecting the roof patches, which were extracted automatically. Corresponding roof peak line pairs were used for horizontal registration. In addition, the tie points were measured manually on planar surface areas and the vertical point-to-plane discrepancy was minimized to solve for the vertical registration ambiguity. Moreover, the horizontal and vertical discrepancies were minimized simultaneously in a LSA. Figurative examples of the registration steps are presented below.




Figure 2.2: Extraction of roof peak lines. Source: Choi *et al.* (2013)

In Figure 2.2, (a) shows the roof peak points derived from the stereo photogrammetry measurements, (b) displays the overlay of the LiDAR points and the image derived points, and (c) displays the roof planes and peak lines before and after the adjustment in green and red, respectively.

Although this registration methodology uses similar registration primitives as the proposed fine registration methods – 3D photogrammetric points and LiDAR planes – it differs from the fine registration approach presented in section 4.4.2 in the following two aspects:

- The image points used for registration were measured manually using stereo photogrammetry

- Rather than utilizing the LiDAR patches themselves, the tie lines were extracted from the LiDAR data

(3) Another closely related work was provided by Makadia *et al.* (2006), which used the spherical harmonics and the Extended Gaussian Histograms (EGI) for the coarse registration of 3D point clouds of close-range laser scans as shown in Figure 2.3. This approach utilized spherical harmonics to fit into the data and to compute the surface normals to derive the surface orientations. Instead of correlating the images via the surface curvatures, the surface orientations were compounded in the spherical histograms to determine the 3D misalignment parameters. It has proved of being more robust with it than with the ICP for close range point clouds with little overlap and coarse initial orientation parameters.




Figure 2.3: Registration of a “Happy Buddha” statue. Source: Makadia *et al.* (2006)

In Figure 2.3, (a) shows the scans in their initial positions and orientations, input to Makadia’s algorithm, color-coded by exposure, (b) shows the crude alignment of ten point-sets from the front, (c) shows the final alignment of all scans from the side after ICP was run, (d) shows a pair of re-projected EGIs from two of the scans, and (e) shows a slice of the correlation grid $G(R)$ at the location of the estimated rotation.

The surfaces of most airborne data are remote from the sensing device and their elevation changes are comparatively small. Thus, the respective orientation histograms are broad and small in magnitude, making successful cross-correlation difficult. Nevertheless, combining the use of spherical harmonics and histograms of surface orientations is an

interesting idea that could perhaps be used to characterize datasets to establish a metric that allows to distinguish between the mountainous and flat datasets.

- (4) An outstanding work in the realm of co-registration is the dissertation of Dr. Devrim Akca (2007). Gruen (1996) extended the least squares 2D matching technique of images for 2.5D and 3D surfaces. Akca achieved a co-registration accuracy superior to those gained with ICP. Like ICP, his “Least Squares 3D” (LS3D) surface matching algorithm evaluated Euclidean distances rather than the height differences, which would be sub-optimal for terrain surface applications. In his dissertation, Akca applied this co-registration method to the close-range 3D point clouds as well as the airborne survey data.

As can be seen from Figure 2.4, the LS3D methodology requires only small overlap for successful co-registration. He also introduced the least squares matching of free-form curves and surface models. Recognizing that some point clouds may not contain sufficient surface information, he integrated the intensity and the color information of the object surface. Like Friess (2006) and Lindenthal *et al.* (2011), Akca generated a color-coded map of the height differences before and after LSA for quick and intuitive quality analysis of the co-registration results. An example of such a height difference plot is given in Figure 2.4.




Figure 2.4: Registered surface models and respective height difference map of the overlap area. Source: Akca (2007)

Akca (2007) also used dense photogrammetric DSMs and LiDAR point clouds for co-registration. Photogrammetric dense point clouds tend to show high noise on flat surface areas compared to the LiDAR point cloud. However, LiDAR point clouds have irregular and often lower point density than the respective photogrammetric dense point clouds. To minimize the impact of these opposing characteristics, he applied a median filter for smoothing the LiDAR data prior to matching.

Despite these great contributions, Akca's work (2007) showed the following significant limitations:

- The LS3D method relies on the photogrammetric surface models, which require computationally expensive dense point matching (Haala, 2013) and are large in data volume. For many airborne photogrammetry applications, generating orthophotos as end-product is sufficient and dense point matching is not part of the data processing workflow. Thus, a co-registration requiring a photogrammetric DSM runs at a high cost, unless the DSM is already available on file.
- The LS3D approach may work with orthophotos. However, depending on the application, the deliverable may just be a set of the color-balanced photos plus the refined EO parameters from aerial triangulation. The computationally intense orthophoto generation – which typically requires another software module in addition to the aerial triangulation module – is not always part of the deliverable.
- Matching intensities, color information, and surface geometries in a combined LSA may work with the dense point surface models but also with orthophotos. However, such an estimation process wouldn't work with the photogrammetric tie point clouds since the traditional photogrammetric software does not report the average color information of the generated tie points.
- LS3D could be extended and used for 3D sparse points, such as ground control points, in combination with a 3D surface. Unfortunately, ground control points may not always be available. Moreover, this extension might also work with photogrammetric tie points. However, the concept of using LS3D for sparse points has only been presented as “further conceptual extension” (Akca, 2007), has not been implemented and tested.

- With the concept of integrating additional functional constraints in his LSA, such constraints typically require some sort of pre-knowledge that is unique to the dataset. Although automatic classification of laser point clouds and dense point clouds into ground- and air points is possible (Soininen and Burman, 2005; Soininen, 2010a; Haala, 2013) and might be a way to introduce such constraints automatically, the tie point clouds are too sparse and too random in distribution for these algorithms to work reliably. Refer to the next section and Chapter 4.0 for more details.
- The stochastic properties of the local surface normals and uncertainty of the points due to the incidence angle were neglected in Akca's work (Grant *et al.*, 2012).
- Akca's method did not address the problem of relief displacement, i.e., the fact that LiDAR data is orthometric per default, but optical imagery – unless converted into an orthophoto – shows effects of perspective geometry (Mishra and Zhang, 2012).
- Although a global matching strategy has been mentioned in (Gruen and Akca, 2005), it was unclear from their publications how accurate the initial approximations need to be for the LS3D adjustment to converge to the correct maximum. As Akca compared his LS3D results primarily against ICP results, the LS3D is seemingly a method intended primarily for fine registration.

In summary, even though Akca's work introduced a powerful new co-registration algorithm, it is unclear how well this method would work when used for coarse co-registration with extremely sparse, unclassified photogrammetric tie points without color information.

- (5) Two years later after we started this research, Gneeniss *et al.* (2014) presented a methodology for fine registration and cross-aided sensor calibration of airborne LiDAR and camera data using the LiDAR point clouds as the ground truth. Their research methodology showed parallels to the fine-registration and cross-aided sensor calibration method (summarized in Chapter 4.4.2), which I proposed in 2012. Their approach used 3D photogrammetric tie points for co-registration and cross-aided sensor calibration, and also considered the LiDAR as reference against which camera calibration and validation can be

performed. They also used a 3D surface matching method utilizing the tie-point-to-laser-plane matches. A high-level overview of Gneeniss' co-registration methodology is illustrated in Figure 2.5.




Figure 2.5: Main logical components of the co-registration and cross-aided sensor calibration method suggested by Gneeniss *et al.* (2015)

The methodology of Gneeniss *et al.* (2015) can be broken down into 9 steps:

- (1) Automatically generate 2D tie points
- (2) Project the tie points on the ground using ISO
- (3) Filter the 3D tie points into
 - ground model points, describing the bare earth surface,
 - points describing vegetation (trees) and buildings, and
 - and apply a thematic image mask to isolate points that fall in between ground and vegetation *and* satisfy a certain vertical threshold
- (4) Automatically extract key model ground model points (DTM) from LiDAR.

(Component (a) in Figure 2.5 above covers steps 1-4)

- (5) Co-register ground model tie points and LiDAR key point DTM surface using a refined version of the LS3D surface matching method (Akca, 2007; Miller *et al.*, 2008) and a 3D conformal transformation model, i.e., a Helmert transformation (*component (b) in Figure 2.5 above*)
- (6) Assess the quality of the registration results by evaluating the elevation difference between the photogrammetric DTM and the LiDAR terrain model
- (7) If the quality test indicates that the surface matching can be improved further: Identify 3D image tie points that meet the following criteria:
 - be located over planar areas that
 - show little slope
 - be defined by more than 3 image rays
 - be in areas with low surface roughness
 - pass an accuracy- and blunder test, and
 - support homogenous spatial distribution across the survey areaand use them as “LiDAR Control Point” (LCP), (block (d) in Figure 2.5).
- (8) Estimate camera calibration parameters by minimizing the 3D distance between the LCPs and the local LiDAR surface (block (d) in Figure 2.5). As can be seen from Figure 2.6, this will lead to an improved fit of the camera and LiDAR data. The process of co-registration and self-calibration can also be conducted simultaneously in one combined LSA that estimates the camera calibration parameters while synchronously the 3D distance is being minimized between the tie points and respective local LiDAR surface.
- (9) Repeat step 6 to 9 and re-assess the vertical difference of the refined alignment between the photogrammetric DTM and the LiDAR terrain model.

To validate this method, the authors used two datasets similar to the test datasets used in this dissertation. One of their datasets was captured over the OEEPE test site in Vaihingen/Enz and an Optech ALTM 2050 LiDAR and a Applanix DSS camera were used for data collection (refer to Section 3.2 and APPENDIX C for details). Figure 2.6 below displays some of Gneeniss’ research outcome.




Figure 2.6: Surface matching residuals of Gneeniss' British data set using 50 LCPs.

Source: Gneeniss *et al.* (2015)

Rows (a) and (b) in Figure 2.6 present the results from the DSS camera flight before and after the self-calibration, respectively. Even after the self-calibration of the DSS data, the co-registered datasets showed a remaining vertical offset of about 20cm which the authors attributed to the differences in surface characteristics and weak block layout leading to limited ability to retrieve the camera calibration parameters reliably.

Gneeniss' procedure showed commonalities with the co-registration and calibration methodology presented in Chapter 4.4.2:

- Both assume the LiDAR to have undergone geometric refinement and provided superior geometric quality.
- Both use the photogrammetric tie points as the tie elements in the co-registration and minimize the sum of the squared distances from the tie points to the LiDAR surface in proximity.
- Both allow for separating the co-registration and the camera calibration or for the simultaneous co-registration and camera calibration in a combined bundle block adjustment in their workflows.
- The data from the DSS medium format camera benefited significantly from their calibration – as can be seen from Figure 2.6, as the co-registration alone did not provide the desired level of decimetric or centimetric registration accuracy. The DiMAC camera data used in this dissertation was affected similarly by the boresights and camera intrinsic offsets, which will be shown in APPENDIX C and G. It is suggested to minimize such effect during the cross-aided sensor calibration.
- Even after the calibration of the DSS camera, the co-registration showed a remaining vertical offset of about 20cm. The authors attributed this relatively large remaining offset to the poor layout of the flight block, which may or may not be true. The magnitude of this offset could be taken to serve as an indicator of what accuracies to expect when co-registering optical and laser point clouds in this dissertation.

Despite these similarities, Gneessniss's work showed some significant differences and limitations compared to our methods:

- Gneessniss' algorithms rely on the dense tie point clouds as their datasets possessed a point density 10-100 times higher than the tie point clouds used in this dissertation.
- Gneessniss developed a sophisticated filter for the tie- and laser points to increase the reliability and accuracy of the registration results. As the classification into ground-, vegetation-, and roof points was executed automatically, the authors also deployed a 2.5m vertical mask to eliminate the chimneys, trees, and unwanted effects of a vineyard. The application of this mask allowed to distill a selection of

suitable points that yielded the robust registration results. However, this filter was data-specific and tuned for the Vaihingen/Enz dataset, which went against the principle of generality and full automation proposed in this dissertation.

- Gneeniss’ concept did not utilize the roof plane patches as registration primitive, which could provide a higher level of redundancy and accuracy than the local surface meshes.
- Since the authors used LS3D for co-registration, their method was affected by the drawbacks mentioned earlier, too.

As the above discussion showed, the co-registration methods presented in these five publications shared similarities with the proposed co-registration methodology. However, the analysis also revealed that these methods had need of either substantial data pre-processing or pre-knowledge about the datasets. Moreover, the conducted tests were only with point clouds associated with similar surface characteristics, which is their most prominent shortcoming.

2.4 Summary

This chapter reviewed the relevant state-of-the-art co-registration techniques in the fusion of airborne laser- and camera data. Despite the large diversity of ideas and concepts, many of the published co-registration techniques required the coarse EOPs and IOPs which made them not practical (Li *et al.*, 2012). Some depended on selecting conjugate pairs interactively, which may lead to unreliable and inaccurate results and worked against the principle of full automation. Most methods were designed for the computer vision applications and assumed to have the high point density, strong variation in elevation, similar footprint, large overlap, same coordinate systems. Besides, they demanded some sort of classification or computationally intense pre-processing. Only one machine learning-based co-registration method has been presented so far could work directly on the raw “sparse” ⁽¹⁾ point clouds. The most important drawback however is that none

¹ The research community concerned with AI-based registration uses the term “sparse point cloud” for point clouds that stem from a close-range scenery and are roughly 100 times less dense than its counterpart. The point cloud with the lower (or higher) density often stems from close-range photogrammetry (a laser scanner).

of the existing methods was designed for and validated with the “extremely-sparse” ⁽²⁾ photogrammetric point clouds.

The ensuing chapters will investigate what can be done outside these boundaries. The focus will be directed towards finding innovative methods that can cope with the stark difference in point density and establishing a novel framework in which they will deliver reliable results.

² Tie point clouds from traditional large- and medium-format airborne mapping applications result a point density of approx. 5,000-10,000 times less than the respective LiDAR data. To discern the tie points targeted in this framework from “sparse” close-range point clouds, they will be referred to as “extremely sparse” whenever their difference needs to be emphasised.

3.0 Test Data and Assumptions

3.1 Introduction

In this framework, two pairs of real-world datasets and two simulated datasets were used for testing. To address the question, to what degree these datasets represent “typical and universal” airborne survey applications, the characteristics of a popular test site are highlighted, for comparison, first.

3.2 Reference Dataset

One of the world’s most flown test sites for evaluating large- and medium format digital airborne sensors is the test field in Vaihingen an der Enz, Baden-Württemberg, Germany. While no dataset of this area was used in this study, the terrain characteristics of this area can serve as a guideline for the simulated dataset and for the comparison against the real-world datasets used in this thesis.

The Vaihingen test site provides the following features:

- High density ground control - almost 200 control points have been spread across the area
- Mostly gently sloped terrain. Few areas with steep slopes
- A mixture of urban scenery and farmland
- Small areas with forest
- Small waterbodies, which typically create a challenge for LiDAR and photogrammetric tie point matching

The test site is enclosed with the red line in Figure 3.1 and Figure 3.2.



Figure 3.1: Aerial view of the OEEPE test site in Vaihingen an der Enz, Germany. The dark green areas depict forest and water bodies. Source: IFP (2017)




Figure 3.2: Digital terrain model of the Vaihingen and der Enz test site. The terrain variations are exaggerated by factor two. Source: IFP (2017)

The central area – enclosed with a purple line in Figure 3.1 – possesses a dense distribution of GCPs and has about 5050m (east-west) by 2750m (north-south) in size. The entire test area – enclosed with a red line in Figure 3.1 – is about 7400m (east-west) by 4700m (north-south) in size. The terrain elevations range between 193m and 356m above sea level. In total, 275 GCPs are available.

3.3 Real-World Datasets

This chapter introduces the real-world datasets that will be used for testing. APPENDIX C presents how these datasets were generated, and a detailed description of their characteristics.

3.3.1 Oshawa Dataset

The Oshawa data was acquired at three flying heights (850m, 600m, 400m above ground level). Only the data collected from the flight lines at an altitude of 850m and 600m agl were selected for co-registration. The flight lines flown with large overlap result in higher laser and tie point densities than the areas with small overlap. The resulting image block comprised 105 photos, 3641 tie points after the outlier removal, and 51 measured control points.

A surface grid was fit into the tie points using MATLAB software (MathWorks, 2019). Figure 3.3 below displays the tie point cloud overlaid with a natural-fit surface grid. For the average point density of 35 tie points per image, MATLAB considered a cell size of 40m by 40m to be optimal. Figure 3.4 displays the respective grid, rendered to an image file.




Figure 3.3: The Oshawa tie point cloud overlaid with the MATLAB-generated surface mesh




Figure 3.4: The Oshawa tie point grid with 81x115 40m cells

The full resolution point cloud had approximately 63 million points and covered an area of approximately 4.5-by-4.5km, yielding an average point density of 4 pts/m². The laser points were down sampled to 8775 points (APPENDIX C), i.e., a point density slightly higher than twice that of the tie point cloud. Figure 3.5 displays the laser point cloud overlaid with a naturally-fit surface grid (MathWorks, 2019). Figure 3.6 displays the respective grid, rendered to an image file.




Figure 3.5: The Oshawa laser point cloud overlaid with the surface mesh




Figure 3.6: The Oshawa laser point grid with 117x138 40m cells

The Oshawa tie and laser points showed a ratio of $3,700/63,000,000 = 1:17,000$, which confirmed the approximate ratio of $1/10,000$ that one would typically obtain from regular survey data. Details on the quality of fit can be found in APPENDIX G.

Figure 3.7 below displays the co-registered point clouds on top of the surface mesh. The laser points covered an area of approx. 4.5km-by-4.5km whereas the tie points covered an area of approx. 3.2km-by-3.2km.




Figure 3.7: Overlay of the Oshawa laser- and tie point cloud and the surface mesh

3.3.2 Keyworth-Nottingham Dataset

The Keyworth-Nottingham (KN) was collected at two flying heights (1250m and 880m agl). The resulting image block comprised 64 photos, 2705 tie points before and 1871 tie points after blunder removal, and 38 measured control points. 38 ground control points had been measured.

Figure 3.8 below displays the tie point cloud overlaid with the “natural-fit” (APPENDIX C) surface grid. For the average point density of 30 tie points per image, MATLAB considered a 40m-by-40m cell size to be optimal. Figure 3.9 plots the matching grid, rendered to an image file.




Figure 3.8: The KN tie point cloud overlaid with the MATLAB-generated surface mesh




Figure 3.9: The KN tie point elevation grid (m) with 36x35 40m cells

The KN laser point cloud had approximately 35 million points covering an area of 2.4km-by-2.4km resulting in a point density of 5pts/m²; the tie point cloud covered an area of 1.5km by 1.5km with 1871 points, resulting in a point density of roughly 0.0008 pts/m². After down-sampling, the laser point cloud contained 2705 regularly distributed points.

Once the LiDAR had been down-sampled, both point clouds were gridded and rendered to images using the algorithm described in APPENDIX C. Figure 3.10 illustrates the down-sampled point clouds overlaid with the grid. Figure 3.11 displays the respective grid, rendered to an image file. The color coding was the same as in Figure 3.3, however the large descent of the grid in the north-east corner caused the center of the grid to occur more yellow. More details can be found in APPENDIX C.




Figure 3.10: The KN laser point cloud overlaid with the MATLAB-generated surface mesh




Figure 3.11: The KN laser point elevation grid (m) with 59x71 40m cells

Figure 3.12 below plots the co-registered point clouds. In this graph, the laser points covered and area of approx. 2.4km by 2.4km and the tie points covered roughly 1.5km by 1.5km ground.




Figure 3.12: Overlay of the KN laser points, 10m elevated tie points, and the surface mesh

3.4 Simulated Dataset

The simulated data set was generated by superimposing five 1D harmonics to a signal that mimicked the elevation magnitudes of the Vaihingen/Enz dataset (see section 3.2 above). To have a similar extend than the KN dataset, a signal length of 2km was chosen. The resulting 1864x1864 “X5d4” DTM grid of 1m pixels is displayed in Figure 3.15 side view and top view below. For the “X5d5” DSM, displayed in the second column of Figure 3.15, A 10m white noise band was added. The X5d4 DTM and X5d5 DSM were used for most of the AFMT testing.




Figure 3.13: The simulated X5d4 DTM (m) and the X5d5 DSM (m) with further modified second peak





Figure 3.14: The padded X5d4 DTM (m) image used for co-registration testing

3.5 Assumptions

When the airborne LiDAR and photo data of a common area have been acquired at different times, and/or from different platforms, and/or at different resolutions, the co-registration may face the following challenges:

- Use of different geodetic datums, map projections, and units, resulting in rotational, translational, and scaling offsets
- Different coverage of the landscape
- Different times of data collection, resulting in capture of incremental changes of the scanned surface: Leaves on or off, de-forestation, new building areas, snow coverage, moving objects, etc.
- Different information-sampling schemes of the surface: LiDAR provides multi-return data whereas the traditional nadir photogrammetry provides only one single Z value associated with the X and Y-coordinates
- Different levels of noises
- Different conventions of the coordinate axes, e.g., flipping of the coordinate axes or changes of direction
- Similarly, different formats of the coordinate input (North values may be the first column in one file, while East values may be the first column in another file)
- Different resolutions: The primary challenge for the co-registration of airborne LiDAR and photogrammetric tie points: In the real-world test datasets used for numerical testing, an area covered by 10,000 laser points was covered roughly only by one photogrammetric tie point.



• tie points ■ laser points

Figure 3.15: Laser points and photogrammetric tie points of one flight line acquired in the Oshawa area

The data plotted in Figure 3.15 had an approximate point ratio of 10,000 laser points to 1 tie point. The laser points captured the surface with great detail (approx. 4 points per m^2) while the 3D tie points recovered only coarse changes in elevation.

The following section examines the impact of different datum, map projection, units, and ground coverage on the co-registration model, and workarounds:

- Different datum: If the search- and target point cloud are referenced in different datum, the resulting offsets can be approximated locally through a translation in X, Y, Z and a change of signs. Rarely does a different datum result in flipped axis. But if it does, the co-registration algorithm can be re-run if/when the first run showed poor correlation results
- Different map projection: Some map projections preserve the size of the mapped area, while others preserve other properties like distance or angular relations. The definition of the orientation of the coordinate axis and the location of the origin can differ as well. However, the Z axis always points upwards, and only the X and Y values can vary in sign. Effects of map flattening in case of cylindrical projects (such as UTM) and shearing effects in case of conic projections (such as US State Plane) and azimuthal projections (which are

a special case of conic projections) are well within the noise level of the discrepancies between perfectly fine-aligned 3D tie points and LIDAR data. To be more specific, the scaling effects of map projections are typically in the range of $<1/10,000$ and can be neglected in the context of coarse co-registration, especially if the datasets cover only few square kilometers of land (Grafarend and Krumm, 2000; Grafarend and Krumm, 2002). Thus, the effect of different map projections can be compensated by cardinal rotation, flipping of axis, and by translation.

- Different units: The difference in units acts like a difference in scale and is typically limited to a scale where 1 international survey foot equals 0.3048 meters. What unit is used depends primarily on the historic conventions of the country, when the data was collected (some parts of the USA for example changed from American units to metric units in the 80s and converted back in the 90s), and requirements of the survey project. Sometimes, especially in border zones, datasets from the same area were generated in different units.
- Different ground coverage: If the datasets cover the same area, vertical differences diminish since the X, Y, Z values are de-measured during the coarse co-registration process. If the datasets overlap to a large degree (80% for example) and the terrain is flat or slightly hilly, the vertical differences will be small compared to the signal of the remaining data and a 2D translation in X,Y will be sufficient to model the spatial discrepancies. Depending on what co-registration method is used, a 3D translation may be necessary if the datasets show little overlap.
- Different format of coordinate input files: Often, coordinate files come with a header that labels each column. However, to support full automation and blind input of coordinate files, the co-registration algorithm can be re-run if the first run showed poor correlation results.

Since contemporary research considers mostly only image-based registration algorithms that model only horizontal translation and datasets with full overlap, it seemed reasonable to initially model X, Y- translations only, and add Z-translation later, if necessary. As shown in the ensuing chapters, this strategy proofed fruitful since the co-registration algorithms that were examined in this dissertation already struggled retrieving larger horizontal translations – adding a Z-translation would have caused unnecessary complexity.

For the reasons listed above, a 7-parameter Helmert transform with

- 3 translations (X, Y, Z axis)
- 3 rotations (around the X, Y, Z axis)
- 1 optional (uniform) scale

was chosen as model for fine registration, and a 4-parameter similarity transform with

- 2 translations (X, Y axis)
- 1 rotation (around the Z axis)
- 1 scale

was selected as mathematical transformation function for coarse co-registration. The geometric relationship of the 4-parameter rigid body transformation is illustrated in Figure 3.16 below.




Figure 3.16: The geometric mapping between point cloud-rendered images in terms of translation, rotation- and scale parameters. Illustration after (Bargoti *et al.*, 2012)

This geometric relationship is mathematically expressed as follows

$$\begin{bmatrix} x_p \\ y_p \\ 1 \end{bmatrix} = s \cdot \begin{bmatrix} \cos \varphi & -\sin \varphi & \Delta x \\ \sin \varphi & \cos \varphi & \Delta y \\ 0 & 0 & 1 \end{bmatrix} \begin{bmatrix} x_q \\ y_q \\ 1 \end{bmatrix} \quad (3.1)$$

wherein φ , s , Δx , and Δy are the parameters of the rigid body transformation, specifically, φ is the rotation angle counter clockwise, s is the scale factor, and $(\Delta x, \Delta y)$ are the translations. For the finite discrete images q and p , equivalently it is true for every pixel to be related to each other

$$q(x, y) = p(\Delta x + s * (x \cos \varphi - y \sin \varphi), \Delta y + s * (x \sin \varphi + y \cos \varphi)) \quad (3.2)$$

Using the traditional nadir-only imagery, the tie point generation delivers only a single Z paired with a given (X, Y) , while can be occasionally more than one Z associated with a (X, Y) location in oblique data (Gruber and Walcher, 2014). Without significant loss of information, it is safe to only use the highest Z value and render the point cloud to a single height image. Therefore, only 2.5D data sets are considered in the context of this framework. Also, both photo- and laser coordinates are assumed to be provided in a projection, since standard photogrammetric software packages require coordinate display in a map projection. Thus, Earth Centered Earth Fixed (ECEF) coordinates will not be considered here since their Z axis may not point up perpendicular to the tangent frame.

In summary, for the reasons above, the following *assumptions and limitations* are made throughout the ensuing chapters: (1) *In all data input, Z points up, i.e., away from the ground towards zenith;* (2) *the only scale difference that needs to be considered during coarse co-registration is the difference of international feet and meter unit, i.e., 0.3048 or 3.2808.*

For airborne survey applications, these assumptions do not limit the general validity of outcome of this thesis work. However, these simplifications should be revisited in case of close range-, street side-, and UAV applications where the point clouds are denser and provide more depth information.

4.0 Coregistration in Space Domain

4.1 Introduction

The research presented in this chapter started out with the proposal that it should be possible to discern if two datasets overlap without prior knowledge and without deploying a computationally intense co-registration procedure. If during this “overlap check” it was found that two point clouds overlap, the point clouds would be aligned coarsely, and the coarse alignment would be refined thereafter. During this stage, a fine registration can be conducted resulting in a final set of refined co-registration parameters, or a full data integration can occur, meaning in case the sensor requires calibration, the data is used for both cross-aided sensor calibration and for fine registration, resulting in both refined sensor and co-registration parameters. The whole process of a complete co-registration could be broken down into the four functional steps as Figure 4.1:




Figure 4.1: The functional steps of the proposed co-registration workflow

Here, the overlap check (step 1) compares the datasets for their similarity. If the datasets are similar enough to likely overlap, the coarse alignment is conducted (step 2). Once the point clouds have been coarsely registered, fine registration (step 3a) and optionally cross-aided sensor

calibration (step 3b) with sensor-driven refinement of the point clouds can occur. Lastly, and the coarse- and fine registration parameters are combined and applied to the search point cloud and optionally, the sensor calibration parameters are refined.

The following sections will introduce space-domain based methodologies that can be used for conducting the overlap-check (Section 4.2), for the coarse co-registration (Section 4.3), and for the fine registration (Section 4.4). As mentioned in the dissertation statement, the goal of this research is to find methodologies that do not require prior knowledge about the point clouds.

4.2 Overlap Check in Space Domain

This section introduces five methods that can be used to characterize point clouds and check them for similarity. Each of these methods extracts the information from the point clouds and generates descriptors to analyse the similarity so to draw a conclusion whether the point clouds may depict the same area or not.

In the context of a coarse overlap check, the following five methods for similarity measure are investigated:

- (1) Global cost functions
- (2) Differential surface characteristics
- (3) Graphical models
- (4) Information content (entropy)
- (5) Fractal structures

After each method is introduced in its own, a review of them in context of their applications for the coarse overlap check of laser and tie point clouds is given at the end.

4.2.1 Global Cost Functions

Typically, a cost function is associated with a ratio or a difference between the input and output of a function. Sometimes, a cost function is also associated with the difference between the targeted output and the actual output of a function (Albert Albiol *et. al.*, 2015). Within the research of cost functions, the following combinations of global extrema were considered: (1) minimum, (2)

maximum, and (3) difference in elevation, (4) width, and (5) height of the survey areas, (6) amount of ground coverage, and (7) shape of the footprint.

To extract the 2D footprint (i.e., (7) above), the border lines of a dataset must be identified. One way to extract the point cloud's outline is to mesh it and then query the nodes for their relationship with neighboring points (Vosselmann, 2000). This concept was implemented in MATLAB as a 2.5D Delaunay triangulation, which is illustrated with a small test patch in Figure 4.2.




Figure 4.2: Delaunay triangulation of a small set of points in MATLAB

By meshing (triangulating in this case) the point cloud, an ID can be assigned to every node and cell in order to query the outmost points, locate the border lines, and derive a shape file of these outmost lines from the point cloud.

Shapefiles can also be used to calculate the ground coverage of the point cloud (i.e., (6)). In case where the datasets overlap nearly 100%, the shape and ground coverage can be used as similarity measure. Similarly, comparing their horizontal or vertical extends (characteristics (3), (4), (5), and (6) above), or their extrema (characteristics (1) and (2) above), may allow to conclude if the datasets overlap (if a large overlap of the datasets is assumed).

Comparing the extrema and the shape of point clouds requires, however, that the thresholds have been established, against which the values can be validated. Not only will these thresholds

be different for different types of the topographies, but deriving these thresholds also requires data evaluation, conditioning, and training, which goes against the concept of overlap check without using prior knowledge or assumptions. Furthermore, comparing extrema and shapes work successfully only if every dataset can be characterised clearly and distinguished reliably from other datasets with these parameters. This condition is, however, nonexistent when the datasets only partially overlap, and/or have complex shape, and/or vary strongly in point density, which is exactly the case with the airborne laser- and tie point clouds. Thus, global cost functions cannot be used to determine if a laser- and a tie point cloud are similar or overlap.

4.2.2 Surface Analysis using Differential Geometry

4.2.2.1 Definition and Characteristics

The use of differential geometry for the characterization of digital terrain data is well-known. Since the differential surface characteristics are like a fingerprint of the data, they could perhaps be used as a tool for overlap check and global co-registration.

Surface geometry can be characterized by the first- and second order derivatives (Grafarend and Schwarz, 2001; Wilson and Gallant, 2000). The first order derivative is the gradient vector that describes the magnitude of the normal vector of a local tangent plane and slope for each data node. To characterise the shape of digital elevation models, Li *et al.* (2004) suggested combining the wavelength of the surface signal and the slope into one descriptor: the roughness vector. The surface roughness contains most of the information crucial to characterize surface models (Grafarend and Schwarz, 2001). Figure 4.3 summarizes the first and second order derivatives.




Figure 4.3: Parameters used for geometric characterization of a topographic surface.

Source: Peckham and Jordan (2007) and (*) Li *et al.* (2004)

According to Li *et al.* (2004), slope, altitude, and wavelength are the most important parameters for terrain description. The slope describes the vertical component of the gradient vector while the aspect describes the horizontal component. For a topographic surface $z = f(x, y)$ where the equipotent surface $u = z - f(x, y)$, these parameters can be defined as (Peckham and Jordan, 2007)

$$\text{grad } u = -\frac{\partial f}{\partial x}i - \frac{\partial f}{\partial y}j + \frac{\partial f}{\partial z}k \quad (4.3)$$

$$= -f_x i - f_y j + f_z k$$

$$\text{gradient} = |\text{grad } z| = \sqrt{f_x^2 + f_y^2} \quad (4.4)$$

$$\text{slope} = \arctan \sqrt{f_x^2 + f_y^2} \quad (4.5)$$

$$\text{aspect} = \arctan \left(\frac{f_y}{f_x} \right) \quad (4.6)$$

where x, y, z are the coordinates of a point, f_x, f_y, f_z are the first order derivatives with respect to x, y, z in the gradient vector in the base i, j, k .

Once these parameters about the surface roughness have been determined for both point clouds, maps of these vector fields can be generated. These maps could then be cross-correlated to conduct an overlap check and to establish a coarse co-registration. Also, the minimum, maximum, and average values of the roughness could be computed from both maps so that their values could be cross-compared or compared against pre-defined thresholds.

Figure 4.4, Figure 4.5, and Figure 4.6 plot the elevation profiles of all the two real datasets and one simulated dataset, described in Chapter 3.3 and Chapter 3.4, respectively, using their gradient magnitudes (slope), gradient directions (aspect), and curvature:





Figure 4.4: Oshawa elevation maps in 40m grid: height (m), slope, aspect, and curvature profiles

In Figure 4.4, Figure 4.5, and Figure 4.6, the slope, aspect and curvature values were obtained from the values of a 3-by-3 Sobel kernel in the original profile, around the respective pixel in the output profile. The plots (B1), (B2), and (D1), (D2), allow for an intuitive interpretation – the larger the magnitude of the gradients, the bigger the slope and curvatures. The plots (C1) and (C2) may be less obvious to read: Figuratively speaking, its values can be interpreted as the direction in which the terrain slopes or water would flow (whereas the magnitude of the gradient describes the speed with which water would flow). The gradient direction profiles (C1) and (C2) contain the angles in degrees within the range of $[-180, 180]$ counter-clockwise from the positive x -axis which points to the East. Overall, the derivatives of the laser- (left) and tie points (right) look quite alike.






Figure 4.5: KN elevation maps in 40m grid: height (m), slope, aspect, and curvature profiles

As with the Oshawa data, the differential profiles of the KN laser- and tie points (- left and right plots in Figure 4.5), resemble each other in appearance. From the plots (C1) and (C2), the same interpretations apply as before.

Figure 4.6 below shows the respective output for the X5d4 and X5d5 height maps.




Figure 4.6: X5d5 and X5d4 simulated elevation maps (pix): height, slope, aspect, and curvature profiles

Figure 4.6 presents the slope, aspect, and curvature of the simulated datasets that mimic data collection with a small image footprint, i.e., a long focal length or low flying height, or imagery that has been tied with a key point-based tie point engine and thus resulted in high point density (see Chapter 3.4). As in Figure 4.4 and Figure 4.5, the slope, aspect and curvature profiles of the X5d4 data (plots on the right in Figure 4.6) are coherent with the elevation plot. However, the differential surface profiles of the X5d5 data (plots on the left in Figure 4.6) showed certain abrupt changes and noise-like characteristics. This is because the elevation noise changes quickly from one cell to the next and the Sobel kernel for gradient and curvature calculation covers only a 3-by-3 neighborhood.

This phenomenon shows how important the size of the filter mask is for characterising the elevation profile. Had a larger-sized kernel been used, the left and right plots would have shown similar appearance. However, the purpose was to find out if the differential surface descriptors could provide a means to conduct an overlap check, without prior knowledge or data training. That also implies that the kernel size of the differential surface operators can be left at one common size. Consequently, the fact that the left and right plots show very different characteristics indicates that the differential surface descriptors are not robust enough to recover this change in resolution.

Like the slope, aspect and curvature, the standard deviation, and the range of the values (max, min) of the surroundings of each pixel could also be used to characterize the roughness of the terrain. By default, MATLAB provides a 9-by-9 kernel for the entropy, standard deviation, and range operator (MathWorks, 2019). Figure 4.7 shows an example of the standard deviation and the range profile gained from a 5x5 kernel.




Figure 4.7: Standard deviation (A1, A2) and range filter (B1, B2) profiles of Oshawa elevation maps (m)

As can be seen from Figure 4.7, the standard deviation and range plots appear very similar to the gradient and curvature plots in Figure 4.4 and do not provide any essential new information. Likewise, the standard deviation and range plots from the KN data and the simulated data looked similar in comparison to their respective gradient magnitude and curvature plots. Since they did not add value to further this discussion, they were left out.

4.2.2.2 Discussion

Differential surface plots derive their information sequentially from small, local environments with a predefined mask, i.e., at one scale. The height variation within such local neighborhood is also referred to as the surface roughness. As will be shown in the ensuing chapters, the surface roughness can also be derived from methods that consider surface information at the different scales.

Furthermore, the determination how similar two height profiles are by comparing the differential surface characteristics, requires that the respective point clouds have similar point density, shape, orientation, and a high degree of overlap. Unfortunately, with the airborne tie- and laser point clouds, these are rarely satisfied. Also, one must establish the thresholds against which the validation for similarity can be introduced. Not only will these thresholds be different for the different types of topography, but the derivation of these thresholds requires the data evaluation, conditioning, and training, which goes against the concept of overlap check without using prior knowledge or assumptions. Thus, comparison of differential surface characteristics cannot be used as a reliable means to determine if a laser point cloud and a tie point cloud are similar, or they overlap.

4.2.3 Fractal Analysis

During this research, the idea to use fractal analysis to conduct the overlap check and global co-registration surfaced in discussions. According to Boeing (2016), a fractal is a geometric figure, also referred to as mathematical set, that exhibits a repeating pattern displayed at every scale. The objects that can be modelled by fractal structures show the geometric self-similarity at the progressively smaller scales. Two popular examples of such structures are snowflakes (Nittmann and Stanley, 1987) and the more complex mathematical set of Mandelbrot (Mandelbrot and Pignoni, 1983). The complexity order of a fractal structure, also referred to as the fractal dimension, describes how often the smallest structure is repeated to assemble the largest representation of the geometric pattern, i.e., how much of the fractal is inherited. The schematic diagram in Figure 4.8 illustrates the idea of the repetitive geometric patterns at different scales within the same object of interest for a complexity order of three:




Figure 4.8: Schematic of a fractal snow crystal. Source: Sakuraba (2013)

Figure 4.8 indicates an applicable rule of the recursive drawing for the structure of a snow crystal. The fractal is repeated on three levels, i.e., the complexity order in this schematic plan is three. The red points mark the center of the fractal pattern while the green points are around the previous green points.

By taking these characteristics into consideration, a coarse co-registration approach using fractal geometry would involve the following computational steps:

- (1) Identify and extract a fractal, i.e., a geometric pattern that repeats itself at different scales, in the reference point cloud.
- (2) Locate this pattern in the search point cloud and compare it for similarity. Since the point clouds can be offset by different scales, and since the fractal pattern itself shows at the different resolutions, the *fractal pattern search must be conducted at different scales*.

Thus, using fractals in co-registrations would require that the point density of both point clouds is high enough to capture, with sufficient detail, the fractal, which is supposedly common to both point clouds. In the example of the snowflake, both datasets must provide a point density high enough to sample the fractal on the “same or similar” order of complexity. The following paragraph will illustrate what is meant by *same or similar order of complexity and its relationship to point density*:

The point clouds (in this case the points describing a snowflake) may overlap only partially (one point cloud may for example sample only 60% of the snowflake) and may be set apart by a

scale difference due to the use of the different units. A difference in unit (meter and feet) would invoke an artificial difference in point density of approximately 3. While the intent is to extract from each terrain a scale-invariant fractal signature and then compare them for similarity, the intent is also recover differences in scale between the terrain models. In case of the snowflake, the upscaled fractal of the third degree in the search point cloud would have approx. the same size as the fractal of the first degree in the template point cloud. The upscaled fractal in the search point cloud has however less details than the same-size fractal in the target point cloud. This difference in detail is the only information that allows for retrieving the scale factor: To isolate the snow fractal of the same (third) order in both point clouds requires that the point densities of the search- and target point cloud differ less than the factor 2 (so that both the fractals of the first and second order can be distinguished from the fractal of third order) and that both point clouds have similar SNR. Since airborne tie- and laser point clouds vary strongly in point density and SNR, this condition is false and the fractal analysis cannot, therefore, be used to conduct a coarse overlap check.

4.2.4 Surface Analysis Using Entropy

4.2.4.1 Definition and Characteristics

The concept of entropy - or theory of information content - was introduced by Shannon and Weaver (1949). Since then, it has found many applications (Studholme *et al.*, 1999). Information entropy measures the uncertainty or surprise associated with a random event or the information content of a message. According to Sabuncu (2006), (Shannon's) entropy of a discrete random variable X with a specific probability density function $p_X = (p_1, \dots, p_n)$ is defined as:

$$H(X) = \sum_{i=1}^n p_i \log(1/p_i) = - \sum_{i=1}^n p_i \log(p_i) \quad (4.7)$$

where $0 \log \infty = 0$ and $\log(1/p_i)$ indicates the amount of the uncertainty associated with the corresponding outcome. Thus, Shannon's entropy can be considered as a statistical average of the uncertainty, or information (Sabuncu, 2006).

Such a discrete random variable X can be considered the information source of an image, too. Entropy values can be calculated locally for small areas or for the entire dataset (and serve in a global cost function for global similarity check, see Chapter 4.2.1). The comparison and matching of the entropy content of local cells are what the MI-based fine registration approaches accomplish.

One drawback of Shannon’s definition of entropy is that it only considers the distribution of the histogram values but ignores the neighborhood relations of the pixels. Its implications are exemplified in Figure 4.9, where the pixel locations of a “house” image were shuffled around, yielding the same histogram, and the same global entropy value.




Figure 4.9: Memory-free characteristic of Shannon’s entropy. Source: Sabuncu (2006)

Obviously, the “house” image contains more structure (i.e., less “uncertainty”). However, treating each pixel intensity as an independent sample is an oversimplification (Sabuncu, 2006). This characteristic coincidence is especially problematic when a co-registration is conducted with optical and LiDAR imagery of natural scenes, where the value of a pixel is likely to be close to the value of its neighbors.

Similarly, as Sabuncu (2006) demonstrated, the pixel locations of the “house” image can be changed systematically, and even though the shape of the histogram changes, the total entropy – according to Shannon – remains the same, up to round-off errors. This phenomenon is highlighted in Figure 4.10.




Figure 4.10: Characteristics of Shannon’s entropy – effects of mapped image intensity.

Source: Sabuncu (2006)

While the image in Figure 4.10(C1) has more natural appearance than in Figure 4.9(A1), it still reveals the same geometric information to the human eyes. If the “house” image represented a DEM, then the bright pixel values would indicate the high peaks and the dark pixels would indicate the low elevations. In Figure 4.10(C1), the intensities are almost inverted, like a photo negative of Figure 4.9(A1). Figuratively speaking, the height map was inverted. Clearly, if two height maps have similar footprint, and one shows a mountain, while the other one describes a valley with similar (yet inverse) magnitude and curvature, the entropy of these substantially different maps would be identical.

As mentioned in Chapter 4.2.2, the entropy can also be used to describe the surface roughness of the DEMs. Similar to Equation (4.8), the entropy $E(x, y)$ of a node in an elevation model is defined as (Sharma et al., 2011)


$$E(x, y) = - \sum_{x,y} (p_{x,y} * \log p_{x,y}) \quad (4.8)$$

where $p_{x,y}$ is the probability of the surface height at location (x,y) . Obviously, the entropy calculates the average of the local statistical average of uncertainty. Consequently, the rougher the surface, the higher the uncertainty about the elevation value of the node, and the more information and energy this node contains.

A reduction of roughness can be interpreted as the information loss about the surface. Thus, the entropy can be regarded as (Hu et al., 2015): (1) a measure of the information loss caused by the aggregation of a DEM and (2) a means of comparing DEMs with the DEM products of varying quality (Wise, 2012).

The entropy increases with increase of the information content. To ensure Shannon's entropy as defined in Equation (4.8), the rarer an event, the higher its surprise, the more importance is associated with its occurrence. According to Hu *et al.* (2015), the entropy is the lowest when there is no spatial variability, such as in a flat area. In contrast, the maximum entropy is obtained in areas where the elevation of the surroundings is not easy to predict, e.g., a surface with strong variation in slope. Consequently, when a DEM is under-sampled, the spatial variation in elevation value is low due to the redundant elevation value and the loss of micro-structures (Vieux, 1999).

Figure 4.11, Figure 4.12, and Figure 4.13 illustrate the entropies of the Oshawa, KN, and simulated point clouds. In each entropy map, the output pixel contains the entropy value of the 9-by-9 neighborhood around the corresponding pixel in the input image.







Figure 4.11: Entropies of the Oshawa elevation maps

In Figure 4.11, the plots (A1) and (A2) are the histogram of the Oshawa laser- and tie point maps, and the plots (B1) and (B2) are the respective entropy maps. The histogram of the laser points shows a much narrower distribution, since – as mentioned in APPENDIX C – the South-East corner showed a strong slope as a by-product of the surface extrapolation by MATLAB’s ‘natural-fit’ function. This local anomaly skewed the distribution of elevation values, giving the laser points a larger range in elevation values and a narrower histogram shape. The slope also skewed the entropy plot in plot (B1) whereas in plot (B2) the entropy scattered quite evenly and showed only a few small peaks in the center. The global entropy of the plots (B1) and B2 resulted in 6.4149 and 7.6184, respectively.







Figure 4.12: Entropies of the KN laser- and tie point elevation maps

In Figure 4.12, the plots (A1) and (A2) display the histogram of the Keyworth-Nottingham elevation profiles, and the plots (B1) and (B2) show their respective entropy maps. The histogram of the laser points shows a wider range of the height values since it covered a larger footprint and - similarly to the Oshawa laser point cloud - since it showed the exaggerated slopes at the borders as a by-product of the surface extrapolation by MATLAB's 'natural-fit' function. Compared to the laser points, the tie point cloud had a smaller footprint and relatively more height variation, which led to the increased entropy values in most areas. Nevertheless, since the laser point cloud covered a considerably larger footprint, it amassed a higher global entropy of 7.7323, in comparison with 7.3802 for the tie points.






Figure 4.13: Entropies of the X5d5 and X5d4 simulated elevation maps

In Figure 4.13, the plots (A1) and (A2) display the histogram of the X5d5 and X5d4 elevation maps. The 10m noise band (see Chapter 3.4) that distinguishes between X5d5 from X5d4 caused the histogram in the plot (A1) to have a smoother surface since the noise values filled in the dents in the histogram shown in the plot (A2). The entropy of the X5d4 profile (B2) showed clear spatial structures. The entropy profile of the X5d5 data however showed the noise-like characteristics, since the entropy mask – which, per MATLAB default, was a 9-by-9 kernel – picked up on the 10m noise band in the X5d5 map. Even though the entropy values in the plot (B1) showed stronger variations locally compared to the plot (B2), X5d5 and X5d4 resulted practically the almost identical global entropy values - 7.3720 for X5d5 and 7.3326 for X5d4. This confirmed that adding the noise band to the X5d4 height map only altered the entropy values locally, but did not alter the global entropy, since the noise band did not alter the complexity of the underlying X5d4 DTM surface, which remained the dominant source of elevation change, globally.

4.2.4.2 Discussion

In Section 4.2.4, it was examined if entropy could be used to assess the similarity of height maps and discern if they overlap, assuming – unlike the MI-based methods (which examine the relative entropy of the overlap area) – there is no prior information available about the height maps. It turned out that the meaning of entropy could be interpreted with multiple ways, yet for airborne

geodetic data, especially intuitive seemed the interpretation that entropy is (1) an indicator of surface roughness, and (2) an indicator of mapped complexity.

Due to the circumstance that Shannon's definition of the entropy does not consider the information of neighborhood relations, it has a drawback, since the location of geographic surface elements with respect to each other is important. As indicated by Mishra and Zang (2012) (who considered image luminance, LiDAR elevations, and LiDAR intensity values in the co-registration), the "information-theoretic methods cannot be applied to natural scenes". Sabuncu (2006) therefore introduced the co-registration methods that integrated the advanced formulations of entropy under the consideration of neighborhood relations, such as Reny's entropy. However, these co-registration methods, similarly to those of Pluim *et al.* (2003) and Fan *et al.*, (2012), were optimized for the medical and close-range data where, knowingly, all data depicts overlapping scenery, and the depicted objects of interest are clearly offset from the background.

From the analysis of entropy plots in this section, the following observations and learnings have been made:

- 1) If the footprint of a point cloud is small compared to its changes in elevation, the artifacts caused by the surface gridding can strongly degrade the accuracy and quality of the final elevation map. To achieve reliable entropy-based overlap checks of real-world data, relative larger footprints and/or high point densities are desirable.
- 2) If the kernel size of the entropy- and differential surface operators is too small compared to the locally mapped information content, then the derived surface plots show artifact-like characteristics (the plots on the left in Figure 4.12 and Figure 4.6). Using the entropy as a surface descriptor proofed not robust enough to recover the magnitudes of the changes in scale and resolution, which are typical with the geodetic point clouds. It also implied that the entropy is not a good means to conduct fully-automated overlap check and/or co-registration without prior knowledge or some form of data training.
- 3) If the noises in the (simulated) data change too quickly from one cell to the next, it causes specific abrupt changes of the gradient and entropy values and an artificial increase in the mapped complexity at small scale. One can conclude that the underlying surface could have been sampled at lower frequency, to imitate the real-world geospatial data even better.

In summary, as it was already the case for the methods introduced in Chapter 4.2.2, to determine how similar two height profiles are by comparing the entropies requires that the respective point clouds have similar point density, shape, orientation, a high degree of overlap. With the airborne tie- and laser point clouds, these conditions are however rarely met. Furthermore, thresholds must be established, against which the validation can occur. Not only will these thresholds be different for different types of the topographies, but their derivation also requires data evaluation, conditioning, and training, which goes against the principle that the overlap check runs without any prior knowledge or assumptions. Thus, the comparison of entropy cannot be used as a reliable means to determine if a laser- and tie point cloud are similar or if they overlap.

4.2.5 Probabilistic Analysis Using Graphical Models

There are many tools within the category of graphical models that can be used for co-registration. One of these tools is called Incremental Learning (IL) using graph structures. The concept of it seems more intuitive than the ones of some other approaches of graphical modelling yet is somewhat representative of most approaches within this category.

IL relies on the graphical modelling of the image or point cloud content. According to Förstner (2011), a graphical model is a stochastic model represented by a graph $G(V, E)$ where the vertices $i \in V$ represent the random variable x_i , and edges E represent the mutual relationships. The directed edges (i, j) model the conditional dependencies $P(x_j|x_i)$. The undirected edges $\{i, j\}$ model the joint probabilities $P(x_i, x_j)$. An example for how images can be used as graphs is illustrated in Figure 4.14.




Figure 4.14: Image used as a graph. Source: Förstner (2011) referring to Roscher (2010)

As can be seen from Figure 4.14, the surface patches and respective border lines are extracted from each image in this type of the image analysis. As illustrated in Figure 4.15, the nodes (pixels, triangles, or surface patches), and the edges (defining neighborhoods) can then be mapped to the graph structures of the different complexity.




Figure 4.15: Graph structures. Source: Förstner (2011) referring to Perez (1998)

Various generative or discriminative stochastic models are available to estimate a signature of each image – refer to Förstner (2011) and Perez (1998) for more details -, which could be used for the coarse overlap check. Sparse graphical models for example allow for an easy interpretation. However, sparse graphical models do not necessarily correspond to the sparsity of the point cloud, but rather to the sparsity of the information content of the surface. Using sparse graphical models as a strategy for overlap check would require that both data sources present conjugate surface elements with similar information content (entropy), and thus compatible resolution, so that the stochastic models can evaluate graphs of comparable content. The airborne laser- and photogrammetric tie point clouds, however, do not meet this condition and make graphical models an inadequate method in conducting a coarse overlap check.

4.3 Coarse Registration in Space Domain

Over the course of this research, many ideas and suggestions for coarse overlap check and coarse co-registration in space domain came forth. While the methods for overlap check introduced

in the previous sections could also be used for coarse co-registration, the methods that will be introduced in the next section are suited for coarse co-registration only.

4.3.1 Kalman Filtering

While the KF is a well-known recursive algorithm for predicting, filtering, and smoothing states over time (Kuhlmann, 2002; Farrell, 2008; Grewal *et al.*, 2007; Schwieger, 2003; Wang, 2008), it has not been used yet for coarse registration, to the author's knowledge. To use KF in the coarse co-registration would require large data overlap, the use of identical units, and reference frames. Alternatively, it would invoke processing the topographic data in the same direction and assume that the point clouds show a similar information content at its ends. Obviously, these conditions do not align with the thesis statement, where the tie- and laser point clouds are allowed to be referenced in different units and reference frames, and it is assumed that no prior knowledge about the origin of these point clouds is available. Thus, Kalman filtering is not a well-suited tool to conduct the coarse co-registration of airborne tie- and laser point clouds without any prior assumptions and/or prior knowledge.

4.3.2 Accelerated Key Point Matching

The cross-correlation is a popular methodology to geometrically align data sets. In its traditional form, cross-correlation of 2D elevation raster data invokes a correlation matrix at least in the size of the anticipated data overlap. Since, in case of a global co-registration, the offsets in overlap, orientation, and scale are unknown, brute force cross-correlation would invoke time-intensive processing of many possible alignments.

During this research, the idea of deploying RANSAC and/or tree structures plus key point matching at different resolution to speed up the process of traditional cross-correlation began to surface. One approach seems particularly interesting, which arrange the key point/salient feature extraction from the elevation images through the fast indexing with SIFT or SURF, the accelerated feature matching, and the removal of outliers through techniques like k-d-tree alongside with RANSAC.




Figure 4.16: Co-registration workflow suggested by Persad and Armenakis (2017)

Figure 4.16 displays the co-registration workflow proposed by Persad *et al.* (2017). The authors use log-polar sampling and correlation in frequency domain, by applying a FFT - locally - to key-point areas, to achieve the invariance to scale, rotation, and translation. This concept will be revisited in Chapter 5.0. The authors also deployed SURF to localize key points in image pyramids and utilized RANSAC and nearest neighbor distance ratio (Lowe, 2004) to iteratively isolate suitable point matches.

If the accelerated key point matching was applied in space domain in this algorithm, then the search point cloud must either be sampled in all directions at different scales, which can be computationally intense, or, alternatively, the information about an approximate relative pose must be made available. This would violate the requirement of having no prior knowledge about the datasets. Furthermore, the algorithms that build on accelerated key point matching, like the one presented above, require the point clouds to have similar and compatible point densities, which is not existent with the coarse airborne tie- and laser point clouds.

4.3.3 Discussion

As highlighted in APPENDIX B, typical image-based tie points are extremely sparse and describe different surface characteristics. Therefore, the constraint of the similar point density and information content on the laser- and tie point cloud does not hold. As the analysis in the previous sections revealed, this detriment made the use of graphical models, differential geometry, entropy, and fractal analysis as a means of coarse overlap check impossible and it eliminated the use of accelerated key-point matching for co-registration.

The use of a KF for co-registration of topographic point clouds or height maps on the other hand requires pre-knowledge about the orientation and sampling distance. At the same time, the use of global cost functions for overlap check requires not only pre-knowledge about the point clouds, but also a training phase to derive thresholds for similarity (when are two point clouds similar enough to possibly overlap?), which may be laborious. Some of these threshold values may be representative for one type of topography, but not for another. Hence, none of the examined methods in this section met the requirements of being automatic, generic (applicable to many different datasets), and reliable; yet these requirements were part of the thesis proposal (Chapter 1.3).

If the tie point densities were high, some of the above methods may anticipate great success. And then, some of the already published methods for global matching – for example, displacement correction based on self-similarity (Huang and You, 2012) or matching of classified points (Li et al., 2012, Wu et al., 2010) – might probably perform well, too.

4.4 Fine Registration in Space Domain

4.4.1 Introduction

As highlighted in Chapter 2.3, the problem of the fully automated fine registration can be considered solved for most applications:

- For the fine registration of point clouds with similar point density, ICP (Besl and McKay, 1992) and refined methods thereof such as ICPatch (Al-Durgham, 2014) have become the standard over the past two decades (Wyawahare *et al.*, 2005).

- For the fine registration of images, many FT-based approaches have been presented that can provide the sub-pixel registration accuracy (Stone *et al.*, 2001; Zhang *et al.*, 2009). While most researchers considered optical images and LiDAR intensity images only (Mishra and Zhang, 2012), some (e.g., Zhang *et al.*, 2009) also considered LiDAR elevation maps (which contain more information about the topography than LiDAR intensity images).

In some cases, these methods may not deliver the optimal solution. Hence, the ICP for example registered the point clouds with a 20cm vertical error (see APPENDIX G for details). The fine registration approaches that are customized to the application may be preferred (Akca, 2005; Gelfand *et al.*, 2005; Weik, 1997). One category of these customized methods is based on minimizing the discrepancies between point cloud surfaces and outperforms ICP for DSM applications (Gruen 1996; Gneeniss *et al.*, 2015). Another category focuses on minimizing the discrepancies between the special non-point tie features such as line intersections and planes (Mitishita and Habib, (2008), Theiler and Schindler, (2012). Also, refer to section 2.3 for more details. Unfortunately, none of these algorithms were designed with using extremely sparse tie points.

These circumstances raise the question how extremely sparse photogrammetric tie points could be used for fine registration, especially in the context of cross-aided sensor calibration, since LiDAR and medium format cameras are typically flown synchronously. As highlighted in Figure 4.17, there are several options where co-registration of airborne LiDAR and camera data can occur, however data alignment during cross-aided sensor calibration leads to the best possible geometric fit (Al-Durgham, 2014; Habib 2006; Moussa, 2014; Niemeier, 2008; Schneider, 2009).




Figure 4.17: Integration of laser points and optical data ⁽³⁾

As highlighted with the dotted lines in Figure 6.102, there are four areas where the co-registration of LiDAR and camera data can occur: (A) During cross-aided sensor-calibration and data refinement, (B) during DSM orthophoto production, (C) during DTM orthophoto production, and (D) when colorizing the laser points. The 3D tie planes from the Lidar, marked with a black (*) in Figure 4.17, have been used only for LiDAR calibration (Lindenthal *et al.*, 2011; Lindenthal *et al.*, 2015), however use with image data will be presented in the following sections.

In Chapter 4.4.2, the methods within category (A) will be investigated. The insights and findings that were made during this research lead to the proposal of four additional methods, which will be presented in Chapter 4.4.3, 4.4.4, 6.5, and 6.6.

³ The (*) in Figure 4.18 indicates that the quality of the sparse photogrammetric DTM may be low

4.4.2 Fine Registration and cross-aided Sensor Calibration using unclassified Tie Points and LiDAR Roof Planes

4.4.2.1 Methodology and Implementation

A) Theoretical Concept

During the rectification of the LiDAR data the laser point cloud planar patches are extracted automatically to serve as the tie features in the block adjustment of. The reader may refer to APPENDIX D for details about the extraction of planar roof patches. Similarly, photogrammetric tie points for the block adjustment of the optical data are extracted automatically from the imagery. Since the extracted planes from the LiDAR data are determined from numerous points, they offer more redundant and more precise registration elements than the single points used in ICP (Al-Durgham, 2014).

Under these considerations, the proposed co-registration algorithm uses unclassified tie points and roof patches as the registration elements. This concept is illustrated in Figure 4.18.




Figure 4.18: Orthogonal minimization of the tie-point-to-roof-plane distance

As illustrated in Figure 4.18, the spatial discrepancy between a tie point and a roof plane is minimized in the normal direction of the plane's normal. The minimization is realized in an LSA in which the incremental improvements to the sensor calibration parameters can be added as the conditional unknowns (Moussa, 2014; Niemeier, 2008; Schneider, 2009).

As Gneeniss *et al.* (2015) suggested intersecting the LiDAR roof polygons to roof peak lines and roof peak points (and use these points to geo-reference the camera data), the use of the roof

polygons themselves (to geo-reference the camera data) is in focus here. In analogy to the co-registration method used by to project LiDAR features from object space into image space (Habib *et al.*, 2005), the roof planes could be projected into the image space by minimizing the discrepancies between the tie features in image space using the LSA. However, as the constraint for co-planarity of the LiDAR planes is formulated in object space during LiDAR adjustment, it makes more sense to conduct the matching of tie points and laser planes in object space, too. The suggested methodology for co-registration and cross-aided sensor calibration in object space can be split in five functional components, which are outlined in Figure 4.19 below:




Figure 4.19: Main components of the proposed co-registration and cross-aided sensor calibration method

As illustrated in Figure 4.19, input into the 5-step process of the co-registration and cross-aided sensor calibration and data refinement are: (a) the 3D photogrammetric tie points and (b) the planar patches from the laser point cloud. In the routine outlined above, the laser point cloud has already undergone calibration and thus the laser patches can serve as the ground truth for the camera. If the camera had been calibrated first, the tie points could serve as the ground truth for the LiDAR. The matching of the tie points with the laser patches, highlighted as step (c) in Figure 4.19, and the LSA, step (d) and (e), were discussed and developed conceptually with Dr. Peter Friess in

2012. The process of the fine registration was covered through steps (c), (d), and (e) where the fine registration can occur simultaneously with the camera self-calibration, indicated by the two arrows at step (d). The process of matching the tie points with the roof planes was implemented in MATLAB (MathWorks, 2018) and will be discussed more in detail in the ensuing paragraphs.

A conditional LSA for a simultaneous estimation of camera- and/or LiDAR calibration and co-registration parameters, subject to minimizing the point-to-plane distances, was formulated in the collaboration with Dr. Peter Friess in 2013. However, since Teledyne Optech planned on extending its LMS software (Friess, 2006; Lindenthal, 2011; Lindenthal et al., 2015) for a new cross-aided camera calibration module at that time, the mathematical formulations of the LSA have been left out here. The reader may refer to Al-Durgham (2014), Moussa (2014) and Schneider (2009), Habib (2006), and Niemeier (2008) for the LSA formula in similar applications. The proposed research method is however outlined in the form of a conceptual extension of the LMS software in Figure 4.20 below.




Figure 4.20: Sensor driven calibration and co-registration of LiDAR and camera using tie points and roof patches as registration elements

In Figure 4.19, the five essential components of this registration method were marked with the letters from a to e. In Figure 4.20 these components are tinted in grey. Whilst Figure 4.19 outlined the generic co-registration concept, Figure 4.20 emphasizes the in- and output of data.

B) Numerical Implementation

Four software suites were used to implement and test the co-registration method outlined in Figure 4.19. As shown in Figure 4.21, 3D tie points were generated using ISAT software (Hexagon Intergraph, 2018). The laser point cloud was geometrically rectified, and planar patches were extracted automatically using Optech LMS software (Friess, 2006; Lindenthal *et al.*, 2015). The tie points and the laser planes were exported in form of shape files.

LMS software extracts both tilted and flat-leveled planes from the laser point cloud. Per default, LMS exported only shape files of sloped planes. Matches of tie points and sloped planes allow for solving all parameters of the rigid body transformation. TerraPhoto software (Soinien and Burman, 2005) was used to visually inspect selected tie-point-plane matches. The relationship between the software components and the data flow is illustrated in Figure 4.21 below:




Figure 4.21: Software components and data flow of the implemented co-registration method

As shown in Figure 4.21, the MATLAB module used to match the tie points and laser planes covered the following steps:

1. Read tie points from text file (ID, X, Y, Z, nr rays, stdX, stdY, stdZ)
2. Read shape files (requires MATLAB Mapping toolbox)
3. Save information from shape files and tie points into the customized structure shown in Figure 4.22 below.

list(1, 1)		list(1, 1).A	
Field ^	Value	Field ^	Value
tiePoint	1x1 struct	Line_ID	13
roofPlanes	2	Channel_ID	'_'
S	1x1 struct	Roof_ID	1141
A	1x1 struct	Plane_ID	8576
		Survey_ID	1
		Mission_ID	'56131_Gemini'
		Points	51
		Slope	24.7278
		Azimuth	252.2549
		Avg_scan	9.0503
		Sigma_0	0.0562
		Center_X	6.6751e+05
		Center_Y	4.8625e+06
		Center_Z	100.1920
		list(1, 1).S	
		Field ^	Value
		Geometry	'Polygon'
		BoundingBox	[6.6751e+05,
		X	1x11 double
		Y	1x11 double
		Z	1x11 double

Figure 4.22: Information of a tie point and proximate planar patch in MATLAB

4. Pair tie points with the laser shape file by checking if the tie point falls within the shape files' boundaries.
 - Initially a horizontal bounding box was used
 - After some testing, this step was refined, and the horizontal bounding box was replaced with a cube that is slightly bigger than the 3D enclosure of the roof polygon.
5. Compute the centroid of the bounding box and the normal vector as described in Lindenthal (2003)
6. Calculate the dropped perpendicular foot location, i.e., the plane intersection point (Figure 4.18), and the 3D tie-point-to-plane distance
7. Minimize the point-to-plane distance of all isolated point-plane pairs while estimating parameters of the rigid body transformation. Alternatively, minimize all point-to-plane distances while estimating camera calibration parameters. As shown in APPENDIX B,

estimating all camera intrinsic parameters and the rigid body parameters simultaneously requires special flight patterns and/or additional constraints.

The implementation of steps 1-6 was completed in MATLAB software. Several unpredicted findings were made when executing step 6 that lead to skipping the implementation of step 7. These findings will be discussed in the next chapter.

4.4.2.2 Results

The above discussed method was tested with both real-world datasets. Image 9, which is in the centre of the Oshawa data set, was studied in great detail. APPENDIX C and G illustrate the tight fit of the tie points and the laser surface. Due to the high overlap between the flight lines, the automated plane extraction produced a much denser flat patches than originally anticipated: 277,447 laser roof planes were identified in the Oshawa area and 126,438 in the KN area. Figure 4.23 below shows an example of another Oshawa dataset that had roughly the same point density of $4/\text{m}^2$ as the Oshawa dataset used in this analysis.



Figure 4.23: Automatically extracted tie planes in LMS software (Friess, 2006). Planes in red were included in the LSA

Within the area of roughly 4-by-8km in Oshawa displayed in Figure 4.23, the LMS software identified about 500,000 planar patches. Automatically extracted planes are depicted in turquoise and planes selected for refinement of the laser point cloud are depicted in red.

The analysis of the customized storage structure (Figure 4.22) of the Oshawa and KN data showed that roughly every tenth roof polygon, that was classified as potential match with a tie point, showed extremely irregular boundaries. Figure 4.24 below shows a few examples. Here, the stars indicate the 2D positions of the tie points that were matched with the roof polygons on display.




Figure 4.24: Examples of very irregularly shaped LiDAR roof polygons (m)

Visual inspection of these oddly shaped roof outlines revealed that following factors contributed to the non-rectangularity of the shapes: (A) complex roof geometry, (B) effects of relatively low point density compared to the size and complexity of the roof top, (C) some of the material on the roof may have absorbed laser pulses while other materials may have provided better reflection.

In image 9 of the Oshawa data, 97 tie points fell within the 3D bounding box of the roof polygons. The closer inspection revealed that the tie points located almost never flat on the roof planes, but rather on roof peak points or corner locations of other objects nearby. This observation was confirmed when analyzing the other tie point locations in the Oshawa and KN dataset. Figure 4.27 shows some examples:




Figure 4.25: Tie points located in corners of surface patches


In Figure 4.25 above, screenshots (a) and (b) were excerpt from image 9 of the Oshawa dataset and screenshots (c) and (d) were taken from image 58 of the KN dataset (see APPENDIX C for more information). Figure 4.26 below shows more examples from the Oshawa dataset where some of the tie points fell at elevated corner locations while others fell flat on the ground.



Figure 4.26: Examples of image-based tie points located at elevated corner locations and flat on the ground

As can be seen from Figure 4.26 above, ISAT-generated (“image-based”) tie points seem to situate mostly at three types of locations: (1) Corner locations of elevated objects, (2) flat on the ground at locations with well-defined color contrast, and (3) on geometrically complex objects near, yet slightly away, from the boundaries of nearby objects.

A closer look at UltraMap generated (“key point-based”) tie points, as illustrated in Figure 4.27 below, revealed similar point characteristics. This confirmed that tie points in general – both image-based and key point-based – tend to locate within these three types of locations. Tie points within category (1) and (3), i.e., tie points that locate at or near border lines or at corner locations of buildings or elsewhere not flat on the ground, make the assignment of point-to-plane correspondence challenging.






Figure 4.27: Key-point generated tie points situated at elevated corner locations and flat on the ground

Figure 4.27 illustrates UltraMap-generated tie point locations. Unfortunately, UltraMap (Wiechert *et al.*, 2012) allowed for collective tie point display only with low resolution imagery in the background. In the scenery above, points that located on the eaves of buildings (category 1) were highlighted in red. They made up roughly 40% of all tie points.

Since the proposed fine registration methodology was intended to function in conjunction with the cross-aided sensor calibration, the data either from the camera or from the LiDAR are still uncalibrated at the time when the co-registration occurs. As a result, either the camera or LiDAR point cloud will be distorted by sensor-intrinsic-, installation-, and other errors. Thus, when testing this co-registration methodology with the Oshawa dataset, the laser- and tie point cloud were shifted with respect to each other, and the camera data was geometrically distorted within itself.

As can be seen from Figure 4.28 below, the laser point cloud was offset by approx. 2m before self-calibration and, as illustrated in Figure 4.29, the camera data was shifted by approx. 3m with respect to the self-calibrated LiDAR.




Figure 4.28: Roof planes and roof peak lines before and after rectification of the laser point cloud

Figure 4.29 above shows the geometric fit of the laser point cloud before and after the LiDAR LSA. Before self-calibration, plot (a), the horizontal roof line separation was approx. 1.5m in North and 2m in East-West direction. After LiDAR sensor calibration and rectification of the point cloud in Optech LMS, plot (b), the separation was within the noise level of few centimeters. As displayed

in plot (c), the vertical separation reduced from roughly 40cm before adjustment to approx. 15cm after adjustment. Figure 4.29 below displays the horizontal discrepancies of the LiDAR data and imagery before co-registration and camera calibration with respect to the laser point cloud.



Figure 4.29: Visual overlay of LMS roof polygons and imagery of the Oshawa dataset

The scenery depicted in Figure 4.29 shows a typical North-American suburban scenery – here, gable roofs (marked with orange numbers) are often the prominent roof type while saddle roof tops (marked with red numbers) tend to be more prominent in European flight data. The roof tops highlighted with an asterisk show that the LMS calibration engine did not take in all roof tops from the scenery. Also, some of the polygons, for example roof top 7 and 9, exceeded the actual physical boundaries of the roof plane slightly.

Figure 6.114 also shows that the self-calibrated LiDAR data and the un-calibrated camera data were shifted by a few meters in East-West direction with respect to each other. This error pattern is typically associated with different use of reference ellipsoid or reference epoch. The earlier mentioned offsets of the uncalibrated sensor-intrinsic parameters could contribute to the spatial

offset as well. As a reference, APPENDIX G shows the quality of fit of the geometrically rectified LiDAR data and tie points after co-registration.

Figure 4.30 below shows examples of some challenging point-plane combinations where tie points fell within the XY boundaries of roof polygons. Yet they were only about 1.5m distant from the roofs in Z. When looking at the point-patch pairs in 3D and considering that the LiDAR and imagery were offset 2-3m in XY, it seemed likely that the tie points located on surface elements other than the displayed roof tops.





Figure 4.30: Tie points that fell into the 2D boundaries and close to roof polygons due to spatial misalignment of the point clouds (m)

The tie point plotted in (a1, a2) in Figure 4.30 above situated within the horizontal roof boundary and approximately 1.5m below the roof polygon. While it is obvious to a human operator that this tie point located at an object somewhere near a wall, accurate knowledge about the vertical offset between the camera and LiDAR data would be required prior to co-registration to automatically exclude this point from the LSA. Similarly, the tie points plotted in (b1, b2) fell within the 2D roof peripheries and showed approximately 1.5m vertical offset yet seemed to locate away from the roof tops.

Figure 4.31 below shows examples where tie points fell just slightly outside the XY circumstance of their associated roof polygon. Yet, when looking at the point-patch pairs in 3D, it seemed that the tie points had been matched correctly with the respective roof tops and that the spatial offsets stemmed from inaccurate pre-alignment of the data.






Figure 4.31: Tie points that fall outside the roof polygon due to spatial misalignment of the point clouds (m)

In Figure 4.31 above, graphs (a1, a2) show a tie point that situated close to the corner of a very large roof top. As in Figure 4.30, the tie points plotted in (b1, b2) and (c1, c2) located roughly 1.5m elevated from their matched roof polygon, however unlike in Figure 4.30, these tie points situated more favorably with respect to the roof peripheries and matched correctly.

Many more tie elements were examined separately in the imagery and in the laser point cloud – see APPENDIX G for some examples. The respective findings are summarized in the next section.

4.4.2.3 Discussion

During this research, hundreds of point-plane matches (roughly 100 per flight strip) in the laser point cloud, and in the imagery were examined using MATKAB (MathWorks, 2019). It turned out that less than 20% of the tie points fell flat onto a sloped roof top – mostly when a change in roofing material or the shadow of a nearby tree created a strong difference in brightness on the roof. In most instances however, the tie points located at corners, at peak points, on chimneys, on trees, and other objects close by that were not part of the roof plane.

Through this research it became clear that matching tie points and roof planes based on distance only leaves a lot of uncertainty associated with every match, because the characteristics of the tie points and the roof planes oppose rather than complement each other. Large spatial discrepancies, in the range of a few meters in magnitude, as displayed in Figure 4.29, are typical for data from the sensors that have not undergone in-situ calibration yet. The close examination of many tie-

point-plane matches revealed that these discrepancies could induce great challenges in the co-registration process:

- Because of the opposing characteristics of tie points and roof polygons, reliable matching is challenging even when the data is pre-aligned within few decimeters
- To compensate for misalignment in the range of few meters (Figure 4.29), the search for potential point-plane matches must be extended considerably beyond the boundaries of the roof patch
- Since tie points tend to locate at object-peak locations, there could be more than one roof plane in proximity and multiple planes could be associated with one peak point. Extending the search space by a few meters can lead to many more potential point-to-plane associations
- The magnitude of misalignment of the datasets is typically unknown at the time of co-registration. In case of co-registration for cross-aided sensor calibration, it can range anywhere from few decimeters up to a few meters in XYZ.

These factors make it difficult to establish a clear, reliable correspondence of tie points and roof planes.

One could oppose these challenges with the following considerations:

- The matching process could be refined, for example by replacing the cubical search space with a 3D search space that extends parallel to the roof periphery and/or by integrating aspects of slope, roof size, and relationships with neighboring planes in the decision making.
- Additionally, one may deploy RANSAC (see APPENDIX A) to establish a reliable set of registration elements quickly and improve the reliability of the matching results.

Such refinement strategies would however complicate the matching process. Also, the assumptions about the level of spatial discrepancy would have to be made. Naturally, these assumptions could be erroneous and may degrade the co-registration results.


During the visual inspection of the tie points in the Oshawa and KN data, it became obvious that many of the tie points were situated flat on the bare earth surface. This raised the idea to not

only consider the sloped planes, but all the planes in the matching process and perhaps to additionally classify the tie points into ground and air points.

4.4.3 Fine Registration using unclassified Tie Points and classified LiDAR Planes

4.4.3.1 Methodology

The fine-registration method in Chapter 4.4.2 considered only the sloped roof polygons as LiDAR-tie features. Subsequent experiments revealed however that many of the tie points situate flat on the ground (Figure 4.32). Thus, horizontally the flat LiDAR polygons could also be included in the data integration.






Figure 4.32: Examples of tie points that situate on bare earth surface

By including the complanate polygons in the parameter estimation, more observations can enter the LSA than in the method introduced in Chapter 4.4.2. However, the slope of the LiDAR patches must be queried since the non-sloped patches can only be used for minimizing vertical discrepancies between point clouds – for example by estimating the elevation component of the rigid body transform, or through estimating the focal length of the camera, or the scale component of the laser scanner. The co-registration process is outlined in Figure 4.33.




Figure 4.33: Data integration using both sloped and complanate polygons as registration elements

Figure 4.33 above illustrates an extended version of the algorithm outlined in Chapter 4.4.2. The (*) indicates that the spatial discrepancies can be minimized by optimizing camera- or LiDAR sensor calibration parameters, also, instead of only fitting a rigid body transform.

4.4.3.2 Results and Discussion

The co-registration introduced in 4.4.3 adds the complanate polygons in addition to the roof polygons to the matching process - to increase the number of observations of the LSA toward improving the reliability of the parameter estimation. However, a closer look at the LMS-extracted planate polygons revealed the opposing characteristics of the tie points. At the same time, the laser patches observed in 4.4.2 affect the sloped and complanate tie-point-polygon pairs alike. This can be seen in Figure 4.32, where tie points situated at corner locations between surfaces of different material: Optech LMS (Friess, 2006) tends to extract planes on roads, driveways, and roof tops where the surface material is smooth and homogenous. As with the roof planes, these planate polygons tend to have very irregular shape. Due to effects of unknown boresight and unknown sensor-intrinsic calibration parameters, the effect of opposing tie-feature characteristics is enhanced by the displacement of the datasets. Thus, while the complanate features increase redundancy, they contribute mostly only problematic observations to the LSA.

4.4.4 Fine Registration using unclassified Tie Points and locally fitted Surface Patches

4.4.4.1 Methodology

The experiments and considerations described in Chapter 4.4.2 and Chapter 4.4.3 led to the conclusion that the LiDAR tie elements should accommodate the characteristics of photogrammetric tie points more closely so that co-registration can occur more effectively. Specifically, the LiDAR features should either wrap around the tie points more closely, or the LiDAR tie features should follow the boundaries of the object, which they describe, more closely, since tie points tend to locate at corners and extrema of objects.

Consequently, during discussion in supervisor meetings, it was proposed to use locally-fitted polynomial patches into the LiDAR and use these as LiDAR tie elements. At first sight, the suggested concept seemed similar to the ICPatch method by Al-Durgham (2014), however, rather

than establishing a point-to-patch correspondence between a point in the search data and the closest TIN in the template data, this approach uses all laser points within a certain range to fit a polynomial patch that is used to minimize the point-to-patch distance. The proposed method is outlined graphically in Figure 4.34 below.




Figure 4.34: Data integration using locally-optimized surface patches as LiDAR registration elements

Figure 4.34 above illustrates an extended version of the algorithm outlined in Chapter 4.4.3. The (*) indicates that the spatial discrepancies can be minimized by optimizing camera- or LiDAR sensor calibration parameters, or by fitting parameters of the rigid body transform.

This method provides the following advantages over ICPatch: LiDAR patches are extracted from numerous points which decreases the impact of noise and allows to model the local environment more accurately, if knowledge about the object on the surface and/or the accuracy of the coarse alignment is given. For example, if it is known that a tie point locates on bare ground in open space, and the point clouds are coarsely aligned within few decimetres, a polynomial surface will allow for a more accurate fit with the tie point than a local TIN cell. Polynomials also provide insight to surface characteristics, such as gradient/slope, curvature, and surface roughness which could be used to classify the surroundings of tie points indirectly from the locally optimized laser patches. These geometric attributes could be used to individually weigh the contribution of each misclosure vector in the LSA of the co-registration parameters.

4.4.4.2 Results

Figure 4.35 and Figure 4.36 below show examples where polynomial surface patches of first and third order were fit into a set of points that covered roughly 10mx10m ground. In these two examples, without restricting the general validity of this experiment, control point locations were used instead of tie point locations since control points are easy to locate in the imagery and typically situate in flat, open areas. Since LiDAR and camera datasets were known to be 2-3m offset in XY (see Figure 4.29), laser points within a search radius of 5m were included in the local point cloud.




Figure 4.35: Locally-fit surface patch of (a) first and (b) third order wrap around ground control point 27 of the Oshawa data (m)

GCP 27, indicated with a green triangle in Figure 4.35, located at the corner of a stop line of a sloped street in Oshawa. This control point looked practically identical to GCP 25 and 261, illustrated in APPENDIX G, and control point 26, illustrated in Figure 4.32 and was surrounded by 69 laser points within a 5m search radius. Knowing that GCP 27 situated in the middle of a slightly sloped road, the surface patch of first order (a) seemed the best representative of the environmental conditions, even though the surface patch of third order (b) showed the smaller residuals.




Figure 4.36: Locally-fit surface patch of (a) first- and (b) third order around ground control point 58 of the Oshawa test site (m)

GCP 58, indicated with a green triangle in Figure 4.36, situated at the corner of a quadratic sewer cover at the side of a road. It looked identical to GCP 64 and 88 in Figure 4.26 and GCP 92 in Figure 4.32 and was surrounded by 34 laser points within the 5m search radius. Knowing that the point situated close to the curb of a street, a surface patch of third order (b) seemed a good representation of the environmental surrounding. As with GCP 27, the surface patch of third order (b) showed smaller residuals than the respective planar patch (a). It also showed smaller residuals than the polynomial of first order (a).

As can be seen from Figure 4.35 and Figure 4.36, both surface fitting methods were able to confirm the vertical offset of 20cm between the laser and optical data that existed prior to co-registration. As shown in APPENDIX G (reference solution generated with TerraPhoto), the point clouds showed a much tighter fit after co-registration.

4.4.4.3 Results and Discussion

As the experiments in Chapter 4.4.2 and 4.4.3 highlighted, the tie points tend to situate at the outskirts and corner locations of the objects whilst the polygons of the automatically-extracted LiDAR planes don't follow the boundaries of the respective surface elements close enough to describe them accurately. This phenomenon turned out to be the most prominent drawback of the aforementioned fine registration methods. Also, before the cross-aided sensor and data calibration, the LiDAR and camera data are typically offset by a few decimeters up to several meters, which worsen this shortcoming.

The method introduced in this section addressed these limitations by replacing the automatically extracted polygons with the locally fitted surface patches. Since the local patches no longer describe planar objects that are adjacent to the tie point locations, but rather sample the laser point cloud around the tie point location, the magnitude of misalignment prior to calibration has less impact on the feature matching (as long as it does not exceed the radius of the locally fitted surface patch).

Choosing the order of the polynomial so that it represents the environmental conditions correctly requires however knowledge about the geometric surface characteristics in proximity of each tie point. Simply choosing a higher order polynomial will lead to smaller residuals in most

cases, however it may also lead to a misrepresentation of the environmental surroundings and degrade the co-registration results.

Classifying both point clouds into, for example, ground, non-ground, vegetation, and hard surface, may help isolate tie points that do not contribute positively to the co-registration and should be excluded from the LSA. These considerations created the intent to automatically classify both point clouds before integration, which will be discussed in 4.4.5.

4.4.5 Fine Registration using classified Tie Points and classified LiDAR Patches

4.4.5.1 Methodology

The fine-registration method suggested in section 6.4 integrated the LiDAR data with the tie points in the form of the locally-fitted surface patches. Lower order polynomials could be used to model ground surface patches while high order polynomials (that nest tie point locations better than planes) could be used to model surface areas above the ground and in transitional zones, where the laser points sample the objects partially on the ground and partially above ground.

Figure 4.37 below illustrates how the classification of the point clouds and point cloud patches could aid the co-registration process. Essentially, the co-registration workflow presented in Chapter 4.4.3 is extended for the DTM- and DSM classification of the tie points:




Figure 4.37: Integration of the classified tie points and the classified LiDAR planes for fine registration and cross-aided sensor calibration excluding the tie points in transitional zones

The co-registration method illustrated in Figure 4.37 matches the coplanate polygons with the tie points that fall on the bare ground and excludes the tie points that situate in transitional areas, i.e., the areas with steep gradient between the ground and horizontal surface areas above the ground. Also, since the magnitude of the data misalignment is typically unknown before data integration, the tie points that are surrounded only by DTM laser points or only by DSM laser points are more likely to make good tie elements. The (*) indicates that the spatial discrepancies can be minimized by optimizing the camera- or LiDAR sensor calibration parameters, also, instead of only fitting the transformation parameters of the rigid body.

Similarly, the co-registration method presented in Chapter 4.4.4 could be extended for an additional classification of the point clouds. Specifically, if the laser points within the proximity of each tie point were classified into ground, non-ground, and optionally other classes, the type of polynomial could be selected for each LiDAR patch according to the local environment. Figure 4.38 summarizes this fine registration concept.




Figure 4.38: Integration of tie points and classified LiDAR patches for fine registration

Figure 4.38 illustrates a refined version of the data integration approach presented in section 4.4.4. The solid lines indicate the classification into ground- and non-ground points while the

dotted lines indicate the optional classification into building- and non-building points (Wu *et al.*, 2010), or vegetation- and hard surface points (Soinien, 2010b) that could aid the co-registration process, too. Where the tie points are situated in transitional zones, the laser points of more than one class can make up the patch which may increase the surface roughness. In practice, this could be accommodated by decreasing the weight of the patch in the LSA.

Rather than classifying the laser points only, as suggested above, both of the laser and tie point clouds could be filtered. Figure 4.39 illustrates an extended version of the co-registration method presented in Chapter 4.4.4 where both point clouds undergo classification routines and terrain points are matched with terrain patches, and the tie points that situate in transitional areas and on surfaces above the ground are matched with the DSM-classified laser patches:




Figure 4.39: Fine registration using the classified tie points and the classified laser point patches

In the above co-registration process (Figure 4.39), both of the point clouds undergo classification into ground- and non-ground points, which allows to weigh the point-patch pairs not only based on their differential surface characteristics (such as gradient and surface roughness), but also on their semantic characteristics. Having the option to group the tie elements based on their semantic properties may lead to a more robust and more accurate convergence in the co-registration. As an example, the tie points that are surround only by the DTM points could be

assigned to a higher weight in the initial LSA than the other tie points. As before, the ^(*) indicates that the spatial discrepancies can be minimized by optimizing camera- or LiDAR sensor calibration parameters, also, instead of only fitting parameters of the rigid body transform.

All three approaches presented in 4.4.5.1 are well supported by specific practical considerations, since an automated classification of point clouds into ground points (DTM) and non-ground points (DSM) is demanded almost always during data post-processing and spatial analysis. The popular LiDAR data processing software products, such as TerraSolid (Burman and Soinien, 2005) and LAStools (Rapidlasso, 2018) allow for the automated extraction of a full-resolution DTM and a “key model” DTM with using a minimal set of points.

4.4.5.2 Results and Discussion

This section introduced three extensions to the integration and co-registration methods presented in 4.4.2, 4.4.3, and 4.4.4. These methods promise to improve the co-registration performance, because they rely on the classification of the point clouds to filter out the tie points that situate in-between ground and surface-top, in order to reduce their negative impact on the co-registration results. Their drawback however is that the classification adds another computational step to the co-registration process. One must ask: How much gain does the additional classification bring to the co-registration in terms of accuracy and reliability? To answer this question, it is crucial to find out how reliable and accurate DTM classification - the most basic and popular fully-automated classification method for airborne geodetic data - of photogrammetric tie points is:

Many methods exist for extracting a bare earth model from DSM data. However, these mechanisms were developed for the dense laser point clouds. When applied to the extremely sparse tie point clouds, these classification methods suffer three aliasing effects (- ordered by their magnitudes below):

- (1) Differences in data acquisition and processing techniques (Maune, 2001). Compared to effects (2) and (3), this effect is negligible.
- (2) The laser point cloud may not be dense enough to sample the surface right next to each tie point and thus won't capture the same parts of the objects on which the tie points are located. On roofs, for example, the tie points tend to locate at the roof corners, on the roof

ridge, or on the top of chimneys, whereas the laser points tend to fall flat on the roof surface away from the corner and peak locations. Compared to effect (3), this effect is however negligible.

- (3) Low point densities and large grid sizes with the tie point clouds lead to the vertical artifacts (Vosselmann, 2000): With increasing of height differences, the important terrain characteristics get lost and the classification errors increase (Carpenter and Hogarty, 2007; Vosselmann, 2000): The more the classification-filter results deteriorate, the more challenging the co-registration becomes. An example is shown in Figure 4.40 (b):




Figure 4.40: Horizontal (a) and vertical (b) differences between 30m resolution SRTM data and 1m resolution reference DEM. Source: Carpenter and Hogarty (2007)

Figure 4.40 (a) shows the horizontal differences between 30m and 1m grid DSM data. As the authors did not provide a legend for the color coding, the plot just serves as an overall qualitative indicator of the information loss of the 30m grid compared to the 1m height map. The graphs in Figure 4.40 (a) show similar resolution to the real-world data sets used in this dissertation where the LiDAR data were collected through multiple flight lines and reached the point densities of 1 to 20pts/m² and where MATLAB suggested a 40m resolution DSM grid for the tie points.

Figure 4.40 (b) describes the vertical errors of a study area that was divided into 16 blocks and the graph in Figure 4.40 (a) shows the height map of one of the 16 blocks. In Figure 4.40 (b), each column shows the average vertical discrepancy of each block. It shows that the height error increases from left to right. The different dots per sub-block represent the values from the different methods used in the evaluation. As can be seen from column 16, a difference in point density of factor 30 (Figure 4.40 (a)) caused an elevation error of 8-12m in magnitude for the DEM data that Carpenter and Hogarty (2007) had on hand.

In view of these findings, it is now examined how the DTM classification methods filter the point clouds and if there exist perhaps specific improved methods that can extract the DTMs from the sparse tie point clouds with higher precision: According to Vosselman (2000) and Ismail and Jaafar (2013) the most common DTM classification methods use either the erosion operators, which rely on uniform point distribution and locally flat, open terrain to extract ground points, or the slope-based filtering that relies on a certain height difference between ground and non-ground points to determine an optimal filter function. Figure 4.41 shows how far apart a potential terrain point can be from its next closest neighboring point, so that the slope-based DTM classification can still deliver a correct result.




Figure 4.41: Thresholds of two slope based DTM filters for the maximum point difference between a terrain point and its nearest neighboring point. Source: Vosselman (2000)

The curves in Figure 4.41 report that the error of the DTM classification and the point distance are related with an approximate ratio of 1/3, which is confirmed in Figure 4.40 (b) as it showed the vertical classification errors of approx. 10m magnitude for 30m grid elevation data.

According to this rule, the DTM-filtering of DSM data with one point per 10m² cell would result in a vertical error of 3-4m, i.e., too large to reliably distinguish between ground and non-ground surface. With the Oshawa and KN data that came with the 40m tie point grids (1m LiDAR grid in native resolution and 40m laser point grid in final resolution), it implies a height error of 8m to 15m caused by the DTM classification.

Figure 4.42 shows two samples of residential scenery of the KN data with relatively high point density. The samples are good representatives for ISAT tie point data and help illustrate how DTM classification can create errors of 10 meters magnitude and more. More examples can be found in APPENDIX G.



Figure 4.42: Close-up view of tie- (turquoise) and control points (red) in KN imagery

In Figure 4.42 (a) some tie points fall flat on street land and others locate on vegetation and corners of small objects in residential yards. The tie point locations in Figure 4.42 (b) are similar, however one tie point locates clearly elevated from the others on a roof top.

Erosion-based estimation of ground surface would start out by locating an opening in the point cloud, using Delaunay triangulation for example, and then fit a horizontal plane. To find such an opening requires however that a small set of points is clearly offset by surrounding elevated points. Similarly, in slope based DTM filtering the height of a point would be compared with the heights of all other points within a certain surrounding using Delaunay triangulation.

In both image scenes of Figure 4.42, regardless which classification method was used, most triangles would cut through buildings and vegetation. The differences in elevation between the tie points (perhaps with exception of the roof tie point in (b)) would be too small – compared to the size of each triangle – to clearly distinguish between elevated points and points that fall flat on ground.

In the data displayed in Figure 4.42 above, the laser point clouds reached point densities of up to 20points per m^2 while the camera data showed less than 0.1points per m^2 . As mentioned in APPENDIX B, most tie point engines provide the option to automatically densify tie points. As well, “image-based” tie point engines like UltraMap (Wiechert *et al.*, 2012; Gruber *et al.*, 2012) natively generate points with higher density, depending on the scenery and processing settings perhaps up to 10 times more than ISAT with default setting. However, a ten time increase in point density may still not suffice for successful DTM classification.

In sight of these shortcomings, it will be examined next, how ICP – probably most prominent and one of the most robust point-based co-registration algorithms – benchmarks with laser point clouds and extremely sparse photogrammetric tie points.

4.5 Conclusions

In this framework, seven methods for assessing the similarity of elevation maps in space domain were researched. Five methods for conducting overlap check were suggested, however the analysis showed they require either similar information content (entropy), or similar degree of mapped complexity (graphical models), or similar point density (differential geometry, fractal dimension), or assumptions about data coverage (extrema), or meta-data or some other form of pre-knowledge about common data characteristics (cost functions). Then, two space domain

methods were for conducting coarse co-registration were analyzed. It became clear, that they required prior knowledge about the reference systems of the data (KF), too, or compatible point densities and surface description (accelerated key-point matching). Another drawback of these methods was that they rely on deriving information sequentially from small, local environments, and that they require similar point densities.

Afterwards, two methods for fine registration and simultaneous cross-aided sensor- and data calibration and three methods for fine registration of laser- and photogrammetric point clouds were proposed and examined. Contrary to the mainstream co-registration algorithms, these concepts were developed to accommodate the characteristics of the extremely sparse tie points and to recycle the by-products of LiDAR calibration and photogrammetric AT. Many of these methods required either the classification of the entire point clouds or the characterization of the laser patches in proximity of the tie points, in the hope that they would add semantic information and increase the accuracy and robustness of the co-registration process.

Each proposed method was evaluated either experimentally or through theoretical analysis and debate. A major finding was that the image-based tie points are too sparse and too irregular for classification methods to work successfully. In particular, it was found out that the DTM classification – the most basic and most popular of point cloud classification methods – requires at least ten times higher point density to be able to reliably distinguish between the ground and air points. In fact, one should expect the height errors of approximately 8-15m in deriving a DTM from the 40m tie point grid of the real-world datasets. Also, the DTM information from the laser data could only aid the fine registration process if the point clouds are already aligned tightly and where the areas with the bare earth exposure dominate the surface data.

In summary, *no method analyzed in this chapter may be preferred in practice and one may have to move to the frequency analysis*. However, the research in this chapter provided insight into what should be investigated next: One drawback common to many of these methods was that they derive information sequentially from small, local environments, and that they require similar point densities. Since frequency analysis considers all the surface information at once, the differences in point density and shape of small environments may be less of a concern. Thus, the possibility to use frequency analysis to achieve automated co-registration will be explored next.

5.0 Coregistration in Frequency Domain: Performance Analysis of Existing Methods

This section examines how methods in the frequency domain can be used to analyze the geometric similarity of the airborne surface point clouds, that have been converted to height raster images. The analysis of these height maps based on their inherent waveforms promises the following advantages over methods in space domain:

- On the contrary to isolating the congruent surface information in space domain, the waveform transforms do not require local sampling and/or feature matching but consider all data of the point clouds in each computational step
- Frequency methods ingest 2.5D raster elevation data instead of 3D vector data
- The frequency analysis using FFT is computationally fast

As shown in APPENDIX B, C and D, both laser- and tie points provided cm-level accuracy, but they have very different characteristics otherwise. As a result, a tie point that situates only few cm apart horizontally from a laser point can differ many meters vertically. This phenomenon, in conjunction with the extremely sparse distribution (a 40m grid size was optimal), limits the level of accuracy at which fine registration can occur (as shown in APPENDIX D, fine registration with ICP resulted in 20cm vertical offset, i.e., an offset of several pixels magnitude, rather than the desired sub-pixel accuracy).

Thus, merely changing from space- to frequency domain will neither change the type of data input nor increase the level of accuracy at which co-registration can occur; rather (as already outlined in section 4.5) would it require another type of photogrammetric point element (such as photogrammetric dense points which assimilate better with laser points). Thus, the research for fine registration methods in frequency domain is not pursued any further, and, as outlined in Figure 4.1 in Chapter 4.0, we first, and only, investigate how coarse co-registration can occur.

While cross-correlation in the space domain is useful in scenarios where the reference signal or desired signal's characteristics are not well-defined or known in advance, and while cross-

correlation can be an optimal tool to provide information about the similarity between the received signal and the reference signal at different spatial positions for spatial analysis applications, the following sections focus on correlation methods in frequency domain. Since the reference signal is known, the frequency methods maximize SNR, providing optimal detection, signal shaping, and efficient implementation make it a powerful tool for signal processing and detection applications.

The frequency-based co-registration methods reviewed in Section 5.1 consider the phase- and amplitude signal of all tie- and laser points at once, which allowed for optimal global co-registration of coarse point clouds. We start with the most basic method, the Phase Only Matched Filter (POMF) (Horner and Gianino, 1984) which recovers translation, followed by the Fourier Mellin Transform (FMT) (Casasent and Psaltis, 1976) which recovers rotation, scale, and translation, and the Analytical Fourier Mellin Transform (AFMT) (Derrode and Ghorbel, 2001).

As the POMF and FMT are well-established, the focus is here on particularly achieving optimal co-registration results with airborne laser- and tie point clouds, which can differ strongly in point density, size of area, shape, and noise level. The following sections will highlight the characteristics of each transform that imply restrictions on their performance with airborne survey data and compare their performance.

All methods ingested 2.5D elevation maps that were gridded from the laser and tie point clouds as described in Chapter 3.0 and APPENDIX C.

5.1 Coarse Registration in Frequency Domain

5.1.1 The Phase Only Matched Filter

5.1.1.1 Definition

Fourier spectra are complex in nature and the geometric information of its image input can be analyzed in terms of its Fourier amplitude and phase. While both elements can be considered when comparing images for similarity, according to both Horner and Gianino (1984) and Oppenheim and Lim (1981), the phase spectrum contained more information being critical for rebuilding the image from the frequency input than the amplitude information of the Fourier spectrum. The filter used by Horner and Gianino (1984), often referred to as POMF (Phase Only Matched Filtering),

or simply the Matched Filter (MF) (Le Moigne *et al.*, 2011), has found its usage in hundreds of image processing applications since then.

By isolating the phase difference and building on the translation invariance of the Discrete Fourier Transform (DFT), the POMF starts out by transforming the two input images f and g which differ only by the displacements $\Delta x, \Delta y$

$$f[x, y] = g[x + \Delta x, y + \Delta y] \quad (5.1)$$

into the Fourier domain. In (5.1), (x, y) are the location of a pixel, and $(\Delta x, \Delta y)$ are the translation. The respective DFTs, denoted by $F[u, v]$ and $G[u, v]$, find

$$\mathcal{F}\{f[x, y]\} = F[u, v] = G[u, v]e^{i(u\Delta x + v\Delta y)} \quad (5.2)$$

$$\mathcal{F}\{f[x + \Delta x, y + \Delta y]\} = F[u, v]e^{i(u\Delta x + v\Delta y)} \quad (5.3)$$

Since the shift between two images is equivalent to a linear phase shift in the Fourier domain (also called Fourier shift property), the translations can be retrieved by applying an adjoint FT to (5.3), which is isolating the term containing the difference in phase:

$$F[u, v]G^*[u, v] = e^{i(u\Delta x + v\Delta y)} \quad (5.4)$$

where $*$ is the complex conjugate. Since the frequency spectra F and G are complex functions, their magnitudes, $|F[u, v]|$ and $|G[u, v]|$, are independent of Δx and Δy . In absence of noise, the phase correlation matrix is divided by its magnitude to normalize the phase difference $e^{i(u\Delta x + v\Delta y)}$ from F and G

$$CPS[u, v] = \frac{F[u, v]G^*[u, v]}{|F[u, v]G[u, v]|} = e^{i(u\Delta x + v\Delta y)} \quad (5.5)$$

The multiplication of the amplitude of two DFT spectra is typically mentioned as the power spectrum. Since the power of a variable y is defined as the average squared amplitude, $C[u, v]$ is often referred to as the cross-power spectrum (CPS), which may mislead, since it describes the normalized phase difference of two DFT signals. Indeed, it is the spectral implementation of the cross-correlation and needs to be more accurately referred to as the cross spectral density matrix (Wilmer, 2012).

The multiplication of the DFT spectra with each other is equivalent to a convolution in object space. Obviously, the convolution of the phase signals will be at a maximum single peak when the signals are identical, i.e., the phase difference is 0. The translative movement between F and G can therefore be determined from the location of the maximum of the phase correlation, by inverting the cross-power spectrum. The inverse of the location where the phase difference is at a minimum is equivalent to the displacements of the two images:

$$[\Delta x, \Delta y] = IFFT\{CPS[u, v]\} \quad (5.6)$$

As will be shown in Chapter 5.1.2, the translation offsets can also be used to recover angular and scale offsets of two images.

5.1.1.2 *Properties*

The purpose of the POMF is to free its data input from translational offsets and align it by cross-correlating the frequency spectra. While this is similar to cross-correlation in space domain, the traditional cross-correlation depends primarily on the energy (square amplitude) of the image rather than its spatial structures, and it leads to a sharp correlation peak only in the absence of noise. Thus, in case where the input DTMs characteristics are not well-defined or known in advance, cross-correlation can provide information about the similarity between the elevation models at different spatial positions, which can be valuable for spatial analysis applications. However, within this given framework, reference signal (laser point cloud) is known in advance. Knowing the reference signal allows to weigh the two signals such that the SNR is maximized, yielding a higher level of noise suppression and sharper correlation peak than what can be achieved traditional correlation techniques (Chen et al., 1994). This can be seen from Equation (5.5), where

the denominator is normalized (division by the amplitude), which makes it robust to noise that is correlated with the images, such as the uniform variations of illumination or offsets in average intensity (which for example can represent the offsets in terrain height) (Chen *et al.*, 1994). The POMF also provides sharper correlation peaks compared to traditional cross-correlation in space domain and the traditional matched filter (Wolberg and Zokai, 2000). Because of this advantageous correlation characteristic and its computational efficiency, the POMF has been widely accepted as the standard algorithm for recovering the translational misalignment.

5.1.1.3 Co-Registration using POMF

In case that the noise level is not strong, and the height profiles are similar, there should be only one single correlation peak (Brown, 1992). However, as highlighted in Chapter 3.0, the LiDAR and photogrammetric data from the same scene are likely to vary in overlap, size, details of the sampled surface (resulting in differences of phase and amplitude), height (small differences in intensity), and level of noise. These differences cause an increase of the noise level of the spectral phase which can result correlation spikes with similar magnitudes rather than one pre-dominant spike. By sorting the IFFT's matrix by magnitude, considering more than one spike, applying the respective translational offsets - one at a time, and feeding the POMF with the respective translation-corrected image anew, it becomes clear which of the correlation candidates reflects the true maximum.






Figure 5.1: The POMF algorithm. Drawing after Bargoti *et al.* (2012)

The generic open-loop POMF workflow is illustrated in Figure 5.1. In step (A), the POMF takes in two image-rendered point clouds and transforms them into the complex domain. The word “Shift” indicates that the quadrants are flipped so that the low frequencies of the spectrum are in the centre of the 2D Fourier transformed elevation image. Step (B) calculates the phase of the Fourier spectrum. Step (C) cross-correlates the phase spectra (CPS) of the first phase quadrant. Step (D) transfers the phase correlation peak(s) from the complex domain back into real realm. Step (E) computes the translational offset(s) from the location of the maximum correlation peak(s). In case the point cloud images show strong changes in magnitude, the CPS will result in one dominant peak that represents the translation offset. Optionally, to optimize results, a sub-pixel refinement of the exact location of the CPS peak of interest (Wisetphanichkij and Dejhan, 2005) can be conducted.


Figure 5.2 below illustrates how steps (A) to (E), outlined in Figure 5.1, look like when the POMF is applied to the simulated X5d4 DTM:




Overlay of original data and partially overlapping data and respective image center points (pix)




Input 1: Zero-padded dataset (pix)




Input 2: Zero-padded dataset translated by 400 pixels




(a1) – Fourier spectrum (cycles/pix), non-logarithmic display




(a2) - Fourier spectrum (cycles/pix), non-logarithmic display




(b1a) - Fourier magnitude spectrum, logarithmic display




(b2a) – Fourier magnitude spectrum, logarithmic display




(b1b) - Fourier magnitude spectrum, logarithmic, top view



(b2b) – Fourier magnitude spectrum, logarithmic, top view



(c1) - Phase spectrum, logarithmic



(c2) - Phase spectrum, logarithmic




Figure 5.2: Graphical illustration of the POMF using translated X5d4 DTM

Figure 5.2 exemplifies each step of the POMF co-registration process. In non-logarithmic display the phase spectra showed a very steep peak (a1, a2). In logarithmic display effects of the angular offset became visible (b1, b2). In the phase spectra the difference in signature seems lost to the human eye (c1, c2). Since the amplitude spectra showed a clearly different signature, the cross-correlation resulted in one dominant CPS peak (d). This peak clearly defined the translational offset ($\Delta i, \Delta j$).

As will be shown in Chapter 5.1.2, the POMF can also be used to recover the rotation and scale offsets of two images.

5.1.1.4 Recovery of the Cartesian Translation using POMF

Since its postulation by Horner and Gianino (1984), many variations of the POMF have been developed that focus on increasing its accuracy and/or robustness. A POMF for recovery of rotation and scale, for example, requires an additional quadrant check and a sophisticated filter strategy (see Chapter 5.1.2.2 and 6.1.4 for details). As shown by Guizar-Sicairos *et al.* (2008), a lot of research has been done on how to recover the Cartesian translations with the POMF accurately, and this problem can be regarded as solved.

In this research, the first goal was to examine how well a generic state-of-the-art POMF performs in recovering the Cartesian shifts between the coarse height maps. MATLAB's Image Processing toolbox (MathWorks, 2020) implemented the POMF in their "imregcorr" phase filter. MathWorks referred their implementation to the work by Reddy and Chatterji (1996) which will be discussed in detail in Section 5.1.2. When the 'window' was set to true, "imregcorr" used a Blackman window (see APPENDIX E) to improve the phase correlation results.

The second goal of this study was to learn what parameter settings are important for successful recovery of the Cartesian translations. This insight may help improve the co-registration results of the FMT (5.1.2) and AFMT (6.1), which execute the POMF many times internally.

Before the POMF was put to the test, the following propositions were made:


- (1) Since the real-world and simulated data show similar height differences, the POMF may show a similar performance. As a result, the datasets may share a certain degree of the overlap at which a successful autocorrelation can occur. If this threshold is similar for all test data, it could be used to characterize the success rate of the POMF and forecast its performance for other data.
- (2) Since the datasets show similar surface characteristics, they should show similar global entropy. The POMF results may reveal a certain threshold for the global entropy at which successful co-registration with the extremely sparse point clouds can occur.

How these expectations match up with the co-registration results, will be explored next.

5.1.1.4.1 Oshawa Data

The real-world and simulated data were processed using POMF in the auto- and cross-correlation mode with the varying translation offsets in East and South. The goal of the autocorrelation tests was to learn how important the similarity of the overlapping portion of the elevation profiles is to the success rate of the co-registration. At each time, the target point cloud was shifted and the entropy of the remaining data was calculated, too, to see if there is a correlation between the overlap, entropy, and success of the POMF.

Figure 5.3 shows results from the Oshawa laser- and tie point clouds through the autocorrelation with their respective translated counterparts.






Figure 5.3: Autocorrelation of Oshawa elevation maps (pix), translated South

Figure 5.3(a0) and (b0) plots the native datasets, and the dotted red lines in (a0) indicate how the tie points situated relative to the laser points, which have a larger footprint than the tie points. In this experiment, the height maps were shifted to the South, successively, by 10 pixels each time. The plots (a1) and (b1) show the footprints of the Oshawa laser points and tie points, respectively, for the first shift. As described in Chapter 3.0, one pixel in the laser- and tie point map here covered 40x40m ground surface.


In Figure 5.3(a2) and (b2), the blue stars indicate the imposed shifts, and the green triangles present the offset that the POMF estimated. Ideally, the green and blue icons should coincide. However, with the increase of the shift, the area of overlap is getting reduced so that autocorrelation became more and more challenging. Eventually, the decrease in common area would reach a point beyond which the POMF was unable to recover the shift correctly. This threshold is indicated in Figure 5.3(a2) and Figure 5.3(b2) where the where the green and blue curves go divergent. The disconnect occurred with around 55pixel shift for the laser points, and 40pixel shift for the tie points, i.e., in both cases the translation reached around 50% of the respective footprint in North-South direction. In case of the tie points, the incorrectly estimated shift formed a tangent aside the induced shifts.

The plots (a3) and (b3) in Figure 5.3 present the relative decrease in size of the shifted data in green, which (since the data has almost quadratic footprint) was almost identical with the decrease in overlap. The red curves in Figure 5.3 (a3) and (b3) show the decrease in the entropy of the non-

padded translated data. The gradual decrease of the entropy was caused by the gradual loss of the South-East area of the laser dataset with the increase of the shift. This area shows the strong height difference and thus a lot of information content. On the contrary, the entropy of the translated padded data, plotted in blue, decreased significantly each time as the dataset were shifted, as the areas where the data lost footprint were re-filled with 0-valued pixels – obviously, these 0-valued areas produced an entropy of zero. Since the Oshawa height maps have nearly quadratic shape, the decrease in the footprint of the remaining data followed a similar trend.

The autocorrelation started to fail once the image size (green curve) reduced to almost 50% and the entropy of the padded data (blue curve) reduced to roughly 60%. Neither the entropy curve nor the footprint curve showed irregularities or other signs that could serve as an indicator that the co-registration failed.

This experiment was repeated whilst the data was shifted in West-East direction. Figure 5.4 presents the results.






Figure 5.4: Autocorrelation of Oshawa elevation maps (pix), translated East (left)

In Figure 5.4, the left and right plots in each pair, respectively, shows the autocorrelation results where the Oshawa tie points and laser points were shifted to the East for 10 pixels, and correlated with their native, non-shifted counterpart. Figure 5.4 (a0) and (b0) illustrate the effect of a 10pixel column shift for each point cloud. As in Figure 5.3, the blue stars in Figure 5.4 (a1) and (b1) indicate the imposed column shifts and the green triangles show the offset that the phase filter estimated. For both elevation maps, the disconnect of imposed and estimated offset happened as around a shift of 45 pixels was reached, which was roughly 40% of the width of the elevation maps. Thus, compared to the recovery of translations in North-South direction (Figure 5.3), the phase filter here failed roughly 10% earlier.


A close look at the Oshawa height maps reveals that the profiles pronounce more in North-South direction than in West-East direction, and the small area in the south-east corner of the LiDAR map shows the strongest decrease in elevation (see Chapter 3.0 for details). Unlike the gradual steady decline in entropy that occurred as the LiDAR data was translated to the South (Figure 5.3), the entropy dipped down 5% within the first 10pixel shift in East direction (red curve) in Figure 5.4 (a3) and (b3).

Surprisingly, the success of the POMF, which started to fail after roughly 35 pixels later, was not negatively impacted by this rapid loss in entropy. Rather, as shown in Figure 5.4 (a3) and (b3), did the POMF fail at a point where the area of overlap had reduced to almost 60% (green curve), which, - since the height maps have nearly quadratic shape -, aligned with roughly 40% width of the maps. However, the phase filter failed when the relative entropy of the padded maps reached

75% (blue curve), i.e., about 10% earlier than when the data was shifted South (Figure 5.3). As before, the incorrectly shift estimates formed a tangent alongside the induced shifts (displayed in Figure 5.4 (a2) and (b2)).

As can be observed from the plot in Figure 5.4 (b2), the entropy of the tie points increased slightly as the data was shifted 5 to 15 pixels to the East. Since the global entropy considers both the number of abnormalities in the elevation profile and the size of the elevation profile, and since the change of the information content in the East end of the Oshawa tie point cloud shows only small gradual change, the remaining non-padded profile contained relatively more information content once the first few data columns in the East end were cut off.

A similar experiment was repeated by shifting the data to the South and East. The results are given in Figure 5.5.






Figure 5.5: Autocorrelation of Oshawa elevation maps (pix), translated South and East

In Figure 5.5, the plots on the left and right show the autocorrelation results where the Oshawa tie- and laser points, respectively, were shifted South-East for 10 pixels each time, and correlated with their native, non-shifted counterpart. The first two plots (a0) and (b0) in Figure 5.5 illustrate the effect of a 10pixel (column) shift for each point cloud. As before, the blue stars in Figure 5.5 (a1) and (b1) indicate the imposed shifts, while the green triangles in Figure 5.5 (a1) and (b1) show the respective shifts estimated by the POMF.

The first deviation between the imposed and estimated shift occurred after the 30pixel shift had been reached. Thus, for the laser- and tie points, this disconnect happened when the translation reached about 25% and 30% of their footprints. The decrease in size of the shifted data, plotted in green in Figure 5.5 (a2) and (b2), and the decrease in entropy of the padded data, plotted in blue in Figure 5.5 (a2) and (b2), followed a similar trend. Here, the autocorrelations failed at the footprints of 60% with the laser points and 50% with the tie points, respectively, so did at an entropy of the padded image (blue curve) of about 70% with the laser points and 60% with the tie points.

The entropy of the non-padded data, plotted in red, changed marginally when the data was translated, and the dip due to the loss of the South-East corner showed up with the smaller magnitude than in the previous experiment. This confirmed again that the relative information content in the data with the smaller footprint did not differ significantly from the information content in the data with the larger footprint.

Next, the three experiments above were repeated whilst the height maps were cross-correlated. The results are given in Figure 5.6, Figure 5.7, and Figure 5.8.





Figure 5.6: Cross-correlation of Oshawa elevation maps (pix), translated South (down)

In Figure 5.6, the plots on the right and the plots on the left show the results where the Oshawa tie points and laser points, respectively, were shifted to the South incrementally by 10 pixels each time, and cross-correlated with their respective non-shifted laser points and tie points, respectively. The thumbnails of the datasets illustrate which data remained stationary and which data was shifted

during the cross-correlation. As before, the blue stars in the plots (a1) and (b1) (Figure 5.6) indicate the imposed shifts, and the green triangles in the same plots show the respective shifts estimated by the phase filter. Clearly, the cross-correlation failed completely in both cases, and the estimated offsets from the phase shift settled on parallels to the induced shifts.

Since the laser point cloud from the site in Oshawa is slightly wider and covers roughly 1/3 more ground in the North-South direction than the tie point cloud from the same site, the decrease in size of the shifted data, plotted in green in (a2) and (b2) (Figure 5.6), started at an entropy offset of approximately 1.6. Since the laser points covered the areas with the stronger elevation contrast, their entropy turned out approximately 20% bigger than that of the tie points. This caused the entropy of the padded data, plotted in blue, and the entropy of the non-padded data, plotted in red, to start at an offset. Both curves decreased only marginally with the increase of the shift, which confirmed that the relative information content remained almost unaffected by the translations.

This cross-correlation experiment was repeated whilst shifting the moving images to the east instead of the south. The results are plotted in Figure 5.7.






Figure 5.7: Cross-correlation of Oshawa elevation maps (pix), translated East

In Figure 5.7, the plots on the right and the plots on the left show the results where the Oshawa tie points and laser points, respectively, were shifted to the South incrementally for 10 pixels each time, and cross-correlated with the respective non-shifted laser points, and tie points, respectively. Obviously, the cross-correlation failed completely in both cases, and the estimated offset from the phase shift settled on parallels to the sequence of induced offsets.

As before, the decrease in size of the shifted data, plotted in green in (a2) and (b2), started with an offset due to the different size of the datasets. The decrease in entropy of the padded data, plotted in blue, also showed the same behavior as before. As in Figure 5.6, the entropy of the non-padded data, plotted in red in (b2), dipped again in the beginning due to the loss of the South-East corner through the shifts. Even though the entropy of the non-padded data increased marginally at a shift of 5-15 pixels (red curve in a2), before it started to decrease, the relative information content remained essentially unaffected by the translations.

Figure 5.8, below, illustrates results where this experiment was repeated while additionally shifting the moving elevation maps South.




Figure 5.8: Cross-correlation of Oshawa elevation maps (pix), translated South-East

As in the preceding figures, the thumbnails in Figure 5.8 illustrate which data remained stationary and which data shifted during cross-correlation. Again, the blue stars in cells (a1) and (b1) indicate the imposed shifts, and the green triangles show the respective shifts estimated by the POMF. As before, the cross-correlation failed completely, and the estimated offset from the phase shift settled on parallels to the induced offsets. The graphs in (a2) and (b2) followed the same trends as in Figure 5.7.

In summary, the POMF showed promising results when the data was auto-correlated and shifted only in one direction, yet the POMF failed entirely when the data was cross-correlated. The change in entropy turned out not to be an indicator for when a successful correlation may occur, which disproved the second proposition introduced in introduction section 5.1.1.4. Thus, the entropy plots were left out from the ensuing sections and added to APPENDIX H.

5.1.1.4.2 Keyworth-Nottingham Data

The experiments described in 5.1.1.4.1 were repeated with the KN data. Since the Oshawa and KN data have a similar height range and footprint (the Oshawa laser point cloud covers 117-by-138 pixels, the Oshawa tie point cloud covers 81-by-115 pixels, the KN laser point cloud covers 59-by-71, the KN tie point cloud covers 36-by-35 pixels), similar co-registration results were expected. As concluded in 5.1.1.4.1, the plotted translation estimates are most important for the interpretation of the co-registration results. They are presented in Figure 5.9.




Figure 5.9: Autocorrelation of KN elevation maps (pix) - expected vs recovered translations

In Figure 5.9, the plots on the right and the plots on the left show the results where the KN tie points and laser points, respectively were shifted incrementally for 10 pixels each time, and autocorrelated with the respective non-shifted laser points, and tie points, respectively. The thumbnails show which data remained unshifted and which data shifted during the correlation. The datasets in (a1) and (b1) (Figure 5.9) illustrate an example where the tie- and laser points were shifted South, by 10 pixels, and plotted against their respective original footprint. The blue stars in the plots (a1) and (b1) indicate the imposed shifts, and the green triangles in the same plots show the respective shifts estimated by the phase filter.

For the laser data, the phase filter was able to recover the translations almost twice as far out than for the tie points (27pix vs 15pix). However, compared to the laser points, the tie points cover only about 2/3 of the ground in North-South direction and they extended only half as far in West-East direction. As could already be observed from the translation experiments with the Oshawa data, once the POMF failed, its translation estimates settled on parallels to the induced shifts.

As displayed in the plots (a2) and (b2) (Figure 5.9), in both cases the phase filter started to fail when the size of the shifted data had reached roughly 40% (30pix/71pix and 13pix/35pix). When the data was shifted South-East, plots (a3) and (b3), the disconnect happened when the translation reached around 30% of their footprint. As shown in APPENDIX H, the decrease in entropy and decrease in size of the shifted data followed a similar trend and the entropy values had reduced approx. 10% less than the respective footprint values when the co-registration failed. These three translation experiments were repeated whilst the laser- and tie point maps were cross-correlated. The results are plotted in Figure 5.10.




Figure 5.10: Cross-correlation of KN elevation maps – expected vs recovered translations

In Figure 5.10, the plots on the right and the plots on the left show the results where the KN tie points and laser points, respectively, were shifted incrementally for 10 pixels each time, and cross-correlated with their respective non-shifted laser points, and tie points, respectively. The thumbnails of the datasets illustrate which data remained unshifted and which data was shifted during the cross-correlation. As before with the Oshawa data, the cross-correlation failed completely, and the phase shift estimated offsets settled on parallels to the induced shifts.

In summary, the KN data followed the same trend as the Oshawa data and the POMF showed promising co-registration results when the data was auto-correlated and shifted only in one direction, yet the POMF failed again entirely when the data was cross-correlated. Overall, the KN data showed less correlation success than the Oshawa data, most likely because the KN topography is slightly flatter. As before, the respective entropy plots can be found in APPENDIX H.

5.1.1.4.3 Simulated Data

The experiments described in 5.1.1.4.1 and 5.1.1.4.2 were repeated with the simulated data. Since the simulated elevation profiles have stronger undulation, a larger footprint (1864-by-1864 pixels), and identical yet much higher point density, better co-registration results were expected. The results of the translation estimates are plotted in Figure 5.9.




Figure 5.11: Autocorrelation of X5d5 and X5d4 (pix) – expected vs recovered translations

In Figure 5.11, the plots on the right, and the plots on the left shows the results where the X5d4 DEM and X5d5 DEM, respectively, were shifted incrementally for 100 pixels each time, and autocorrelated with the respective non-shifted X5d4 DEM, and X5d5 DEM, respectively. The thumbnails show how the data was shifted during the correlation. The datasets in plot (a1) and plot (b1) (Figure 5.11) illustrate an example where the DEMs were shifted to the South, by a few 100 pixels, and plotted against their respective original footprint. The blue stars in the plots indicate the imposed shifts, and the green triangles in the same plots show the respective estimated shifts from the phase filter.

For the X5d5 DSM, the phase filter was able to recover the translations further out than for the X5d4 DTM (1600pix vs 1300pix when the data was translated to the South and 1600pix vs 1200pix when the height maps were translated to the East). In both cases, the phase filter started to fail when the size of the shifted data had reached roughly 85% (65%). When the data was shifted to the South-East, plot (A3) and plot (B3) (Figure 5.11), respectively, the disconnect happened when the translation reached around 50% of their respective footprint (700pixels vs 1000pixels).

These translation experiments were repeated whilst the DSM and DTM were cross-correlated. The results are plotted in Figure 5.10.




Figure 5.12: Cross-correlation of X5d4 and X5d5 (pix) – expected vs recovered translations

In Figure 5.12, the plots on the right and the plots on the left show the results where the X5d4 DTM and X5D5 DSM, respectively, were shifted incrementally for 100 pixels each time, and cross-correlated with their respective non-shifted counterpart. The thumbnails of the datasets illustrate which data remained non-shifted and which data was shifted during the cross-correlation. As before with the Oshawa and KN data, the cross-correlation failed completely, and the estimated offsets from the phase shift settled on parallels to the induced shifts.

In summary, the co-registration with the POMF of the simulated DEMs followed the same trend as the co-registration with the POMF of the Oshawa and KN data. It showed promising results when the data was auto-correlated and shifted only in one direction, yet the POMF failed again entirely when the data was cross-correlated/. Overall, the POMF showed more correlation success with the simulated data than with the real-world data because the simulated data has a more pronounced height profile and a higher point density than the Oshawa and KN topographies. This finding disproves the first proposition formulated in section 5.1.1.4, which anticipated that there may be a common threshold for the data overlap at which co-registration can occur successfully.

5.1.1.5 Summary

In section 5.1.1 and its subsections, the co-registration performance of the POMF, as provided by MATLAB, was investigated using the Oshawa, KN, and simulated height maps in autocorrelation and cross-correlation mode. Understanding the performance of the POMF is important because the POMF is part of the FMT, AFMT, and QAFMT. While the recovery of the cartesian shifts with the POMF is a well-known process, our focus was on how well it could perform with the low-resolution airborne point clouds with highly irregular point distribution. In these experiments, the “imregcorr” phase filter in MATLAB’ image processing toolbox (MathWorks, 2019) was used for the POMF. This POMF version was able to recover the translation offsets up to 85% when the height maps were self-registered, however it failed completely to cross-correlate the DEMs successfully.

The results revealed that there is neither a common threshold for overlap at which successful auto-correlation can occur (first proposition, section 5.1.1.4), nor do they share a common entropy at which a successful co-registration can occur (second proposition, section 5.1.1.4).

In the above experiments, all three datasets showed similar POMF output, and their graphs followed similar trends. Whenever the data were put in the cross-correlation mode, the estimated offsets by MATLAB settled primarily on parallels to the induced shifts. This observation leads to the (third) proposition, that the MATLAB POMF settled on the optima of the second or higher order. Hence, it may not always be the tallest CPS peak that represents the best solution, but rather a CPS peak that has the smaller magnitude. Consequently, as pointed out by Wilmer (2012), more than one peak should be considered in the search for the optimal translation offset. Figure 5.13 illustrates how the POMF can be expanded to evaluate more than one CPS peak.




Figure 5.13: POMF that considers N CPS peaks in the search for the correct translation offset

As illustrated in Figure 5.13, in correspondence to Wilmer (2012), more than one CPS peak can be evaluated in the search for the translation offset. Here, the translation offsets associated with N peaks are assessed in Step (D) and Step (E). Thus, the translation offset for each of the N peaks is applied to the floating image, which is then re-ingested into the POMF. Then, the translation offset that leads to the highest CPS peak, highlighted in Step (F), is considered as the final translation estimate.

This expanded POMF model requires to decide how many CPS peaks should be considered, whether a window filter needs to be applied in space domain (and if so, which one) or a high pass filter (HF) in frequency domain (and if so, what form should it have). This expanded POMF will be used in Chapter 5.1.2.2 and 6.1.4 where one needs to evaluate multiple peaks for the recovery of the translation offsets that represent angular and scalar unknowns.

5.1.2 The Fourier Mellin Transform

In section 5.1.1, the POMF was introduced as a DFT-based method to efficiently retrieve the translation offsets between images. This section presents the mathematical extension of the POMF for the recovery of angular and scale offsets – the Fourier Mellin Transform (FMT) (Horner and Gianino, 1984). The FMT will be introduced along with their characteristics and limitations that affect the success rate of the co-registration of laser- and photogrammetric point clouds.

5.1.2.1 Definition

Normally, the FT output of a two-dimensional signal is not invariant to rotation and scale. A function's rotation in two dimensions by an angle φ will have its FT rotated by the same angle (Lyons, 1997; Le Maistre Pratt, 2000). In the ensuing paragraphs, it will be shown how log- polar sampling can be used to achieve both the rotational and scaling invariance, which, when applied inside the POMF, eventually, extends to the FMT.

The one-dimensional continuous FT is defined as

$$\mathcal{F}\{f(x)\} = F(\omega) = \int_{-\infty}^{\infty} f(x)e^{-i\omega x} dx \quad (5.7)$$

where $\omega = 2\pi u$. According to (Bracewell, 1986), the 1D Mellin Transform (MT) formulates as

$$M(s) = \int_0^{\infty} f(x)x^{s-1}dx \quad (5.8)$$

By replacing the variable x with e^{ξ} and s with $-j\omega$, the MT of $f(x)$ on the imaginary axis is the Fourier Transform $f(e^{\xi})$

$$M(j\omega) = \int_{-\infty}^{\infty} f(e^{\xi})e^{(-j\omega\xi)}d\xi \quad (5.9)$$

Hence, the MT can be realized by logarithmically scaling the coordinates (e^{ξ}) of the input function and Fourier transforming the resulting scale function.

The two-dimensional MT of a function $f(x, y)$ is defined along the imaginary axis by (Casasent and Psaltis, 1976)

$$M(ju, jv) = \int_0^{\infty} \int_0^{\infty} f(x, y)x^{-ju-1}y^{-jv-1}dx dy \quad (5.10)$$

In analogy to the substitution of the variables in (5.9), replacing x with e^{ξ} , where $\xi = \log(r)$, and y with φ , and dividing by j , the 2D MT in *polar coordinates* finds (Casasent and Psaltis, 1976):

$$M_f(u, v) = \frac{1}{2\pi} \int_0^{\infty} \int_0^{2\pi} f(r, \varphi)r^{-iu}e^{-iv\varphi}d\varphi \frac{dr}{r} \quad (5.11)$$

for $\forall(u, v) \in \mathbb{Z} \times \mathbb{R}$.

In (5.11), u and v are the MT parameters affecting r and φ , respectively. As with the 1D example in (5.9) and (5.10), by comparing this expression with the formulation of the continuous FFT, the FMT reveals as a FT with the logarithmically scaled input $f(e^{\xi}, \varphi)$ with $\xi = \log(r)$:

$$M_f(u, v) = \frac{1}{2\pi} \int_0^\infty \int_0^{2\pi} f(e^\xi, \varphi) r^{-iu\xi} e^{-iv\varphi} d\varphi d\xi \quad (5.12)$$

As can be seen from (5.12), the FMT is defined as FT applied to the observations in the log-polar base. This means that the FMT can be accomplished by performing a 2D FFT after a polar conversion and logarithmic scaling.

Under the consideration of two images F and G mentioned earlier, $F(r, \varphi)$ and $G(s * r, \theta + \varphi)$, varying only in the scale s and rotation φ , the FMTs are related to them as follows:

$$M_f(u, v) = s^{-iu} e^{iv\varphi} M_g(u, v) \quad (5.13)$$

Thus, the angle φ appears as a phase change (Euler' cofactor $e^{iv\varphi}$) and the scaling as a magnitude change (appearing as cofactor s^{-iu}). To make it a pure phase change, a derivative is built and r is re-mapped to $r = e^p$. Thus, by remapping the original Fourier Transform into the log-polar coordinates, the images F and G differ only in translation. The horizontal and vertical axes of the log-polar images now represent the rotation and scaling, respectively.

5.1.2.2 Co-registration using FMT

5.1.2.2.1 The FMT Algorithm

There is more than one way in which a FMT can be carried out. The most popular method is to apply the log-polar transform (the Mellin part of the FMT) to the Fourier magnitudes. In the following sections, this method will be referred to as the standard FMT, or just FMT, which is outlined in Figure 5.14.




Figure 5.14: Computational steps of the standard FMT of image spectra. Illustration after Bargoti *et al.* (2012)

As illustrated in Figure 5.14, the standard FMT algorithm comprises six steps:

1. Input are two image-rendered point clouds
2. Step (A) conducts an FFT on the two images, transforms them into the complex domain. The quadrants are flipped so that the low frequencies of the spectrum are in the centre.

Optionally, a blur filter (Kazik and Göktoğan, 2011) such as a Hamming window, Bartlett, Blackman, Boxcar, Hann, MATLAB’s “Chebwin” Chebyshev, or similar window filter (Wilmer, 2012; MathWorks, 2019) can be applied to the input images to reduce the high-frequency effects of tiling (Bargoti *et al.*, 2012).

3. Step (B): The amplitude of the Fourier spectra is generated for the following reasons:
 - The POMF can only ingest signals in real domain

- By operating on the amplitude spectrum, the translational differences are circumvented (since the magnitude spectrum of an image and its translated counterpart are identical) and only their phase spectra are different.

By doing so however, the phase information of the Fourier spectra gets lost where *only the magnitude information* of the Fourier spectra $\|F_1\|$ and $\|F_2\|$ (*and not the amplitude information*) are considered further.

4. Step (C): The spectral magnitude is mapped into a log-polar base. This log-polar spectral (magnitude) image is also referred to as FMT spectrum (Wolberg and Zokai, 2000). In log-polar base, rotation and scale parameters are de-coupled and associated with the x- and y-axis.

By converting the Cartesian Fourier magnitudes into log-polar base, the FMT exploits the Fourier shift theorem a second time and allows the POMF to recover XY displacements, which in this case are the rotational and the scalar offset of the FMT spectra. Indeed, the only reason for mapping the image spectra rather than the images themselves into log-polar base is that it allows for estimating linear displacements (since the Fourier spectra are freed from translational offsets) in a second POMF thereafter (step E), in addition to rotational and scaling offsets in the first POMF.

Optionally, a High-pass Filter (HF) is applied to the amplitude spectra prior to sampling it into log-polar base to remove low frequency components before the second POMF. As suggested by Kazik and Göktoğan (2011), who refer to Reddy and Chatterji (1996), the HF has the transfer function

$$H(\xi, \eta) = (1.0 - X(\xi, \eta)) * (2.0 - X(\xi, \eta)) \quad (5.14)$$

where

$$X(\xi, \eta) = [\cos(\pi\xi) \cos(\pi\eta)] \quad (5.15)$$

with $\xi \leq 0.5$ and $\eta \leq 0.5$.

The plots C1,1 and C2,1 in Figure 5.15 illustrate the shape of this HF as it was used during FMT testing. Here $\xi = \eta = [-0.5 : incr : 0.5]$ with $incr = 1/(\text{number of pixels})$.

- Step (D): The log-polar representations of the Fourier magnitudes are free of the translational offsets by taking the FFT a second time inside the POMF.

Here, the only remaining offsets to determine are rotation and scaling, which themselves appear now as translations. One can argue, that since the standard FMT applies the FT twice to the input data, – the POMF inside the standard FMT observes $FT(\|FT(image)\|)$ –, the product of (step D) can be considered the derivative of second order in frequency domain (Gruber, 2017). Sellami and Ghorbel (2012) found that for this indirect phase correlation (the phase of the Fourier magnitude is correlated and not the phase of the images), the relationship results

$$\|F_{2D}(f \circ (\rho, \varphi))\| = \|(F_{2D}(f))(\rho^{-2}, \varphi)\| \quad (5.16)$$

Thus, if the image f has been rotated with φ and scaled with a factor ρ , then the classical FT (F_{2D}) of this image is rotated with the same angle φ and scaled with ρ^{-2} .

Since the spectral density plot and the magnitude plot are symmetric and repetitive with a phase of 180 degrees, the location of the maximum correlation peak refers to two potential rotation candidates: φ and $\varphi + 180^\circ$. To recover from this ambiguity, the original search image is scaled and back-rotated one time by φ and one time by $\varphi + 180^\circ$ (Chen *et al.*, 1994). Since the POMF outputs two estimates for the angular offset, the retrieval of the translation offset occurs twice – one time for the image pair where the moving image was rotated by φ and one time where the moving image was rotated by $\Delta\varphi+180$.

- Step (E): The rotation angles and scale are applied to the search image and the POMF is re-applied to both image pairs, in native non-frequency format, (steps (E1) and (E2) – reference image plus the search image rotated by φ and the reference image plus the search image rotated by $\varphi + 180$, respectively). In this process, in step (E) the rotation angle $\Delta\varphi$

and scale Δs can be extracted from the location of the peak of the inverse FT of the CPS. According to Kazik and Göktoğan (2011) the conversion formula finds

$$\Delta\varphi = \frac{360y}{s} \quad (5.17)$$

$$\Delta s = \rho^{2x} \quad (5.18)$$

where


$\Delta\varphi$	the angular offset in degrees
Δs	the scale
x, y	the log-polar location of the peak
s	the length of the image on the angular axis
ρ	the base according to the length of the image on the polar axis


The translation parameters $\Delta x, \Delta y$ are returned directly by the POMF, without using the conversion formula.

- Step (F): The image pair that results in the highest correlation peak inside the second POMF, i.e., after correction of scale and rotation, is then chosen as final candidate. Since the translation parameters $\Delta x, \Delta y$ are returned directly by the POMF from the very same correlation peak, all four parameters $\Delta x, \Delta y, \Delta\varphi, \Delta s$ are determined by this correlation peak. Thus, in case of a single, clear, isolated peak, the search is over.


In applications, where subtle changes in respective image phase are expected, many of the correlation peaks found in step (D) may be considered in the search for the best solution. In that case, the calculations that tail the ^(*) in Figure 5.14 have to be executed for each correlation peak: For each of the N correlation peaks in step (D), where $N \in \mathbb{N}$ is a user specified number as for example $N = 20$, the search image is first back-rotated by the respective two angles, then the POMF is re-applied to each image pair (step E), and the image pair which results in the highest CPS peak, overall, is then considered to represent the final set of $\Delta x, \Delta y, \Delta\varphi, \Delta s$.

Figure 5.15 illustrates the results from Steps (A) to (D). In this instance the FMT was used to auto-register the X5d4 DTM (introduced in section 3.4) with its 30deg rotated counterpart. By complying with Wilmer's version of the FMT (Wilmer, 2012), the window filter was applied to the log-polar Fourier magnitudes (step B) instead to the input imagery:







(a1,2): Real part of amplitude spectrum, native PCL, logarithmic display (cycles/pix)




(a2,2): Real part of amplitude spectrum, 30deg rotated PCL, logarithmic display (cycles/pix)




(b1,1): Magnitude spectrum, native PCL, non-logarithmic display (cycles/pix)




(b2,1): Magnitude spectrum, 30deg rotated PCL, non-logarithmic display (cycles/pix)




(b1,2): Magnitude spectrum, native PCL, logarithmic display, side view




(b2,2): Magnitude spectrum, 30deg rotated PCL, logarithmic display, side view







(c1,1): High-pass filter, side view




(c2,1): High-pass filter, top view




(c1,2): HF'ed magnitude spectrum, native PCL, side view




(c2,2): HF'ed magnitude spectrum, 30deg rotated PCL, side view



(c1,3): HF'ed magnitude spectrum, native PCL, quadrants flipped, side view



(c2,3): HF'ed magnitude spectrum, 30deg rotated PCL, quadrants flipped, side view






Figure 5.15: Graphical illustration of the standard FMT workflow using rotated X5d4 DTM (pix)

Figure 5.15 illustrates each step of the standard FMT with the X5d4 DTM. In non-logarithmic display (the plots (a1,1) and (a2,1)) the phase spectra showed a very steep peak. In logarithmic display (the plots (a2,1) and (a2,2)), the effects of the angular offset became visible.

The non-logarithmic magnitude spectrum (the plots (b1,1) and (b2,2)) showed a singular steep peak, also. As before, in logarithmic base, the effects of rotation became visible in the phase spectra of the Fourier magnitudes (b1,2 and b2,2) and, like the rotation angle, the "+" artifact appears at a 30deg angle. This signature became even more obvious in (b1,3 and b2,3). In log-polar base, the lobes of the "+" showed the 30deg rotation in form of a shift (b1,4 and b1,5; b2,4 and b2,5).

The HF (c1,1 and c2,1) suppressed the lower magnitudes at the centre of symmetry. The HF is applied to the magnitude spectrum before its quadrants are flipped and the quadrants are flipped back after HF: (c1,2 and c2,2) show the magnitude plot where the high frequencies are located at the outskirts and (c1,3 and c1,4; c2,3 and c2,4) show the HF'ed version where the high frequencies are in the centre. As with (b1,4 and b1,5; b2,4 and b2,5), in log-polar base, the lobes of the "+" showed the 30deg rotation in form of a shift (c1,5 and c1,6; c2,5 and c2,6). After the FT had been applied a second time, the triple "+" artifact became obvious in the Mellin spectra (c1,7 and c1,8; c2,7 and c2,8).

Cell (d) displays the CPS graph of the cross-correlated Mellin spectra with four dominant peaks. In this example, it was the second highest peak that corresponded to the correct 30deg angular offset.

5.1.2.2.2 Properties

The benefit of the standard FMT over the AFTM (which will be introduced in Chapter 6.1) is that it can recover directional offsets in addition to rotation and scale offsets. While this FMT has become a standard for image registration in the computer vision domain, it does have downsides, too. Some drawbacks are caused by the concatenation of DFT which invokes loss of information through taking the (first and) second derivative, whilst other drawbacks are inherited from the FT and from log-polar mapping. These drawbacks and their respective workarounds are:

- (1) *Resampling error*: Also referred to as aliasing. Logarithmic and polar sampling are an essential element of the FMT. The log-polar resampling causes a redistribution of elevation

information in such a way that a lot of the original signal content is re-located near the centre point of the log-polar axis. In a mathematically continuous case, the point density becomes infinity towards the centre (Derrode and Ghorbel, 2001). In digital applications, the logarithmic base is typically chosen such that the number of samples along the axis is the same as the number of pixels of the image (Kazik and Göktoğan, 2011). A strict one-to-one assignment of the Cartesian image pixels to the log-polar image pixels can however cause under-sampling in the outskirts of the log-polar image which can cause severe interference with the resulting FMT transformation parameters (Mittal and Sharma, 2008).

Le Maistre Pratt (2000), referring to Thornton and Sangwine (1998), discussed a spatially variant filter to remove possible aliasing due to the under-sampling inherent in log-polar transformation (Nyquist theorem). Whilst in regular log-polar transformation the outermost area is the most under-sampled, the method of Thornton's and Sangwine (1998) increases the sampling interval such that the image is oversampled approximately three-fold and later interpolated (in the oversampled areas) and down-sampled to re-create an image with original resolution. The result is an image in log-polar form without the under-sampling artifacts and loss of information. Similar numerical approximation strategies have been suggested by Kunis (2006).

Keller *et al.* (2005) proposed the Pseudo-Polar FMT (PPFMT) that avoids aliasing by applying a special, quasi-log-polar sampling structure. Keller claimed numerical superiority of the PPFMT over the Non-equispaced Fast Fourier Transform (NFFT) implementation presented by Keiner (Keiner *et al.*, 2009) who built on the work of (Kunis, 2006). Similarly, Guo (Guo *et al.*, 2005), referring to Gho's work (Gho, 1985), proposed applying the DFT on the input images (this idea will be revisited in Chapter 6.1) and approximating the log-polar conversion by quasi-polar summation of the Fourier magnitudes in a Cartesian grid. Aside from increased numerical efficiency, this method seemed to decrease sensitivity to white noise in the images.

- (2) *Interpolation error*: [Here, interpolation is defined as resampling whilst maintaining the number of samples]. The loss of information introduced by interpolation may reduce the accuracy of the co-registration. In the FMT workflow, interpolation occurs when (a)

gridding the laser and tie point data, (b) rendering the search image, and (c) log-polar sampling close to the origin of the coordinate axis.

Loss of information through interpolation needs to be evaluated against other characteristics of the interpolation methods: The experiments described in APPENDIX E showed that MATLAB's "natural" gridding function, though lower in order, yielded more realistic-looking surface curvature than bicubic and spline interpolation when coarse elevation data was interpolated.

- (3) *Tiling Error*: The DFT perceives the finite input image as sample of a continuous infinite signal of images. Thus, the input image appears tiled infinitely in the plane, where the top (left) end of an image is bordered directly by the bottom (right) end of the image. This tiling effect is illustrated in Figure 5.16 below.




Figure 5.16: Periodicity of 2D DFT transformation and illustration of discontinuities where the repeated image meets. After Le Maistre Pratt (2000)

The resulting jumps in elevation at the image borders introduce a “+”-shaped artifact in the Fourier domain. There are the following workarounds for reducing this unwanted effect:

- Reduce all pixel values by the minimum value of the image. As illustrated on the left side of Figure 5.17 (a) below, this method tends to produce only small reduction of discrepancies at the image borders.
- Reduce the pixel values by a best fitting plane. This concept is illustrated for one image line in Figure 5.17 (b) below.





Figure 5.17: Decrease of discontinuities at image borders through reduction of minimum value and best fitting plane

Figure 5.17a illustrates how differences in elevation at the image borders cause artificial step signals. Though more efficient than the first method (reduction by the minimum), the second method, displayed in Figure 5.17b, can only reduce but not entirely annul the step effect, since the edges of point cloud images may fluctuate in elevation along the border. Thus, without application of a window filter, the FT will frequently observe image edges, especially if an image is offset by a rotation angle to its partner image.

- Blur the image borders in the pixel domain against the opposite border.
- Similarly, apply a window or weigh function in space domain (Le Maistre Pratt, 2000), or by applying a high pass filter (HF) in frequency domain (Douillard *et al.*, 2013; Wang and Liao, 2013). While the HF reduces the corresponding dominant side lobes in frequency domain, the window filters soften the step effect of the image borders and make them gradually reach a common value (Mather and Koch, 2011). Le Maistre Pratt (2000) for example suggested applying a Bartlett or a Hamming window. As will be shown further below in this section, MathWorks (2019) and Wilmer (2012) offer additional window masks.

(4) *Padding error (“+”-effect)*: Since the FFT and FMT require image input of same size, the elevation profiles can either be cropped or padded to generate input data of the same size. Image padding can create similar artifacts as tiling - padding with zeros is likely to cause abrupt, step-like changes in pixel values at the image-edges, which leads to a “+”-shaped overlay of energy on top of the point cloud’s own phase signal in Fourier domain.

To reduce padding artifacts, the image can be padded in such a way that the last pixels are repeated. Gradually transcending from the last value to the first value (on the adjacent image border) is another padding option that reduces artifacts. However, regardless of what padding function is chosen, some residual step effects will likely remain at the border. If padded with zeros, the magnitude of the “+” artifact becomes stronger with increasing width of the padded border lines. This phenomenon is illustrated in Figure 5.18 below.






Figure 5.18: Zero-padded image border (pix) creating a “+” artifact in the Fourier spectrum

As can be seen from Figure 5.18, expanding the zero-padded areas causes the “+” effect to increase magnitude. The “+” effect distributes along the axis of the phase signal and takes away from the portion of the phase signal that is of interest. In case of 90^0 (180^0 , 270^0) degree rotation, i.e., multiples of $\frac{\pi}{2}$, only one “+” is created as the border lines coincide with the image boundaries in native orientation. Empirical tests showed however that increasing the width of the padding zone does not degrade the registration performance in case there is only one “+” visible.

As indicated in equation (5.16), in case of a 45deg (and 135deg, 225deg, 315deg) point cloud rotation, i.e., multiples of $\frac{\pi}{4}$, the padding areas have borders of equal length and the “+” shows up at 45deg angles, too. In case of rotation angles that are not a multiple of $\frac{\pi}{4}$, as for example 30deg or 80deg, the image rotation yields padding areas with unequal sides. As can be seen from Figure 5.19 below, the unequal length of the border lines materializes as multi- “+” effect.




Figure 5.19: Multi- “+” artifact in FMT phase spectra and corresponding CPS of a 30deg and 80deg rotated X5d4 DTM (cycles/pix)

As Figure 5.19 illustrates, the wide-spread dispersion of energy across multiple “+” lobes distract from the actual phase information of the point clouds, which – as will be discussed in more detail in the ensuing sections – is the reason why the FMT has more difficulty resolving angular offsets that are not a multiple of $\frac{\pi}{4}$. For the 80deg rotated image in Figure 5.19 for example, it was not the first but the second highest peak in the CPS plot that corresponded to an 80deg rotation.

As before with the tiling error, the unwanted “+” effect can be reduced by applying a window or weigh function in space domain and/or by applying a HF in frequency domain

- (5) *Leakage error*: Leakage is caused by discontinuities in Z in the DFT input (McGuire, 1998; Le Maistre Pratt, 2000), which are often an effect of *tiling* and *padding* (see above). They

can however also be caused by abrupt elevation changes in the surface scan, for example when an object clearly stands out from its background. Such jumps in magnitude act like a step function. As a result, the spectrum signature of smaller signals can get lost (McGuire, 1998).

To reduce the leakage from step-like objects in high resolution data, Le Maistre Pratt (2000) and Lyons (2010) suggested softening the scene, for example with a window filter, so that object gradually faces a common value with its background. Since in low-resolution airborne geodetic data the discontinuities in elevation are small and gradual, leakage is mostly only a by-product of tiling and padding.

(6) *Loss of (the input image's) phase information:* As discussed earlier in this section and displayed in Figure 5.14, it is *not the amplitude* of the Fourier spectra, *but the magnitude* of the Fourier spectra that gets mapped into log-polar base inside the FMT. This is because only real numbers can be carried over into log-polar base.

When taking the magnitude, the phase information of the Fourier spectra gets lost. The POMF can therefore only take the phase information of the Fourier magnitude into account. It is however the phase information of the Fourier (amplitude) spectrum that contains crucial information about the location of objects and fine structures of the image scenery. In fact, the phase information carries more information about the image content than the magnitude information (Oppenheim and Lim, 1981).

In summary, the standard FMT inheres six challenging characteristics. The first two, - loss of information through resampling and interpolation -, can be minimized by deploying specially adapted algorithms. Similarly, artifacts from tiling, padding, and leakage can be reduced by applying a window or weigh function in space domain and/or by applying an HF in frequency domain. However, unless one modifies the standard FMT algorithm itself, there is no mitigation strategy for the loss of phase information that occurs inside the standard FMT. In theory, the loss of phase information should have a strongly degrading effect on the co-registration results. The impact of this detriment and the other above-mentioned drawbacks will be examined next.

5.1.2.2.3 Considerations about the Design Space and Angular- and Scale Offsets

As indicated in Chapter 5.1.2.1 and Figure 5.14, the FMT contains many parameters that can be adjusted. It is important to understand the function of each parameter and be aware of potential cross-correlation, so that the FMT and its co-registration performance can be clearly characterized. An overview of all FMT parameters, arranged in form of a Pugh diagram (Kacker *et al.*, 1991), is shown in Figure 5.20 below.




Figure 5.20: Pugh diagram of the standard FMT

Figure 5.20 gives an overview of the variables that can be adjusted in the standard FMT:

- Primary input to the FMT are height maps in image format, displayed in green on the left. Since FMT co-registration keeps one of the two input images static, as reference, while the other “floating” image is transformed onto the reference image, there are only two possible data combinations per survey site: In case of the real-world data, these are (1) static laser points with floating tie points, or (2) static tie points with floating laser points. During performance evaluation however, two additional data combinations can occur: (3) static laser points with moving laser points, or (4) static tie points with moving tie points.

- The output of the FMT, displayed in purple on the right, are the co-registration parameters. Most of the tests described in the ensuing paragraphs focused on the recovery of rotation and scale only. Therefore, the observation of translation offsets was greyed out.
- Encircled in orange are all FMT-internal parameters that can be adjusted. Displayed in grey are control parameters that were set to a constant value or reduced to a reduced sub-selection in the tests described in the ensuing discussions.
- The group of variables in the blue box describe environmental factors, or noise factors (Kacker *et al.*, 1991), that are typically out of the user's control. In the context of co-registration, these are the unknown co-registration parameters we are trying to recover. As before, the parameter values that were tested less frequently than others in the subsequent experiments were marked in grey.
- The "Symptoms" describe output attributes that can only be observed but cannot be controlled. Examples are robustness, SNR, performance characteristics like convergency behaviour (predicted vs actual, stability, artifacts related to the outcome like how many correct solution candidates stem from a certain amount of CPS peaks), and the cost of computational time.

Related to the symptoms is the recording, evaluation, and feedback of FMT results: From conducting experiments with different data input- and processing settings, the range of correctly solved scale-, translation-, and rotational offsets can be identified and looped back. Wilmer (2012) for example found that his FMT implementation could resolve a maximum scale factor of 5.0 when applied to images that show close-range scenery with people in them.

As Figure 5.20 shows, there are many noise factors and control factors that need to be considered when analyzing the FMT's co-registration performance. This raises the question how these factors should be conditioned to explore the test space, or "design" space, efficiently and effectively. APPENDIX E shows how the design space was reduced from 1156 tests for "full factorial" testing (all test combinations), to less than 50 experiments.

Another consideration for the test design was that in co-registration of airborne geodetic data, there are only two ways how a difference in scale can occur: 1) when the target point cloud is referenced in feet and the search point cloud in meters, the search point cloud appears to be scaled by 0.3048, 2) when the target point cloud is referenced in meters and the search point cloud in feet,

the search point cloud appears to be scaled by 3.2808. Since the point clouds are very coarse in resolution, small scale differences due to map projection can be ignored.

5.1.2.2.4 Numerical Results

The experiments presented in APPENDIX E helped identify the variables that have a major influence on the co-registration results and isolate settings which allow for successful co-registration. Most importantly, it was found that it is sufficient to limit the search for the correct set of co-registration parameters to the 20 tallest CPS peaks. In addition, nine filter masks were introduced, and two test configurations were shown where the application of the Chebyshev mask degraded the co-registration results.

In this chapter, it is examined how well scale and rotation offsets can be recovered from real-world data with the FMT. In addition, the impact of interpolation methods and filter masks is studied in more detail. To keep the test configuration simple, only cardinal angles were chosen as angular offsets and all tests were conducted with real-world data.

The results were summarized in Table 5.2 to Table 5.7 below. To allow for accurate comparison, all scale estimates were rounded to four decimals. The rotation estimates were rounded to one decimal since the real-world datasets provided only few pixels ground coverage. In all tests, the search image was rotated in 90deg increments with respect to the template image and the images were padded to have the same size. The FMT results were evaluated using 20% and 50% thresholds for the deviation of estimated and actual parameters. Scale offsets of 0.3048, 1.0, and 3.2808 were applied. Table 5.1 below illustrates the respective thresholds for each co-registration parameter.

Table 5.1: Thresholds used to evaluate FMT-estimated co-registration parameters of real-world height maps

	Acceptable Error Margin	
	20% deviation	50% deviation
Acceptable deviation from expected scale estimate 3.2808	$ds = \pm 0.7: 2.6 < s < 4.0$	$ds = \pm 1.7: 1.6 < s < 5.0$
Acceptable deviation from expected scale estimate 1.0	$ds = \pm 0.2: 0.8 < s < 1.2$	$ds = \pm 0.5: 0.5 < s < 1.5$
Acceptable deviation from expected scale estimate 0.3048	$ds = \pm 0.06: 0.24 < s < 0.37$	$ds = \pm 0.15: 0.15 < s < 0.46$
Acceptable deviation from expected rotation estimate	$dr = \pm 18\text{deg}$, e.g., for 90deg i.e., $72\text{deg} < r < 108\text{deg}$	$dr = \pm 45\text{deg}$, e.g., for 90deg i.e., $45\text{deg} < r < 135\text{deg}$
Acceptable deviation from expected translation estimates in x and y for Oshawa data (81-by-115 pixels)	$dx = \pm 16\text{pix}$ $dy = \pm 23\text{pix}$	$dx = \pm 40\text{pix}$ $dy = \pm 57\text{pix}$
Acceptable deviation from expected translation estimates in x and y for KN data (51-by-71 pixels)	$dx = \pm 10\text{pix}$ $dy = \pm 14\text{pix}$	$dx = \pm 25\text{pix}$ $dy = \pm 35\text{pix}$

In the subsequent tables, the FMT results were evaluated against the 20% and 50% thresholds. Parameter estimates marked in **green** indicate that the FMT terminated **correctly**, and the estimates deviated less than **20%** from their respective nominal value. Parameter estimates that deviated more than 20% but less than **50%** were considered incorrect yet **acceptable** and were highlighted in **blue**, since they led in the right direction (and could therefore serve as a starting point for fine registration). Estimates with more than 50% deviation from the actual values were considered **wrong** and colour coded in **red**.

Thus, for example, all translation estimates that deviated less than 16pix in x and 21 in y from the expected value of 0pix were considered correct estimates and were marked in green. Translation estimates that deviated more than 20%, but less than 50%, i.e., $16\text{pix} < \text{estimated } dx < 40\text{pix}$ and $21\text{pix} < \text{estimated } dy < 57\text{pix}$ were highlighted in blue. Scale estimates outside that boundary were marked in red. Cases where the image footprint increases slightly through padding for certain rotation angles was accounted for in a conservative manner by maintaining the thresholds that were derived for regularly sized image footprints.

Out of all window filters shown in APPENDIX E, the Chebyshev window showed the strongest damping effect; the Bartlett window displayed low damping power. Since the Bartlett mask seemed to represent a good compromise between applying no filter and the Chebyshev filter, it was chosen in most of the following experiments that examined the effect of window filtering. To get an impression of how good a representative the Bartlett filter is for window masks in general, and how much impact differences in damping power have, one of the test cases was repeated with the Chebyshev filter.

Table 5.2, Table 5.3, and Table 5.4 below illustrate co-registration results for the Oshawa height maps using Wilmer's implementation of the standard FMT (Wilmer, 2012). Table 5.5, Table 5.6, and Table 5.7 illustrate the respective results for the Keyworth-Nottingham data.

Table 5.2: FMT self-registration results of Oshawa laser point cloud at scale 0.3048

		Oshawa laser point cloud, Scale = 0.3048															
		0°				90°				180°				270°			
		rot	scale	dx	dy	rot	scale	dx	dy	rot	scale	dx	dy	rot	scale	dx	dy
Near neighbour	No filter	180.0	1	0	46	180.0	1	0	43	0	1	0	46	0	1	0	42
	Bartlett	0	0.3054	0	1	275.6	1	0	24	160.3	0.9696	41	5	264.4	1	38	34
Bilinear	No filter	182.8	1.0330	1	45	180.0	1	0	43	0	1	46	1	0	1	0	42
	Bartlett	0	0.3005	0	0	175.8	1	0	41	355.8	1	0	45	357.2	1	0	41
	Chebwin	7.0	0.9272	39	0	251.5	1	21	25	180.0	1.1678	0	22	74.5	1	21	25
Cubic	No filter	182.8	1.0330	1	45	180.0	1	43	1	90.0	1.0156	30	4	0	1	0	42
	Bartlett	0	0.3005	3	0	243.3	1	0	35	355.8	1	0	45	354.4	0.9546	0	42

Many co-registration parameters in Table 5.2 are highlighted in red (false parameter estimate) and blue (inaccurate estimate); it shows that the standard FMT performed poorly at a scale offset of 0.3048. A closer look revealed that it was mostly the translation estimates that solved correctly. In total the standard FMT was able to recover all co-registration parameters in only three out of 28 tests. These three tests had 0deg rotation offset.

In comparison to using no window filter, the application of the Bartlett mask improved the results: with nearest neighbour interpolation the FMT delivered seven correct parameter estimates whereas only four corrections were obtained without filter. Seven co-registration parameters were

also obtained with bilinear and cubic interpolation and, again, only four were obtained without filter mask.

The Bartlett window gave slightly better results than the Chebyshev window, and the worst results were obtained when no window filter was applied. Overall, for a scale offset of 0.3408 (i.e., template image referenced in meters and floating image in feet), the FMT estimated 38 parameters (34%) correctly, 25 parameters (22%) acceptable, and failed on 49 parameters (44%).

Table 5.3: Recovery of scale, rotation, and translation using Oshawa laser point cloud at scale 1.0

		Oshawa laser point cloud, Scale = 1.0															
		0°				90°				180°				270°			
		rot	scale	dx	dy	rot	scale	dx	dy	rot	scale	dx	dy	rot	scale	dx	dy
Near neig- bour	No filter	0	1	9	0	90.0	1.0800	9	10	180.0	1	0	0	180.0	1.0800	9	9
	Bartlett	0	1	0	0	90.0	0.9100	13	12	180.0	1	0	0	270.0	0.9100	13	12
Bilinear	No filter	0	1	0	0	88.5	1	12	0	180.0	1	0	0	88.5	1	12	0
	Bartlett	0	1	0	0	90.0	1	0	0	180.0	1	0	0	270.0	1	0	0
	Chebwin	0	1	0	0	271.4	1	15	26	180.0	1	0	0	91.4	1	17	26
Cubic	No filter	0	1	0	0	271.4	1	1	0	180.0	1	0	0	91.4	1	3	0
	Bartlett	0	1	0	0	90.0	1	0	0	180.0	1	0	0	265.8	0.9400	9	16

As Table 5.3 shows, compared to scale offset 0.3048 (Table 5.2), the standard FMT delivered much better results when no scale offset was applied. The co-registration failed only in five out of 112 instances: Here, it settled in the opposite quadrant – 270deg and 90deg rotation estimate, instead of the actual 90deg and 270deg, respectively in four cases, and in one case it settled incorrectly in the neighboring quadrant. In total, the FMT was able to recover the full set of co-registration parameters in 22 out of 28 tests.

The use of the Bartlett window resulted in twelve correct parameter sets whereas only eight correct sets were obtained without window filter. As before when scale offset 0.3048 was applied (Table 5.2), the Bartlett window gave slightly better results than the use of the Chebyshev window, and the worst results were obtained when no window filter was applied. Overall, without scale offset (reference image and floating image are both either in feet or meters), the FMT estimated 104 parameters (93%) correctly, two parameters (2%) acceptable, and failed only 6 estimates (5%).

Table 5.4: Recovery of scale, rotation, and translation using Oshawa data at scale 3.2808

		Oshawa laser point cloud, Scale = 3.2808															
		0°				90°				180°				270°			
		rot	scale	dx	dy	rot	scale	dx	dy	rot	scale	dx	dy	rot	Scale	dx	dy
Near neighbour	No filter	180.0	1.0212	107	42	180.0	1	0	31	270.0	1	0	35	0	1	0	36
	Bartlett	0	3.2973	0	0	180.0	1.0625	0	43	180.0	3.2925	0	0	180.0	1.0625	0	35
Bilinear	No filter	0	1	127	0	180.0	1	0	32	270.0	1	0	35	0	1	0	35
	Bartlett	0	3.2973	0	0	270.0	0.9412	127	37	180.0	3.2925	0	0	90.0	0.9412	127	37
	Chebwin	0	3.2973	0	0	192.7	1.0204	70	40	180.0	3.2925	0	0	14.1	0.8859	66	0
Cubic	No filter	0	1	127	0	180.0	1	0	34	270.0	1	0	35	0	1	0	36
	Bartlett	0	3.2973	0	0	180.0	1.0625	0	32	180.0	3.2925	0	0	180.0	1.0625	0	35

At first sight, Table 5.4 shows a lot of commonalities with Table 5.2 (FMT results for scale offset 0.3048) as both tables contain many false rotation and translation estimates. However, a closer look reveals that the co-registration worked better at scale 3.2808 and resulted more correct rotation and scale estimates. In total, the standard FMT was able to recover the full set of four co-registration parameters in only eight out of 28 tests. These eight tests had either 0deg or 180deg angular offset.

The application of the Bartlett window improved the results: with nearest neighbour interpolation the FMT delivered nine correct parameter estimates whereas only three corrections were obtained without filter. Similarly, with bilinear interpolation, eight sets of correct co-registration parameters were obtained with Bartlett filter and only five without filter mask. And with cubic interpolation, ten correct parameter sets were obtained with mask and only five without.

The Chebyshev window gave slightly better results than the use of the Bartlett window, and the worst results were obtained when no window filter was applied. Overall, for a scale offset of 3.2808 (target image referenced in feet and floating image in meter), the FMT estimated 51 parameters (46%) correctly, 17 parameters (15%) acceptable, and failed in 44 instances (39%).

Table 5.5: Recovery of scale, rotation, and translation using KN data at scale 0.3048

		KN laser point cloud, Scale = 0.3048															
		0°				90°				180°				270°			
		rot	scale	dx	dy	rot	Scale	dx	dy	rot	scale	dx	dy	rot	scale	dx	dy
Near neig- bour	No filter	82.1	1	11	2	0	1.1850	17	9	180.0	1	25	0	180.0	1	30	35
	Bartlett	82.1	1	11	2	181.3	0.9527	7	9	165.9	1	0	23	171.6	1	28	23
Bilinear	No filter	91.4	1	8	2	180.8	1	1	1	180.0	1	0	0	180.0	1	0	1
	Bartlett	347.3	1	16	13	188.4	1	1	1	165.9	1	23	0	68.9	1	11	11
	Chebwin	347.3	1	16	13	251.5	1	21	25	180.0	1.1678	0	22	74.5	1	21	25
Cub ic	No filter	21.1	1	15	10	180.8	1	1	1	180.0	1	25	0	180.0	1	25	36
	Bartlett	22.5	1	15	10	109.7	1	32	11	165.9	1	0	23	171.6	1	28	23

Similar to Table 5.2, the majority of estimated co-registration parameters for scale offset 0.3048 in Table 5.5 turned out either false (red) or inaccurate (blue). As before, it was mostly the translation estimates that solved correctly. Even though the standard FMT was able to recover all 180deg rotational offsets, it could not recover even a single set of co-registration parameters.

The application of the Bartlett and MATLAB's Chebyshev "Chebwin" mask (MathWorks, 2019) delivered better results than standard FMT without window filter, but none of the test combinations yielded correct scale estimates. Overall, for a scale offset of 0.3408 (template image referenced in meters and floating image in feet), the FMT estimated 38 parameters (34%) correctly, 20 parameters (18%) acceptable, and failed on 54 parameters (48%).

Table 5.6: Recovery of scale, rotation, and translation using KN data at scale 1.0

		KN laser point cloud, Scale = 1.0															
		0°				90°				180°				270°			
		rot	scale	dx	dy	rot	scale	dx	dy	rot	scale	dx	dy	rot	scale	dx	dy
Near neig- bour	No filter	0	1	0	0	23.9	1	9	40	180.0	1	0	0	203.9	1	0	40
	Bartlett	0	1	0	0	23.9	1	9	40	180.0	1	0	0	203.9	1	0	40
Bilinear	No filter	0	1	0	0	90.0	1	0	0	180.0	1	0	0	270.0	1	0	0
	Bartlett	0	1	0	0	271.4	0.9332	0	11	180.0	1	0	0	115.3	1	18	3
	Chebwin	0	1	0	0	118.1	1	0	1	180.0	1	0	0	208.1	1	5	0
Cub ic	No filter	0	1	0	0	90.0	1	0	0	180.0	1	0	0	270.0	1	0	0
	Bartlett	0	1	0	0	77.3	1	13	0	180.0	1	0	0	257.3	1	13	0

Table 5.6, similarly to Table 5.2, shows the standard FMT delivered mostly good results when no scale offset was applied. As before with the Oshawa data, the standard FMT performed also better with the KN data and delivered much better results without scale offset compared to scale 0.3048 (Table 5.5). It worked best for 0 degree and 180 degree rotation, which gives much better SNR. In total, the standard FMT was able to recover the full set of co-registration parameters in five out of 28 tests.

The use of the Bartlett window resulted in eight correct parameter sets whereas twelve correct sets were obtained without window filter. Unlike before when scale offset 0.3048 was applied (Table 5.5), the Chebyshev window gave slightly better results than the use of the Bartlett window. Overall, without scale offset (reference image and floating image are both either in feet or meters), the FMT estimated 97 parameters (87%) correctly, 3 parameters (2%) acceptable, and failed only 12 estimates (11%).

Table 5.7: Recovery of scale, rotation, and translation using KN data at 3.2808

		KN laser point cloud, Scale = 3.2808															
		0°				90°				180°				270°			
		rot	scale	dx	dy	rot	scale	dx	dy	rot	scale	dx	dy	rot	scale	dx	dy
Near neig bour	No filter	0	3.2755	0	0	180.0	1	33	60	180.0	1	0	81	270.0	1	1	100
	Bartlett	0	3.2755	0	0	277.0	1	14	68	196.9	1	35	81	105.5	1	0	60
Bilinear	No filter	0	3.2755	0	0	0	1	50	1	180.0	1	81	1	0	1	21	1
	Bartlett	0	3.2755	0	0	90.0	1	23	1	15.5	1	33	81	15.5	1	33	81
	Chebwin	0	3.2755	0	0	208.1	0.9418	50	81	180.0	3.2851	0	0	358.6	0.9315	0	110
Cub ic	No filter	0	3.2755	0	0	180.0	1	0	10	180.0	1	0	81	180.0	1	45	40
	Bartlett	0	3.2755	0	0	21.1	1.6731	26	78	180.0	3.2851	0	0	28.1	1.5043	37	81

Table 5.7 shows that the FMT performed especially well when no rotation was applied. In total, the standard FMT was able to recover the full set of four co-registration parameters in only nine out of 28 tests. The FMT showed similar performance with application of the Bartlett window and without window mask: with nearest neighbour interpolation the FMT delivered six correct parameter estimates whereas eight were obtained without filter. With bilinear interpolation, six sets of correct co-registration parameters were obtained both with Bartlett filter and without filter

mask. And with cubic interpolation, eight correct parameter sets were obtained with mask and only seven without mask.

As before, the Chebyshev mask gave slightly better results than the use of the Bartlett window, and the worst results were obtained when no window filter was applied. Overall, for the scale offset of 3.2808 (target image referenced in feet and floating image in meter), the FMT estimated 63 parameters (56%) correctly, 3 parameters (3%) acceptable, and failed in 46 instances (41%).

A detailed discussion of the real-world tests is given in APPENDIX C.

5.1.2.3 Summary

In section 5.1.2.2 and its subsections, the standard FMT was reviewed as another tool to co-register tie- and laser point clouds. Unlike the methods examined earlier, the FMT was developed specifically for co-registration. However, it was developed for computer vision and medical applications.

After introducing the standard FMT workflow, properties that could limit its performance with geodetic data, such as effects of padding and loss of phase information, were highlighted. To better study the FMT performance with geodetic data, the standard FMT was analyzed from the viewpoint of DOE. This allowed to get a thorough overview of the design space, led to efficient ways to reduce the design space, allowed to study the sensitivity of the FMT to certain parameters, and it allowed to identify optimal parameter settings quickly. Using theoretical reasoning and experiments with real-world and simulated DEM data, the following factors were analyzed:

- (1) Number of correlation peaks
- (2) Cropping
- (3) Window filters
- (4) High pass filter
- (5) Interpolation method
- (6) Magnitudes of the elevation map
- (7) Sensitivity to noise in the elevation map (SNR)
- (8) Image rendering
- (9) Convergence behavior and reliability of the FMT results

The experiments showed that it is sufficient to consider the twenty tallest correlation peaks in the search for the correct solution candidate (1 above), that image padding should be used instead of cropping (2 above), that application of an image mask with medium damping power can improve the co-registration results in cases where the input data shows strong elevation changes and it can degrade the results for data with low SNR (where reduction in information content outweighs the benefits of damping the image borders and reduction of the “+” effect) (3 above), and that application of a standardized HF mask advances the results (4 above). Changing the interpolation method (5 above) showed little effect on the outcome. Changes in magnitude in the elevation map (6 above) in conjunction with the SNR (7 above) turned out to be the most crucial success-driving factors – the FMT solved correctly mostly for input data with strong elevation changes and low levels of noise.

This research also uncovered two unexpected symptoms, in that effects of image rendering can create artifacts and distort the parameter estimates (8 above), and in that there is no relationship between the shape of the parameter curves associated with the 20 tallest CPS peaks and how accurate and reliable the final solution is (9 above). For the examination of the convergency behaviour, it turned out especially beneficial to replace the CPS graphs by the more meaningful curve of the twenty tallest CPS peaks (see for example APPENDIX D).

The experiments also showed that FMT can recover all cardinal rotations $\pm 10^\circ$, where only one “+” effect occurs, when simulated elevation data was used. Degraded results were observed in case of a 45 (135,225,315) degree rotation where two “+” overshadow the elevation signal. In case of other rotational offsets, the FMT failed often, because the changes in elevation were too small to compensate for information overlay of the multi- “+” effect.

The results provided important insights about the relation of data input, noise parameters, and the success rate of the standard FMT, and led to development of new concepts:

1. Large scale offsets turned out to be the most difficult parameter to solve. When paired with an angular offset, the FMT was unable to co-register the real-world data in many cases. Closer investigation showed that one root cause was the loss of information from down sampling the search image.

This finding led to the idea of circumventing the reduction of the search image and starting a polar FMT with an enlarged template image instead whenever the log-polar FMT estimates a scale offset significantly smaller than 1.0. So, when the target point cloud is referenced in meters and the search point cloud is referenced in feet, and the log-polar FMT estimates a scale offset near 0.3, the scale could be fixed at 0.3048, and a polar FMT could be run using an upscaled target image instead of a downscaled search image. That way, the polar FMT could estimate the remaining rotation and translation offsets separately, and with more accuracy since the difference in scale between search and target point cloud has been removed.

However, this workaround may not always improve the co-registration results since the rotation and scale offsets could correlate too much and cause the scale estimate to settle far away from the true offset.

2. In theory, since rotation and scale are orthogonal in log-polar base, the FMT should result the same rotation and scale estimates when the data is set apart by a certain angle only, or by a certain scale only, or by both offsets at the same time. The research showed however that the FMT performed better when the point clouds were set apart by either rotation *or* scale, and it performed worse, when the point clouds were offset by both rotation *and* scale at the same time.

Had the phase spectra of the elevation maps been rendered into log-polar base, then this correlation would be a clear indicator that the point clouds themselves did not provide enough elevation difference for full discernment of rotation- and scale offsets in frequency domain. However, as discussed in APPENDIX E, it is magnitude of the height maps' phase spectra that gets rendered into log-polar base inside the FMT. Therefore, this correlation could be caused by *either* by (A) insufficient elevation changes in the height map (low SNR), *or* (B) too much loss of the height map's phase information by mapping the magnitude (instead of the amplitude) of the phase spectra into log-polar base, *or both*.

From the above experiments it is clear the FMT struggled with large scale offsets due to loss of surface information and respective phase information through image scaling. It became also clear that the FMT results showed correlation between scale and rotation offset. However, it

remained unknown how much of this correlation is introduced by the lack of undulation or by the loss of phase information of the FMT algorithm.

In the 2000s, this detriment was noticed by a few European scientists, who proposed the AFMT, a modified FMT that maintains the phase information. Thus, repeating some of the experiments above using the same data and same noise parameters with the AFMT should lead to a better understanding of the impact of loss of phase information versus the impact of elevation profile on the registration results. This concept and the earlier introduced idea of applying polar rendering instead of log-polar mapping will be examined further in the next chapter.

5.2 Conclusions

The exploration of time domain-based co-registration algorithms started with the POMF, the most basic frequency-based co-registration method presented in this framework. The POMF can only recover translation offsets. MATLAB's "imregcorr" POMF function (MathWorks, 2019) was able to recover translation offsets up to 85% when the height maps were self-registered, however it failed completely to cross-correlate the DEMs successfully. This insight led to an expanded POMF model where the POMF with multi-peak evaluation was used for recovery of translation offsets that represent angular and scalar unknowns. This improved POMF was then used inside the subsequent frequency-based co-registration methods. The POMF tests also revealed that there is neither a common threshold for overlap at which successful auto-correlation can occur (first proposition, Section 5.1.1.4), nor is there a threshold for entropy at which successful co-registration can occur (second proposition, Section 5.1.1.4).

Then, the standard FMT, developed specifically for co-registration applications, was reviewed as another tool to co-register sparse tie- and laser point clouds. First, the standard FMT was analyzed from the viewpoint of DOE, to reduce the design space effectively and identify optimal parameter settings quickly. Using theoretical reasoning and experiments with real-world and simulated DEM data, the impact of 12 FMT parameters were analyzed. These experiments provided important insights about the relation of data input, noise parameters, and success rate and

ultimately, the best parameter settings. These proved to be applicable to the AFMT algorithms, too, as they recycle some of the FMT functionality.

Amongst the most important setting-related findings was that it is sufficient to consider the twenty tallest correlation peaks in the search for the correct solution candidate, and that application of an image mask with medium damping power should be applied, and that changes in magnitude in the elevation map in conjunction with the SNR are the most crucial success-driving factors – the FMT solved correctly mostly only for input data with strong elevation changes and low levels of noise. The latter finding led to the idea of circumventing the reduction of the search image by using an enlarged template image instead whenever the algorithm would estimate a scale offset significantly smaller 1.0. This workaround was applied in the consecutive AFMT experiments, too.

The experiments also showed that the FMT performed better when the point clouds were set apart by either rotation or scale, and it performed worse, when the point clouds were offset by both rotation and scale at the same time. It became also clear that the FMT results showed correlation between scale and rotation offset. However, it remained unknown how much of this correlation was caused by insufficient elevation changes in the height map (low SNR), or too much loss of the height map's phase information by mapping the magnitude (instead of the amplitude) of the phase spectra into log-polar base, or both.

6.0 Coregistration in Frequency Domain: Expanded Existing and Novel Methods

This section examines the performance of more advanced, rarely used frequency-based co-registration methods and a newly proposed method, in the end. This Chapter starts with the review of the Analytical Fourier Mellin Transform (AFMT) (Derrode and Ghorbel, 2001) with its three versions – the Polar Analytical Fourier Mellin Transform (PAFMT), the Log-Polar Analytical Fourier Mellin Transform (LAFMT), and the Extended Log-polar Analytical Fourier Mellin Transform (EAFMT), which recover rotation, rotation and scale, and rotation and scale and translation, respectively. Afterwards, it ends with the review of the novel Quasi Log-Polar Analytical Fourier Mellin Transform (QAFMT) which recovers rotation, scale, and translation. While the use of the Log-polar Analytical Fourier Mellin is well known, it is the unique combination of the LAFMT and PAFMT in the QAFMT that is novel and that allows for improved coarse registration performance. Also, the use of these algorithms with extremely sparse point clouds is novel

As in Chapter 5.0, the focus is here on particularly achieving optimal co-registration results with airborne laser- and tie point clouds, which can differ strongly in point density, size of area, shape, and noise level. The following sections will highlight the characteristics of each transform that imply restrictions on their performance with airborne survey data and compare their performance.

6.1 The Analytical Fourier Mellin Transform

6.1.1 Introduction

The traditional FMT co-registration algorithm, which applies the POMF to log-polar Fourier magnitudes, is widely used in computer vision since it allows for full recovery of all four registration parameters (rotation, scale, translation in x,y). In computer vision applications, the object of interest is typically clearly offset by a homogenous background and the Fourier magnitude contains enough information about the object to make up for the loss of phase

information that occurs inside the traditional FMT. However, as concluded in Chapter 5.1.2.3, the change in terrain elevation in airborne surface scans does not always build up enough signal strength to allow for recovery of all registration parameters from the Fourier magnitude only.

Approximately ten years after the introduction of the traditional FMT by Horner and Gianino (1984) and De Castro and Morandi (1987), a group of French scientists (Derrode and Ghorbel, 2001) tried to co-register close range geospatial imagery and faced the same problem, such that the image content provided low contrast, yet the entire scenery acted as object of interest. In 2001, they proposed a modified version of the FMT that maintains both the image's phase and amplitude information - the Analytical Fourier Mellin Transform (AFMT).

6.1.2 Definition

As illustrated in Chapter 5.1.2.2, the traditional FMT encompasses cascaded application of the DFT and log-polar sampling of frequency spectra. Derrode and Ghorbel (2001) were concerned that the FM integral does not fully converge at the image centre, which is also the centre of the log-polar coordinate system. This shortcoming could evoke critical loss of image information in both the DFT of log-polar images and the DEFT of log-polar image spectra.

As a workaround, Derrode and Ghorbel (2001) suggested cutting out a small disc around the image centre. They published the respective modified algorithm under the name “Analytical Fourier Mellin Transform” (AFMT). Strictly speaking, it is only this treatment of the image centre and respective modification of the FM integral what distinguishes the AFMT from the FMT. However, since this analytical refinement of the Fourier Mellin integral was suggested along with log-polar sampling of the input images themselves (instead of phase spectra), the term AFMT has been associated primarily with this altered sequence in the FMT workflow.

The AFMT of an image is defined as

$$M_{f_\sigma}(k, v) = \frac{1}{2\pi} \int_0^\infty \int_0^{2\pi} f(r, \varphi) e^{-ik\varphi} \cdot r^{\sigma-iv} \cdot \frac{dr}{r} d\varphi \quad (6.1)$$

where

$k \in \mathbb{Z}$ typically, with $k \in [0, \dots, N]$ where N is the number of log-polar increments which equal the width and length of the input image in pixels

$v \in \mathbb{R}$, ranging $[0 \dots 2\pi]$

$f(r, \varphi)$ the log-polar sampled image

$\sigma > 0$ is a constant number with $\sigma \in \mathbb{R}^+$

The cut of a small disc is obvious in equations (6.1) and (6.2) where the image input f into the AFMT is not only a function of r and φ but also of σ . The parameter σ allows to spare a small disc around the origin of the coordinate axis, thus avoiding the singularity that would arise if regular log-polar sampling would occur near 0,0.

The inverse AFMT finds

$$IAFMT\{M_{f_\sigma}(k, v)\} = f_\sigma(r, \varphi) = \int_0^\infty \sum_Z M_{f_\sigma}(k, v) \cdot r^{-\sigma+iv} \cdot e^{-ik\varphi} dv \quad (6.2)$$

where $M_{f_\sigma}(k, v)$ is the image in AFMT representation, i.e., the “analytic prolongation” (Chen, 1993; Chen *et al.*, 1994) or the FT’ed log-polar sampled image.

In analogy to equation (5.5), the CPS finds

$$CPS(r, \varphi) = \int_0^\infty \sum_Z M_{f_{1\sigma}} M_{f_{2\sigma}} \cdot r^{-\sigma+iv} e^{ik\varphi} dv \quad (6.3)$$

Since the rotational and scalar offsets appear as translative movement between $M_{f_{1\sigma}}$ and $M_{f_{2\sigma}}$, they can be determined from the inverse of the CPS. In analogy to equation (5.6), the inverse CPS finds

$$[\Delta x, \Delta y] = IFFT\{CPS[k, v]\} \quad (6.4)$$

where $[\Delta x, \Delta y]$ are the peak locations that relate to the looked-for offsets in rotation and scale, that can be computed from equation (5.17).

Since its postulation in 2001, three forms of the AFMT evolved: The “standard” log-polar AFMT which was introduced first and is most commonly used, the polar AFMT which is the simplest AFMT variant but rarely used, and the extended AFMT. These AFTM variants and their respective properties and co-registration results will be presented in the ensuing chapters. The results will be cross-compared to find out which AFMT version is the most promising for co-registration of airborne geodetic data. A summary of the findings will be given at the end.

6.1.3 Co-Registration using Log-Polar AFMT

The following sections introduce the Log-polar AFMT (LAFMT), which is the most used form of the AFMT, along with its properties and respective co-registration results.

6.1.3.1 *The LAFMT Algorithm*

As shown in Chapter 6.1.2, AFMT algorithms are a modified version of the traditional FMT algorithm. In the LAFMT, which is the “standard”, most often used form of the AFTM, the POMF is applied on log-polar imagery rather than on the log-polar image of FMT amplitude spectra. Figure 6.1 and Figure 6.2 below illustrate how the LAFMT can be derived from the traditional FMT algorithm.




Figure 6.1: Modification of the FMT into the LAFMT algorithm

Figure 6.1 outlines how the components of the LAFMT compare to the standard FMT. As can be seen, steps B and C (which cover the first FT and computation of magnitude in the standard FMT) are left out in the LAFMT. Thus, the property of translation invariance is lost, which is highlighted with red crosses at steps E and F.

The disembowelled version of Figure 6.1 is illustrated below.





Figure 6.2: The LAFMT algorithm


As outlined in Figure 6.2, the log-polar, “standard” AFMT algorithm comprises three steps:

1. Input of two image-rendered point clouds
2. Step A: The images are rendered into log-polar base where rotation and scale are decoupled and associated the x- and y- axis, respectively. Optionally, if the data has quadratic shape, a window filter can be applied before log-polar mapping to soften the image borders. The same window masks as used in the FMT (APPENDIX E) can be deployed.
3. Step B: The log-polar images are ingested into a POMF to isolate the rotational and scalar offsets. The FFT inside the POMF perceives the incoming images as infinite periodic signal and thus must consider the ambiguous location of the correlation peak as referenced to the first and third quadrant. Thus, the POMF outputs two estimates for the angular offset, where one estimate offsets the other by 180 degrees.
4. Step C: The original “moving” input image is scaled and rotated one time by φ and one time by $\varphi+180$ to locate the best estimate for the set of $\Delta\varphi, \Delta s$.

If only subtle changes in image amplitude are expected, i.e., if the SNR of the phase spectrum is high, many of the CPS correlation peaks found inside the POMF (step B) may be considered. In that case, steps B and C are repeated for each candidate and the $\Delta\varphi, \Delta s$ combination that results the highest correlation peak in Step C is chosen as final candidate.

Figure 6.3 exemplifies each step of the LAFMT algorithm as the X5d4 DTM is self-registered with its 30deg rotated counterpart.








Figure 6.3: Self-registration of a 30deg-offset X5d4 DTM (pix) using LAFMT

Figure 6.3 illustrates results of each step of the LAFMT as suggested by Derrode and Ghorbel (2001). In this example, the X5d4 DTM was self-registered with an angular offset of 30degrees. Cells a1,1-3 and a2,1-3 show the height maps in log-polar rendering. In log-polar view, the angular offset appears as translation along the x-axis. “Peak B”, highlighted with a red dot, shows an approximate x-coordinate of 2150 before rotation (a1,1 and a1,3) and 180 after rotation (a2,1 and a2,3). Accordingly, the CPS peak highlighted with a red arrow, showed a shift of roughly 200pixels. As per equation (5.17), the angular offset of 30deg in Cartesian base relates to a full rotation of 360deg the same way as the translation of 200pixels in log-polar base relates to the full length of 2500pixels. Indeed, $30deg/360deg \approx 200pix/2500pix$ confirms the 30deg rotation. Since the similarity of the log-polar images was very pronounced, the CPS resulted a dominant first peak.

While as aforementioned, polar rendering would have been sufficient to recover the angular offset, log-polar rendering is required to recover scale. Figure 6.4 below exemplifies the effect of a scalar offset of 2.0 for the X5d4 DTM in Cartesian and log-polar rendering.




Figure 6.4: Effects of rotation and scale for the X5d4 DTM (pix) in Cartesian (a,b) and log-polar (c,d) base

In Figure 6.4, cell (a) shows the X5d4 DTM in its original orientation and resolution. Cell (b) shows the same image after it was enlarged by factor 3.2808 and rotated counter-clockwise by 90 degrees. In both cells, the location of “peak B” of the X5d4 DTM was highlighted with a red dot. Cells (c) and (d) present the respective log-polar rendering. Comparing the locations of “peak B” shows that the offset in rotation and scale translated into an offset in the x- and y- axis in the log-polar plots, where the x- (column) and y- (row) axis represent rotation and scale, respectively.

Based on visual inspection the centre of “peak B” located approx. 1200 pixels further left and approx. 500 pixels further down in cell (d) than in cell (c). The calculated values confirmed that the centre of “peak B” in cells (c) and (d) were apart by $dcol = 1530$ pixels and $drow = 452$ pixels, respectively. Based on equations (5.17) and (5.18), the change in rotation and scale calculated as

$$size = \log_{base} \frac{rows}{2} \quad (6.5)$$

$$base = e^{\frac{\log \frac{rows}{2}}{size}} = e^{\frac{\log \frac{6116}{2}}{6116}} = 1.0013 \quad (6.6)$$

$$\Delta\phi = \frac{dcols}{cols} * 360deg = \frac{1530}{6116} * 360deg = 90.0589deg \quad (6.7)$$

$$\Delta s = base^{2*drow} = 1.0013^{2*452} = 3.2363 \quad (6.8)$$

This confirmed the actual offsets of $\Delta\phi = 90deg$ and $\Delta s = 3.2808$ with high precision.

Lastly, it should be noted that the recovery of flipped data (the X- and Y- axis are swapped) has not been covered in the body of LAFMT-related publications available to the author. Recovery of axis swap could occur through an additional loop in the AFMT algorithm or perhaps through query of the CPS graph against predefined thresholds. Figure 6.5 below shows an example of what the CPS of a flipped height map looks like. While this topic is interesting, it adds only limited value, so the primary focus will remain on the recovery of the standard AFMT in this framework.




Figure 6.5: CPS of self-registered KN tie point cloud (pix) in (a) original orientation and (b) with flipped axis


6.1.3.2 Properties

While the LAFMT algorithm is less complex than the standard FMT, it does have downsides, too. Some of its drawbacks are a consequence of only considering the first derivative in frequency domain and other drawbacks are inherited from log-polar mapping. Its positive and negative properties and respective workarounds summarize as:

- Since the LAFMT considers both the amplitude and phase information of the input images, it should, in theory, be less susceptible to noise and yield more robust registration results than the FMT – especially when low-resolution DEMs with low elevation contrast are co-registered.
- "Since the spectral phase preserves the location of objects but is insensitive to the image energy, the application of the POMF to a pair of identical images under the constraint of translation yields a much sharper peak than the classical FMT" (Chen *et al.*, 1994). In context of the LAFMT, the POMF retrieves translational offsets from log-polar imagery, which resembles scale and rotation factors. In comparison to the FMT, the LAFMT should

therefore, in theory, be able to retrieve scale and rotation offsets within a larger range and/or with higher precision and/or with increased robustness.

- Since, unlike the FMT, the AFMT encompasses only one DFT, it is free of “+”-effects.
- As with the FMT, a window filter can be used to reduce the effects of image-tiling (see Chapter 5.1.2.2.2), if the data has a quadratic footprint. Since the window-filtered images undergo only one DFT, all of the window effect enters the amplitude-based CPS (in comparison to the FMT where parts of the window effect has been lost at the stage of the magnitude-based CPS).
- Since the FT is applied only once, the log-polar images themselves are not invariant to translation. Therefore, the LAFMT, in its native form, can only recover scale and rotation but not translational offsets. Figure 6.6 illustrates this drawback and shows an example where the X5d4 DTM was rendered in log-polar space after it had been offset by translation:






Figure 6.6: Effects of horizontal translation of the X5d4 DTM (pix) in cartesian (a,b) and log-polar (c,d) base

In Figure 6.6, cell (a) displays the X5d4 DTM in its original configuration. Cell (b) shows the DTM after it was shifted by 400 pixels to the right. The location of “peak B” was highlighted with a red dot and the location of the main “peak A” was marked with a blue cross. The removal of 400 pixels caused “peak A” and “peak B” to locate closer to the right border of the height map. It also added a zero-padded band of 200 pixels on both the left and right side. Cell (c) and (d) display the log-polar rendering of cell (a) and (b), respectively. In cell (d), the additional padding became obvious in form of two bulges. And while in cartesian space the translation caused “peak A” and “peak B” to move in the same direction, in log-polar space the translation caused “peak A” to locate further down further and further to the right, yet “peak B” to locate further down and further to the left. Since the peaks moved in opposite direction, they no longer aligned accurately in the CPS.

- Log-polar rendering infers a singularity at 0,0. When the FMT is used for co-registration of elevation profiles, re-sampling around 0,0 may invoke loss of high frequencies (that situate around 0,0 in the phase plot) which could stem from strongly pronounced changes in the elevation profiles. In case of the standard AFMT however, log-polar sampling occurs directly on the height imagery. Thus, the information loss from neglecting few pixels

around 0,0 is minimal, especially since the elevation values in proximity of the origin (or any other location within the height map) have usually similar values than their surrounding pixels. Even if the number of pixels in log-polar and cartesian space were different, the authors of the log-polar AFMT (Derrode and Ghorbel, 2000) foresaw cutting out a small disc around 0,0 to minimize the information loss.

- The accuracy at which the rotation angle and scaling factors can be estimated depends on the log-polar sampling (Chen, 1993). If the number of pixels in log-polar space is chosen to be equal the number of pixels of the cartesian image, then the information content is maintained, the angular coordinate is sampled uniformly, and the precision of the angle is uniform. However, the logarithmic distortion along the radial coordinate leads to a non-uniform precision of the scale estimation (Chen, 1993).
- As with the FMT, in cases where the height maps are noisy and provide only small change in elevation, cross-correlation between the rotation and scale estimates may occur in the LAFMT, too. In this circumstance, an inaccurate scale (rotation) estimate may invoke an inaccurate rotation (scale) estimate. On the other hand, there may be instances where the CPS search settles on a peak that is driven by a strong, correct scale (rotation) estimate, and this scale (rotation) estimate may “pull” the rotation (scale) estimate favorably to a more accurate rotation (scale) estimate.
- Log-polar sampling yields high point density around the origin of the coordinate system, and lower densities away from the centre. This makes sampling of the input image close to the origin of the coordinate system challenging. For image pairs that are only set apart by rotation, polar sampling can be used instead (Sellami and Ghorbel, 2012). By using polar mapping of the images, the aliasing error typically associated with the non-linear distribution of image samples along the scale axis can be reduced. This workaround will be explored in Chapter 6.1.4.1

Implications of these properties in conjunction with the review of numerical results will be discussed in the next chapter.

6.1.3.3 Numerical Results

In analogy to the empirical evaluation of the FMT, the performance of the LAFMT was first analyzed in autocorrelation (the most promising data combination) and subsequently in cross-correlation. Both the real-world and the simulated data were used. As it turned out, since the simulated elevation maps have large grid size, executing a LAFMT with the simulated data and scale factor 3.2808 required many hours. Therefore, as before for the FMT in Chapter 5.1.2, the LAFMT results were created for cardinal rotations only (90deg rotation increments), rather than, for example, in 10deg increments.

As with the FMT, only scale values 1.0 and 3.2808 were tested when correlating the height maps, since - as elaborated in Chapter 5.1.2.3 - this allowed for circumventing the reduction of the search image and conducting the co-registration with an enlarged template image instead, whenever the log-polar AFMT estimated a scale offset significantly smaller than 1.0. Thus, in cases where the target point cloud is referenced in meters and the search point cloud is referenced in feet, and the LAFMT estimates a scale offset near 0.3, the scale can be fixed internally at 0.3048, and the LAFMT can be run using a target image upscaled by 3.2808 instead of a downscaled search image.

The results of the LAFMT experiments are summarized in Table 6.1 below. In these and subsequent tables, the LAFMT-estimated scale and rotation values were evaluated against the same thresholds of 20% and 50% deviation defined in Table 5.1 that were used to evaluate the FMT results. Thus, parameter estimates marked in **green** indicate that the LAFMT terminated **correctly** and that the estimates deviated less than **20%** from their respective nominal value. Parameter estimates that deviated more than 20% but less than **50%** were considered **acceptable** and were highlighted in **blue**, since they led in the right direction (and could therefore serve as a starting point for fine registration). Estimates with more than 50% deviation from the actual value were considered **incorrect** and colour coded in **red**.

Table 6.1: Registration of Oshawa height maps using LAFMT

		Oshawa laser and tie PCL (117-by-136, 82-by-117), Scale = 1.0							
		0°		90°		180°		270°	
Laser – Laser	No filter	2.6	1.40	90.0	1.00	180.0	1.00	90.0	1.00
	Bartlett	-177.4	1.00	90.0	1.00	2.6	1.00	90.0	1.00
	Chebwin	2.6	1.00	90.0	1.00	180.0	1.00	272.6	1.00
Tie -Tie	No filter	3.0	1.03	95.2	1.00	150.2	1.00	-87.8	1.03
	Bartlett	3.0	1.00	92.2	1.00	-177.0	1.00	-87.8	1.00
	Chebwin	3.0	1.00	92.2	1.07	-177.0	1.00	267.7	1.03
Laser -Tie	No filter	280.3	1.03	-66.9	1.10	267.4	1.03	2.6	1.00
	Bartlett	0.0	1.40	2.6	2.89	2.6	2.89	185.1	2.34
	Chebwin	15.4	1.44	-2.6	1.83	105.4	1.58	20.6	1.35
Tie -Laser	No filter	169.5	1.45	267.4	1.40	349.5	1.45	84.9	1.27
	Bartlett	357.4	1.40	36.0	2.50	2.6	1.40	-144.0	2.50
	Chebwin	-140.6	1.03	321.4	1.13	2.6	1.03	-164.6	1.28

(a)

		Oshawa DEMs, Scale = 3.2808							
		0°		90°		180°		270°	
Laser – Laser	No filter	360.0	3.27	179.6	2.78	267.6	2.73	-89.8	3.29
	Bartlett	0.0	3.31	90.2	3.29	-179.2	3.27	-89.8	3.29
	Chebwin	1.6	3.27	-178.4	1.14	1.6	3.27	0.8	1.35
Tie -Tie	No filter	358.1	3.27	90.0	3.27	179.1	3.31	270.0	3.27
	Bartlett	359.1	3.31	90.0	3.27	179.1	3.31	270.0	3.27
	Chebwin	1.9	3.22	90.0	2.06	-177.2	3.27	270.9	3.18
Laser -Tie	No filter	270.0	2.00	180.9	2.00	270.0	2.00	360.0	1.97
	Bartlett	90.0	1.97	180.9	2.00	270.0	2.00	360.0	1.97
	Chebwin	90.9	1.97	180.9	2.00	90.9	2.92	0.9	1.63
Tie -Laser	No filter	-179.2	4.60	-89.8	4.54	0.8	4.60	-89.8	4.54
	Bartlett	360.0	4.71	-89.8	4.54	180.0	4.71	-89.8	4.54
	Chebwin	-179.2	1.60	0.8	13.54	270.6	1.23	-179.2	13.54

(b)

Table 6.1a shows the LAFMT co-registration results of the Oshawa tie- and laser point clouds without scale offset. The upper half shows the self-registration results. During autocorrelation, the LAFMT failed only in 4 out of 24 instances, i.e., it converged correctly in 83% of the test cases. The lower part shows the respective cross-correlation results, which reflect practical geodetic application. Here, the LAFMT converged correctly in only 3 out of 24 cases, i.e., it showed a success rate of only 13%. Table 6.1b shows the LAFMT co-registration results with scale offset 3.2808. The upper half shows the results where the point clouds were self-registered. Here, the LAFMT failed only in 3 out of 24 instances, i.e., it converged correctly in 87% of the test cases. When the Oshawa tie- and laser point clouds were co-registered (lower half of Table 6.1b), the LAFMT converged correctly in only 4 out of 24 cases, i.e., it showed a success rate of only 17%.

Overall, the co-registration of Oshawa data with LAFMT succeeded in only 15% of the test cases.

Table 6.2a, below, shows the LAFMT co-registration results of the Keyworth-Nottingham tie- and laser point cloud without scale offset. The upper half shows the self-registration results. Here, the LAFMT showed 100% success. The lower half shows the respective cross-correlation results. Here, the LAFMT converged correctly in 10 out of 24 cases, i.e., it showed a success rate of 42%. Table 6.2b, below shows the co-registration results for the KN tie- and laser point cloud with scale offset 3.2808. When the KN data was self-registered, the LAFMT failed only in 1 instance, i.e., it converged correctly in 96% of the test cases. When the KN data was co-registered, the LAFMT converged correctly in 16 cases, i.e., it showed a success rate of 67%.

The overall success rate (with and without scale offset, self- and co-registration) of the registration of KN data with LAFMT was 55%.

Table 6.2: Registration of KN height maps using LAFMT

		KN DEMs (61x73, 40x37), Scale = 1.0							
		0 ⁰		90 ⁰		180 ⁰		270 ⁰	
Laser - Laser	No filter	9.9	1.00	88.8	1.00	-150.4	1.00	286.0	1.00
	Bartlett	4.9	1.05	93.7	1.00	180.0	1.05	276.2	1.00
	Chebwin	4.9	1.05	93.7	1.00	180.0	1.05	276.2	1.00
Tie - Tie	No filter	24.5	1.00	114.1	1.00	-163.6	1.33	281.0	1.00
	Bartlett	8.2	1.00	96.6	1.00	-171.8	1.00	272.2	1.08
	Chebwin	8.2	1.00	101.0	1.00	-171.8	1.00	272.2	1.08
Laser - Tie	No filter	360.0	1.53	37.4	1.05	-133.2	1.53	243.1	1.40
	Bartlett	9.4	1.00	72.5	1.70	-175.3	1.00	275.8	1.33
	Chebwin	4.7	1.00	95.8	1.00	-175.3	1.00	275.8	1.00
Tie - Laser	No filter	-175.3	1.60	360.0	1.15	360.0	1.94	180.0	1.85
	Bartlett	4.7	1.00	93.5	3.94	166.0	1.53	271.2	4.33
	Chebwin	346.0	1.39	84.2	1.60	166.0	1.53	-105.2	1.53

(a)

		KN DEMs, Scale = 3.2808							
		0 ⁰		90 ⁰		180 ⁰		270 ⁰	
Laser - Laser	No filter	360.0	3.30	89.6	3.30	182.2	3.30	-6.7	2.93
	Bartlett	360.0	3.30	92.6	3.30	182.2	3.30	271.9	3.30
	Chebwin	3.0	2.30	91.1	3.11	-178.5	2.93	271.9	3.30
Tie - Tie	No filter	2.6	3.18	90.0	3.18	-177.4	3.18	272.6	3.28
	Bartlett	10.6	3.28	92.6	3.18	-177.4	3.18	-87.4	3.08
	Chebwin	5.3	3.49	92.6	3.08	-177.4	2.99	270.0	3.08
Laser - Tie	No filter	344.1	1.59	62.3	1.66	232.9	1.50	-44.7	1.51
	Bartlett	2.6	1.92	90.7	3.42	272.6	3.15	-79.8	3.32
	Chebwin	2.6	2.31	105.6	2.07	-177.4	1.92	-90.7	1.88
Tie - Laser	No filter	360.0	4.91	1.5	22.73	180.0	4.91	-178.5	22.73
	Bartlett	2.2	4.91	90.0	2.53	182.2	4.91	270.0	3.09
	Chebwin	-85.9	2.45	90.0	2.53	-148.6	1.64	-178.5	2.39

(b)

Table 6.3a, below, illustrates the LAFMT co-registration results of the simulated tie- and laser point cloud without scale offset. When the data was self-registered, the LAFMT failed only in 2 instances and showed a 92% success rate. The lower half shows the respective cross-correlation results. Here, the LAFMT converged correctly in 8 out of 24 cases, i.e., it showed a success rate of 34%.

Table 6.3b shows the co-registration results of the simulated tie- and laser point cloud with scale offset 3.2808. When the data was self-registered, the LAFMT showed 100% success. When the simulated data was co-registered, the LAFMT converged correctly in only 4 cases, i.e., it showed a success rate of 17%.

The overall success rate of the co-registration of simulated data with LAFMT was 26%.

Table 6.3: Registration of simulated height maps using LAFMT

		Simulated DEMs, Scale = 1.0							
		0 ⁰		90 ⁰		180 ⁰		270 ⁰	
DTM - DTM	No filter	10.2	1.00	90.0	1.00	179.8	1.00	270.0	1.00
	Bartlett	10.2	1.00	91.2	1.00	180.0	1.01	270.4	1.02
	Chebwin	10.2	1.00	0.2	1.00	180.0	1.01	89.4	1.01
DSM - DSM	No filter	0.2	1.30	90.0	1.00	183.3	1.00	270.0	1.00
	Bartlett	10.2	1.00	90.0	1.00	180.2	1.00	270.0	1.00
	Chebwin	10.2	1.00	90.0	1.00	180.2	1.00	270.0	1.00
DTM - DSM	No filter	11.4	1.00	100.6	1.00	187.9	1.00	280.2	1.00
	Bartlett	87.5	2.49	79.0	2.21	-88.3	4.29	7.3	2.92
	Chebwin	261.7	2.71	173.4	3.70	179.2	2.67	-78.6	2.91
DTM - DSM	No filter	6.4	1.00	97.9	1.00	187.9	1.00	281.8	1.00
	Bartlett	103.3	1.40	-179.2	2.34	-62.6	2.47	-56.6	2.53
	Chebwin	69.3	1.02	270.6	1.65	158.9	2.49	33.0	2.43

(a)

		Simulated DEMs, Scale = 3.2808							
		0 ⁰		90 ⁰		180 ⁰		270 ⁰	
DTM - DTM	No filter	0.1	3.28	89.9	3.28	180.1	3.28	269.9	3.28
	Bartlett	360.0	3.27	89.9	3.27	180.0	3.27	269.9	3.27
	Chebwin	360.0	3.27	89.9	3.27	180.1	3.27	269.9	3.27
DSM - DSM	No filter	0.1	3.28	90.0	3.28	180.1	3.28	270.0	3.28
	Bartlett	360.0	3.28	90.0	3.28	180.1	3.28	270.0	3.28
	Chebwin	360.0	3.28	90.1	3.27	180.1	3.28	269.9	3.27
DTM - DSM	No filter	0.1	266.00	180.1	35.12	-89.9	74.49	180.0	36.53
	Bartlett	-179.9	266.00	0.1	372.68	0.1	266.00	180.1	266.00
	Chebwin	0.1	266.00	-90.0	266.00	180.0	266.00	90.0	266.00
DTM - DSM	No filter	-90.0	2.95	-89.9	1.55	270.0	2.95	270.1	2.95
	Bartlett	154.2	2.87	89.5	2.72	92.6	605.61	90.1	266.00
	Chebwin	-6.2	1.68	-89.9	266.00	182.7	2.15	90.1	266.00

(b)

6.1.3.4 Discussion

The test results presented in Chapter 6.1.3.3 provided the following insight into the performance of the LAFMT and how it compares to the standard FMT:

- The LAFMT delivered best results when the datasets were self-registered – all height maps auto-registered successfully in almost all test cases. The lowest success rate, 83%, was obtained when the Oshawa height maps were auto-registered without application of a scale offset.
- While the FMT failed to co-register the Oshawa and the simulated data entirely, the LAFMT succeeded in 15% and 26% of the test cases, respectively. However, the KN data co-registered correctly in 55% of the test cases.
- The LAFMT converged correctly slightly more often when the data input had the same scale in case of the Oshawa and simulated data. On the contrary, it succeeded more often when the data was scaled in case of the KN data. This confirms that the KN data provides the strongest phase signal which supports strong decorrelation of scale and angle in the CPS.
- Overall, the LAFMT performed best with the KN data while the FMT performed worst with the KN data. This confirms how important it is that the phase of the input signal be maintained inside the co-registration algorithm – the KN height maps provide a slightly smaller range of elevation (magnitudes) than the other two datasets, however the change in elevation occurs over shorter distances and thus creates a stronger phase signal than in the other datasets. Since the FMT considers only elevation magnitudes, it performed poorly with the KN signal, and slightly better with the Oshawa and simulated data, which provide a bigger range in elevation magnitude.
- While the application of a window filter improved the majority of FMT results, it improved the LAFMT results only in some cases and degraded it in other cases. On average, the LAFMT performed roughly the same with or without application of a window filter. This

finding confirms the primary purpose of the window filter, which is to minimize the “+” effect in the Fourier magnitude spectrum, which is only relevant for the FMT. The secondary purpose of the window filter is to reduce effects of tiling, which are small in case of the real-world point clouds.

- While the FMT succeeded to auto-register the Oshawa data in only in 50% of the cases when the data was scaled (Table 5.4), yet the LAFMT succeeded 100% (when the Bartlett filter was applied). Similarly, the FMT failed to register the KN data when a 90deg and 270deg and scale offset were applied (Table 5.7), yet the LAFMT succeed in all self-registration test cases.

In summary, the LAFMT performed much better than the FMT. This was expected since the height maps used in this dissertation (introduced in Chapter 3.0) are relatively flat and noisy and since the FMT considers only the magnitude of their phase spectra.

6.1.4 Co-Registration using Polar AFMT

While the LAFMT is the most frequently used “standard” version of the AFMT, it has its downsides, too. The main drawback is that log-polar sampling yields high point density around the origin of the coordinate system, and lower densities away from the centre. As Sellami and Ghorbel (2012) suggested, for image pairs that are only set apart by rotation, polar sampling can be used instead. By using polar mapping of the images, the Polar AFMT (PAFMT) reduces the aliasing error typically associated with the non-linear distribution of image samples along the scale axis.

6.1.4.1 The PAFMT Algorithm

The PAFMT is a simplified form of the LAFMT. Figure 6.7 below illustrates how the PAFMT can be derived from the LAFMT (illustrated in Figure 6.2).




Figure 6.7: Modification of the LAFMT into the PAFMT algorithm

The disembowelled version of Figure 6.7 is illustrated below.




Figure 6.8: The PAFMT algorithm


As outlined in Figure 6.8, the PAFMT algorithm comprises four steps:


1. Input of two image-rendered point clouds
2. Step A: The images are rendered into polar base where rotation is associated the x-axis. Optionally, if the surface profiles have quadratic footprint, a window filter can be applied before polar mapping to soften the image borders and reduce tiling effects. The same window masks as used in the FMT can be deployed.
3. Step B: The polar images are ingested into a POMF to isolate the rotational offset. The FFT inside the POMF perceives the incoming images as infinite periodic signal and thus must consider the ambiguous location of the correlation peak as referenced to the first and third quadrant. Thus, the POMF outputs two estimates for the angular offset, where one estimate offsets the other by 180 degrees.
4. Step C: The original “moving” input image is scaled and rotated one time by φ and one time by $\varphi+180$ to determine the best estimate for the set of $\Delta\varphi$.

Obviously, the property of scale invariance is lost in the PAFMT. Thus, in steps B and C, only the x-value of the CPS peaks (indicating the rotational offset) is considered.

If only subtle changes in image magnitude are expected, i.e., if the SNR of the phase spectrum is high, more than one of the CPS correlation peaks found inside the POMF (step B) may be considered. In that case, steps B and C are repeated for each candidate and the $\Delta\varphi$ that results the highest correlation peak in Step C is chosen as final candidate.

Figure 6.9 below exemplifies the three steps of the PAFMT algorithm:








Figure 6.9: Self-registration of a 30deg offset X5d4 DTM (pix) using PAFMT

Figure 6.9 exemplifies each step of the PAFMT algorithm for the self-registration of the X5d4 DTM with an angular offset of 30degrees. Cells a1,1 and 2 and a2,1 and 2 show the height maps in polar rendering. Here, the angular offset appears as translation along the x-axis. “Peak B” showed an approximate x-coordinate of 2150 before rotation (a1,1 and a1,2) and 1850 after rotation (a2,1 and a2,2). The CPS peaks highlighted with green arrows confirmed a shift of roughly 200pixels. As per equation (5.17), the angular offset of 30deg in Cartesian base relates to a full rotation of 360deg the same way as the translation of 200pixels in polar base relates to the full length of 2500pixels. Indeed, $30deg/360deg \approx 200pix/2500pix$ confirmed the 30deg rotation.

6.1.4.2 Properties

The PAFMT is the simplest form of the AFMT. It has the following properties

- Like the LAFMT, the PAFMT considers the phase information of the height maps. Therefore, it is less susceptible to noise and yield more robust registration results than the FMT, especially when low-resolution DEMs with low elevation contrast are co-registered.

- Like the LAFMT, the PAFMT yields a much sharper peak than the classical FMT "since the spectral phase preserves the location of objects but is insensitive to the image energy" (Chen *et al.*, 1994). Therefore, it should be able to retrieve rotation offsets within a larger range and/or with higher precision and/or with increased robustness.
- Since the PAFMT encompasses only one DFT, it is free of “+”-effects.
- As with the FMT, window filters can be used to reduce the effects of image-tiling (see 5.1.2.2.2), if the height maps have a square footprint. However, effects of image tiling should be small for height maps with small changes in elevation.
- Unlike log-polar sampling, polar sampling provides homogenous sampling all across the image which results in less information loss around the origin of the coordinate system.
- If the number of pixels in polar space is chosen to be equal the number of pixels of the cartesian image, then the angular coordinate is sampled uniformly and the precision on the angle is therefore uniform (Chen, 1993).
- Since the PAFMT provides relatively dense resampling of the ingested elevation data, it should perform better in recovery of angular offsets than the LAFMT.
- Since the CPS is no longer invariant to scale, multiple peaks that share an x-value (that presents the angular offset) can lead to the same solution.
- Since the CPS energy tends to spread across multiple valid peaks (illustrated in left column of Figure 6.10 below), the CPS plot of polar images tends to look noisier than the CPS plot of respective log-polar images (illustrated in the right column of Figure 6.10).




Figure 6.10: Co-registration of Oshawa and KN point clouds (pix), native scale and orientation, using PAFMT (left) and LAFMT (right)

As can be seen from Figure 6.10, the correlation energy distributes across more valid peaks in the CPS of polar-sampled images (left column) compared to the CPS graph of log-polar sampled imagery (right column): Since in this example both height maps have the same unit and orientation, multiple peaks along the first and last row of y-values in the CPS of polar images can represent the correct candidate (indicated by yellow ellipses), while

only one peak along the last row of y-values qualifies as potential candidate in the CPS of log-polar images.

The implications of these properties in conjunction with the review of numerical results will be discussed in the next chapter.

6.1.4.3 Numerical Results

Since log-polar rendering causes more agglomeration of map surface around the origin of the coordinate axis, and since multiple CPS peaks in polar rendering can represent the correct solution, it is expected that the PAFMT achieves the same or a better success rate in retrieving angular offsets than the LAFMT. To test this proposition, the real-world and simulated height maps were co-registered using the PAFMT. Table 6.4 summarizes the results below.

In Table 6.4, the same color coding was applied as before for the LAFMT results (6.1.3.3 above).

Table 6.4: Registration of real-world and simulated elevation profiles using PAFMT

		Oshawa Height Maps				KN Height Maps				Simulated Height Maps			
		0 ⁰	90 ⁰	180 ⁰	270 ⁰	0 ⁰	90 ⁰	180 ⁰	270 ⁰	0 ⁰	90 ⁰	180 ⁰	270 ⁰
Laser – Laser	No filter	0.0	90.0	-180.0	90.0	0.0	91.2	-180.0	271.2	0.0	0.0	90.0	180.0
	Bartlett	0.0	90.0	180.0	0.0	0.0	88.8	-165.2	271.2	0.0	90.0	180.0	270.0
	Chebwin	0.0	0.0	-180.0	0.0	0.0	88.8	-165.2	271.2	0.0	90.0	180.2	270.0
Tie - Tie	No filter	0.0	89.3	-180.0	270.7	0.0	92.2	-180.0	92.2	0.0	0.0	180.0	-180.0
	Bartlett	0.0	89.3	-180.0	273.7	0.0	87.8	-180.0	272.2	0.0	89.8	180.2	271.2
	Chebwin	0.0	92.2	-180.0	273.7	0.0	87.8	-180.0	272.2	0.0	-90.2	-0.2	271.2
Laser - Tie	No filter	352.3	-180.0	100.3	0.0	9.3	-16.4	166.0	14.0	0.0	0.0	-180.0	-180.0
	Bartlett	0.0	-180.0	0.0	0.0	0.0	177.7	-180.0	0.0	0.0	90	180.0	-49.2
	Chebwin	0.0	-180.0	0.0	0.0	0.0	-177.7	-180.0	0.0	0.0	90	180.0	270.0
Tie - Laser	No filter	90.7	-180.0	90.7	0.0	350.6	79.5	9.4	261.8	0.0	0.0	-180.0	-180.0
	Bartlett	175.7	0.0	0.0	-180.0	0.0	341.3	-147.3	-175.3	0.0	90	180.0	270.0
	Chebwin	180.0	0.0	0.0	-180.0	0.0	-180.0	-180.0	-180.0	0.0	90	180.0	270.0

The upper half of Table 6.4 shows the results where the PAFMT was used to self-register tie- and laser point clouds. The PAFMT failed in twelve instances (20+23+18 correct angles), i.e., in autocorrelation mode it converged correctly in 85% of the test cases. In comparison, the LAFMT converged correctly in 89% of the test cases (20+24+20 acceptable results) in auto-correlation mode when no scale was applied.

The lower half of Table 6.4 shows the cross-correlation results, which resemble practical geodetic applications, and are therefore more important in this assessment than the outcome of auto-correlation. Here, the PAFMT converged correctly in only 31 cases (3+13+20 acceptable angles), i.e., it showed a success rate of 50%. In comparison, the LAFMT converged correctly in 35% of the test cases without scale offset (3+10+12 acceptable pairs of scale and angle) and, considering its angle estimates only, the LAFMT showed a success rate of 36% (3+18+5 acceptable angles).

Table 6.5: Registration of real-world and simulated elevation profiles using PAFMT to scale 3.2808

		Oshawa Height Maps				KN Height Maps				Simulated Height Maps			
		0°	90°	180°	270°	0°	90°	180°	270°	0°	90°	180°	270°
Laser – Laser	No filter	0.0	-89.8	-180.0	-89.8	0.0	90.0	-180.0	90.0	0.0	-180.0	-180.0	-180.0
	Bartlett	0.0	0.0	-180.0	0.0	0.0	88.6	180.0	90.0	0.0	90.0	180.0	270.0
	Chebwin	0.0	0.0	-180.0	-180.0	0.0	88.5	180.0	90.0	0.0	90.0	180.0	270.0
Tie -Tie	No filter	0.0	-89.8	-180.0	-89.8	0.0	90.0	-180.0	270.0	0.0	-180.0	180.0	-180.0
	Bartlett	0.0	0.0	-180.0	0.0	0.0	90.0	180.0	270.0	0.0	90.0	0.0	270.0
	Chebwin	0.0	0.0	-180.0	-180.0	0.0	90.0	0.0	270.0	0.0	90.0	0.0	270.0
Laser - Tie	No filter	-90.0	0.0	90.0	-180.0	-166.7	11.8	168.2	-168.2	0.0	0.0	-180.0	-180.0
	Bartlett	90.0	-180.0	-90.0	0.0	0.0	-180.0	-180.0	175.6	0.0	0.0	180.0	-180.0
	Chebwin	0.0	90.0	-180.0	-90.0	0.0	0.0	180.0	180.0	0.0	-180.0	0.0	0.0
Tie -Laser	No filter	90.0	-180.0	270.0	0.0	-166.7	101.8	11.8	-10.3	0.0	-180.0	-180.0	-180.0
	Bartlett	270.0	0.0	90.0	-180.0	0.0	-97.4	180.0	0.0	0.0	-180.0	180.0	0.0
	Chebwin	0.0	270.0	0.0	-90.0	0.0	180.0	180.0	0.0	0.0	0	180.0	0.0

Table 6.5 shows the PAFMT results where both the template- and search point cloud were scaled by 3.2808. This case occurs when the search point cloud is referenced in feet and the

template point cloud is referenced in meters, and the template point cloud is upscaled so that the PAFMT can be run on two datasets of the same size. Ideally, the results should be identical to those obtained when both are referenced in the same unit. However, the interpolation induces loss of information, the padded areas can change size and the application of the window filter can therefore have slightly different effects. Thus, the results in Table 6.5 and Table 6.4 deviate from each other.

When the scaled point clouds were self-registered, the PAFMT failed in 15 instances (19+20+18 correct angles), i.e., in autocorrelation mode it converged correctly in 80% of the test cases. When the point clouds were co-registered, the PAFMT converged correctly in only 24 cases (4+10+11 acceptable angles), i.e., it showed a success rate of 33%. Thus, the PAFMT performed slightly worse when the point clouds were referenced in different units.

6.1.4.4 Comparison of LAFMT and PAFMT Angle Estimates

To better understand the difference of LAFMT and PAFMT results, the LAFMT results were replotted below and colour coded to highlight the differences. LAFMT rotation estimates that paired with a correct scale estimate and were superior to the respective angular estimate of the PAFMT were underlined blue with “_”. Likewise, instances where the PAFMT estimated an angular offset with more accuracy than the LAFMT, and the PAFMT paired with a correct LAFMT scale estimate, were highlighted with black “_”. Since for successful co-registration both the rotation and scale estimate must be correct, only rotation estimates that paired with a correct scale estimate were considered in the comparison.

Table 6.6: Comparison of LAFMT and PAFMT results for Oshawa height maps

		Oshawa DEMs, Scale = 1.0							
		0 ⁰		90 ⁰		180 ⁰		270 ⁰	
Laser - Laser	No filter	2.6	1.40	90.0	1.00	180.0	1.00	90.0	1.00
	Bartlett	-177.4	1.00	90.0	1.00	2.6	1.00	90.0	1.00
	Chebwin	2.6	1.00	90.0	1.00	180.0	1.00	272.6	1.00
Tie - Tie	No filter	3.0	1.03	95.2	1.00	150.2	1.00	-87.8	1.03
	Bartlett	3.0	1.00	92.2	1.00	-177.0	1.00	-87.8	1.00
	Chebwin	3.0	1.00	92.2	1.07	-177.0	1.00	267.7	1.03
Laser - Tie	No filter	280.3	1.03	-66.9	1.10	267.4	1.03	2.6	1.00
	Bartlett	0.0	1.40	2.6	2.89	2.6	2.89	185.1	2.34
	Chebwin	15.4	1.44	-2.6	1.83	105.4	1.58	20.6	1.35
Tie - Laser	No filter	169.5	1.45	267.4	1.40	349.5	1.45	84.9	1.27
	Bartlett	357.4	1.40	36.0	2.50	2.6	1.40	-144.0	2.50
	Chebwin	-140.6	1.03	321.4	1.13	2.6	1.03	-164.6	1.28

(a)

		Oshawa DEMs, Scale = 3.2808							
		0 ⁰		90 ⁰		180 ⁰		270 ⁰	
Laser - Laser	No filter	360.0	3.27	179.6	2.78	267.6	2.73	-89.8	3.29
	Bartlett	0.0	3.31	90.2	3.29	-179.2	3.27	-89.8	3.29
	Chebwin	1.6	3.27	-178.4	1.14	1.6	3.27	0.8	1.35
Tie - Tie/	No filter	358.1	3.27	90.0	3.27	179.1	3.31	270.0	3.27
	Bartlett	359.1	3.31	90.0	3.27	179.1	3.31	270.0	3.27
	Chebwin	1.9	3.22	90.0	2.06	-177.2	3.27	270.9	3.18
Laser - Tie	No filter	270.0	2.00	180.9	2.00	270.0	2.00	360.0	1.97
	Bartlett	90.0	1.97	180.9	2.00	270.0	2.00	360.0	1.97
	Chebwin	90.9	1.97	180.9	2.00	90.9	2.92	0.9	1.63
Tie - Laser	No filter	-179.2	4.60	-89.8	4.54	0.8	4.60	-89.8	4.54
	Bartlett	360.0	4.71	-89.8	4.54	180.0	4.71	-89.8	4.54
	Chebwin	-179.2	1.60	0.8	13.54	270.6	1.23	-179.2	13.54

(b)

Without scale offset, Table 6.6a, the LAFMT found 3 (1+2) angles that the PAFMT missed (Table 6.4), yet the PAFMT produced 8 (4+4) angles that were missed by the LAFMT for the Oshawa point clouds. Thus, the PAFMT outperformed the LAFMT by estimating an additional 5 angles (28 angles versus 23 angles) correctly.

With application of scale offset, Table 6.6b, the LAFMT found 4 (0+4) angles that the PAFMT missed (Table 6.5), yet the PAFMT located 6 (2+4) angles correctly that were estimated incorrectly by the LAFMT.

Overall, the PAFMT outperformed the LAFMT slightly by estimating 2 (5) more angles correctly when (no) scale offset was applied, respectively.

Without scale offset, Table 6.7(a), below, the LAFMT estimated 8 (1+7) angles correctly, where the PAFMT failed for the KN data. The PAFMT performed better in only 1 instance. So, without scale offset the LAFMT outperformed the PAFMT by estimating an additional 7 angles (38 angles versus 31 angles) correctly. With scale offset, Table 6.7b, below, the LAFMT resulted 12 (1+11) angles correctly that the PAFMT missed, and the PAFMT found 3 (1+2) angles that were missed by the LAFMT. Thus, with scale offset, the LAFMT outperformed the PAFMT by estimating an additional 9 angles (39 angles versus 30 angles) correctly.

In total, the LAFMT outperformed the PAFMT and estimated an additional 9 (8) angles, that also paired with a correct scale estimate, when (no) scale offset was applied, respectively.

Table 6.7: Comparison of LAFMT and PAFMT results for KN height maps

		KN DEMs, Scale = 1.0							
		0°		90°		180°		270°	
Laser – Laser	No filter	9.9	1.00	88.8	1.00	-150.4	1.00	286.0	1.00
	Bartlett	4.9	1.05	93.7	1.00	180.0	1.05	276.2	1.00
	Chebwin	4.9	1.05	93.7	1.00	180.0	1.05	276.2	1.00
Tie -Tie	No filter	24.5	1.00	114.1	1.00	-163.6	1.33	281.0	1.00
	Bartlett	8.2	1.00	96.6	1.00	-171.8	1.00	272.2	1.08
	Chebwin	8.2	1.00	101.0	1.00	-171.8	1.00	272.2	1.08
Laser -Tie	No filter	360.0	1.53	37.4	1.05	-133.2	1.53	243.1	1.40
	Bartlett	9.4	1.00	72.5	1.70	-175.3	1.00	275.8	1.33
	Chebwin	4.7	1.00	95.8	1.00	-175.3	1.00	275.8	1.00
Tie -Laser	No filter	-175.3	1.60	360.0	1.15	360.0	1.94	180.0	1.85
	Bartlett	4.7	1.00	93.5	3.94	166.0	1.53	271.2	4.33
	Chebwin	346.0	1.39	84.2	1.60	166.0	1.53	-105.2	1.53

(a)

		KN DEMs, Scale = 3.2808							
		0°		90°		180°		270°	
Laser – Laser	No filter	360.0	3.30	89.6	3.30	182.2	3.30	-6.7	2.93
	Bartlett	360.0	3.30	92.6	3.30	182.2	3.30	271.9	3.30
	Chebwin	3.0	2.30	91.1	3.11	-178.5	2.93	271.9	3.30
Tie -Tie	No filter	2.6	3.18	90.0	3.18	-177.4	3.18	272.6	3.28
	Bartlett	10.6	3.28	92.6	3.18	-177.4	3.18	-87.4	3.08
	Chebwin	5.3	3.49	92.6	3.08	-177.4	2.99	270.0	3.08
Laser -Tie	No filter	344.1	1.59	62.3	1.66	232.9	1.50	-44.7	1.51
	Bartlett	2.6	1.92	90.7	3.42	272.6	3.15	-79.8	3.32
	Chebwin	2.6	2.31	105.6	2.07	-177.4	1.92	-90.7	1.88
Tie -Laser	No filter	360.0	4.91	1.5	22.73	180.0	4.91	-178.5	22.73
	Bartlett	2.2	4.91	90.0	2.53	182.2	4.91	270.0	3.09
	Chebwin	-85.9	2.45	90.0	2.53	-148.6	1.64	-178.5	2.39

(b)

Table 6.8: LAFMT results, in comparison to the PAFMT results for simulated height maps

		Simulated DEMs, Scale = 1.0							
		0 ⁰		90 ⁰		180 ⁰		270 ⁰	
DTM - DTM	No filter	10.2	1.00	90.0	1.00	179.8	1.00	270.0	1.00
	Bartlett	10.2	1.00	91.2	1.00	180.0	1.01	270.4	1.02
	Chebwin	10.2	1.00	0.2	1.00	180.0	1.01	89.4	1.01
DSM - DSM	No filter	0.2	1.30	90.0	1.00	183.3	1.00	270.0	1.00
	Bartlett	10.2	1.00	90.0	1.00	180.2	1.00	270.0	1.00
	Chebwin	10.2	1.00	90.0	1.00	180.2	1.00	270.0	1.00
DTM - DSM	No filter	11.4	1.00	100.6	1.00	187.9	1.00	280.2	1.00
	Bartlett	87.5	2.49	79.0	2.21	-88.3	4.29	7.3	2.92
	Chebwin	261.7	2.71	173.4	3.70	179.2	2.67	-78.6	2.91
DTM - DSM	No filter	6.4	1.00	97.9	1.00	187.9	1.00	281.8	1.00
	Bartlett	103.3	1.40	-179.2	2.34	-62.6	2.47	-56.6	2.53
	Chebwin	69.3	1.02	270.6	1.65	158.9	2.49	33.0	2.43

(a)

		Simulated DEMs, Scale = 3.2808							
		0 ⁰		90 ⁰		180 ⁰		270 ⁰	
DTM - DTM	No filter	0.1	3.28	89.9	3.28	180.1	3.28	269.9	3.28
	Bartlett	360.0	3.27	89.9	3.27	180.0	3.27	269.9	3.27
	Chebwin	360.0	3.27	89.9	3.27	180.1	3.27	269.9	3.27
DSM - DSM	No filter	0.1	3.28	90.0	3.28	180.1	3.28	270.0	3.28
	Bartlett	360.0	3.28	90.0	3.28	180.1	3.28	270.0	3.28
	Chebwin	360.0	3.28	90.1	3.27	180.1	3.28	269.9	3.27
DTM - DSM	No filter	0.1	266.00	180.1	35.12	-89.9	74.49	180.0	36.53
	Bartlett	-179.9	266.00	0.1	372.68	0.1	266.00	180.1	266.00
	Chebwin	0.1	266.00	-90.0	266.00	180.0	266.00	90.0	266.00
DTM - DSM	No filter	-90.0	2.95	-89.9	1.55	270.0	2.95	270.1	2.95
	Bartlett	154.2	2.87	89.5	2.72	92.6	605.61	90.1	266.00
	Chebwin	-6.2	1.68	-89.9	266.00	182.7	2.15	90.1	266.00

(b)

Without scale offset, Table 6.8a, above, the LAFMT failed only in 2 instances where the PAFMT converged correctly. When the data was co-registered, the PAFMT resulted 3 (2+1) angles correctly that the LAFMT estimated incorrectly. Thus, the PAFMT outperformed the LAFMT and estimated an additional 13 angles (43 angles versus 30 angles) correctly.

When the simulated data was self-registered with scale offset, Table 6.8b, above, both LAFMT and PAFMT showed 100% success. When it was co-registered, the LAFMT resulted 2 angles correctly that the PAFMT missed, and the PAFMT produced 3 angles correctly that were estimated incorrectly by the LAFMT.

In total, the PAFMT outperformed the LAFMT angle estimates slightly as it estimated an additional 3 (3) angles, that also paired with a correct LAFMT scale estimate, when (no) scale offset was applied, respectively.

6.1.4.5 Discussion

At first sight, considering the success rate of all three datasets combined, the PAFMT angle estimates seemed to outperform the LAFMT slightly (47% versus 44% correct co-registration angles). However, as the detailed comparison above showed, the PAFMT did not deliver better rotation estimates in all scenarios:

- For the co-registration of the Oshawa data, the PAFMT delivered only 5 correct angles, and the LAFMT succeeded in 3 instances, i.e., the PAFMT performed only slightly better than the LAFMT.
- In the co-registration tests with simulated data, the PAFMT delivered 18 correct angles, whereas the LAFMT succeeded in only 13 instances. Here, the PAFMT recovered only 1 angle that the LAFMT had missed, and the remaining 17 and 12 angles that the PAFMT and LAFMT had estimated correctly, respectively, resulted from different test combinations.

- In case of the KN data, surprisingly, the LAFMT clearly showed better co-registration performance: the LAFMT delivered a total of 18 correct angles, whereas the PAFMT delivered only 8 correct rotation estimates. Also, the LAFMT estimated 7 angles correctly that the PAFMT had missed, yet the PAFMT recovered only 1 angle that the LAFMT missed. Thus, contrary to the co-registration results gained with simulated data, the polar and log-polar angle estimates did not complement each other, but rather did the LAFMT clearly outperform its polar counterpart with the KN data.

Thus, one can conclude that the LAFMT performed best with the KN point clouds, whereas the PAFMT performed better in case of the Oshawa and simulated data. From these observations, one can conclude further that the equal sampling distance of polar rendering (versus non-even distance of log-polar sampling) is not the primary success-driving factor for PAFMT co-registration, but rather

- 1) the characteristics of the height profile (low SNR) and
- 2) the characteristics of the CPS (polar vs log-polar CPS).

RE 1: While the KN height maps have the smallest range in elevation, the change in elevation occurs over shorter distances than in the other test sites. This circumstance yields a stronger phase signal, which allows for stronger cross-correlation. Similarly, the simulated data provides higher point density than the real-world data, which results in more distinct pronunciation of the surface signal.

RE 2: Since in polar rendering the y-component does not provide a direct physical interpretation, multiple peaks that share the x-value can yield the correct solution, and thus, may increase the chance for success when datasets don't produce strong correlation ("hypothesis 1"). Contrary, there may be instances where the PAFMT fails to deliver an accurate rotation estimate, yet the LAFMT may succeed since it considers the combination of scale and rotation in the CPS, and if the data has low SNR (as many large-scale airborne geodetic high maps have), cross-correlation may occur in the CPS, and a strong scale signature may "pull" the rotation estimate favorably to a more accurate rotation offset ("hypothesis 2").

Many datasets would have to be evaluated to create strong proof both hypotheses. Especially “hypothesis 1” is difficult to prove, since, as shown in Chapter 5.1.2 and APPENDIX F, stronger correlation of the data may not always lead to a CPS with fewer, more assertive CPS peaks, nor does the pronunciation of the 20 tallest CPS peaks allow to make conclusions about successful convergency. However, for the given datasets, in self-registration mode, the PAFMT solved only 78% of all angles (Oshawa 73%, KN 90%, Simulated 73%) while the LAFMT delivered 92% correct angle estimates and all of them paired with a correct scale estimate (Oshawa 81%, KN 98%, Simulated 96%). While these results may not allow for strong proof of the potential “pull”-effect of y- and x-components (scale and rotation), they support “hypothesis 2” and allow for good explanation and characterization of the test results.

In summary, the most important discovery of this analysis was that the PAFMT co-registration outperformed the LAFMT rotation estimates slightly for the Oshawa and simulated data (very low SNR), yet the LAFMT delivered slightly better rotation estimates for the KN data which had stronger phase signal, while it could also recover scale.

6.1.5 Co-Registration using Extended AFMT

The following sections introduce the Extended version of the AFMT (EAFMT), which can recover translation offsets in addition to rotation-, and scale- offsets.

6.1.5.1 The EAFMT Algorithm

In its native formulation, the standard AFMT can only recover scale and rotation. In 2005, Zokai and Wolberg presented an Extended version of the log-polar AFMT (EAFMT) that applies a LAFMT and POMF recursively to multiscale image pyramids to recover the complete set of similarity parameters. This iterative concept can be applied to both LAFMT and PAFMT, however, for simplification, only the Extension of the LAFMT will be discussed further; Also, the

expression EAFMT will be used synchronously with the extended version of the LAFMT. The EAFMT is outlined Figure 6.11 below:




Figure 6.11: The EAFMT algorithm

As outlined in Figure 6.11, the EAFMT algorithm for full recovery of all four similarity parameters (scale, rotation, and x-, y- translation) comprises five steps:

1. Input of two low-resolution image-rendered point clouds
2. Step A: The low-resolution images are transferred into log-polar space where rotation and scale de-couple and associate with the x- and y- axis, respectively. Optionally, if the data has a square footprint, a window filter can be applied before the log-polar mapping to reduce effects of tiling. The same window masks as used in the FMT (introduced in Chapter 5.1.2 and described in detail in APPENDIX E) can be deployed.
3. Step B: The log-polar images are ingested into a POMF to retrieve rotation and scale, which appear as a set of translation parameters. The FFT inside the POMF perceives the incoming

images as infinite periodic signal and thus must consider the ambiguous location of the correlation peak as referenced to the first and third quadrant. Thus, the POMF outputs two estimates for the angular offset, where one estimate offsets the other by 180 degrees.

4. Step C: The “moving” input image is scaled and rotated one time by φ and one time by $\varphi+180$.
5. Step D: The images are ingested into a POMF in cartesian presentation to retrieve the translation offsets.
6. Step E: The parameter set that results the highest CPS peak in step D is applied to the imagery at next higher resolution and the co-registration process (steps 1-5) is repeated. The co-registration parameters are refined incrementally until the imagery has been co-registered in its original resolution.

In short, the EAFMT, as proposed by Zokai and Wolberg (2005), is a recursive application of the AFMT and POMF on multiscale image pyramids.

6.1.5.2 Discussion

As shown in 6.1.3.3, and 6.1.4.3, the AFMT delivered better co-registration results than the standard FMT (5.1.2.2.4). However, unlike the FMT, the AFMT cannot recover translation by default. Therefore, the EAFMT was introduced to show that this restriction can be overcome, if necessary.

However, unlike the AFMT, the FMT was purposely designed to recover all four similarity parameters. So even though the EAFMT aims at the same, it requires that the datasets overlap to a large degree, and the range of translations that can be recovered is therefore smaller. Also, for it to work successfully, the EAFMT requires strongly pronounced elevation profiles so that height differences are clearly present at every layer of the image pyramid. There are doubts as to whether the EAFMT is appropriate for co-registering low-resolution airborne tie point- and laser point cloud data, particularly when the height maps exhibit minimal elevation contrast. The same reservations would extend to using an extended POMF (see Chapter 5.1.1.4). While the recovery of translation offsets using the existing EAFMT, or the extended POMF, is comparably of minor importance, it adds unnecessary complexity when combined with the study of scale and rotation

recovery. Moreover, the relatively high failure rate of the PAFMT and LAFMT carry over into the EAFMT and limit its overall performance. It is therefore important to understand how the LAFMT performance could be improved, before investigating the numerical capabilities of the EAFMT. Thus, new ways to improve the performance of the LAFMT will be explored next in Chapter 6.1.3.

6.1.6 Conclusions

In this section, the AFMT was introduced and tested as a potential alternative to the FMT. Since the AFMT preserves the phase information of the height maps, and the Fourier phase carries more information than the Fourier magnitude, it was expected that the AFMT delivers superior results.

Two forms of the AFMT were examined: The PAFMT and the LAFMT. As with the FMT, only scale values 1.0 and 3.2808 were tested when correlating the height maps, (since - as elaborated in Chapter 5.1.2.3 - this allowed for circumventing the reduction of the search image and conducting the co-registration with an enlarged template image instead, whenever the log-polar AFMT estimated a scale offset significantly smaller than 1.0.)

As expected, both the PAFMT and LAFMT showed a higher success rate than the FMT, since the height maps introduced in Chapter 3.0 are relatively flat and noisy and since the FMT considers only the magnitude of their phase spectra:

The LAFMT delivered better rotation estimates than the PAMFT for the KN data, which could be attributed to the stronger phase signal of the KN data, and the PAFMT surpassed the LAFMT with the Oshawa and simulated data. Overall, the LAFMT proofed to be the preferred method for co-registration since it delivers both rotation and scale and since its rotation estimates lagged the PAFMT angle estimates mostly only slightly.

Lastly, the EAFMT was introduced as an extended AFMT that can – like the FMT – recover not only rotation and scale, but also translation. But since the translations are recovered with Cartesian height images in an internal POMF, the range of translations that can be recovered with the EAFMT is smaller than that of the FMT. The most important drawback of the EAFMT, and also LAFMT and PAFMT, is however that they were not optimized for the presence of two scale offsets (3.2808 or its inverse), i.e., the presence of two metric units (feet and meter). Knowing that the laser and tie point clouds can only be either referenced in the same unit or offset by a scale of

3.2808 (or its inverse), better optimized algorithms could be designed that account for this condition. Thus, new ways to improve the performance of the AFMT will be explored next.

6.2 The Quasi Log-Polar Analytical Fourier Mellin Transform

In the late 1970s, the FMT was introduced for pattern recognition and image processing, however the numerical estimation of the Mellin integral implied numerical difficulties and the FMT did not hold the phase information. In 1994, Ghorbel proposed the AFMT to solve both of the problems, but it became no longer invariant to the translation. Soon, the EAFMT was further introduced to gain back the translation invariance (Zokai and Wolberg, 2005).

With focusing on the recovery of the rotation and scale – as mentioned in Chapter 6.1.5.2, only if the recovery of rotation and scale offset look promising, further investigation with EAFMT could occur – , the co-registration experiments in Sections 6.1.3 and 6.1.4 examined two forms of the AFMT: the LAFMT and PAFMT. Based on the design of the PAFMT, it should deliver better rotation estimates than the LAFMT. The tests confirmed that the PAFMT resulted overall slightly more correct rotation estimates in the co-registration than the LAFMT from the height maps with very low SNR (PAFMT: 14.6% (Oshawa), 60.4% (KN), 64.5% (Simulated); LAFMT: 14.6% (Oshawa), 75.0% (KN), 41.7% (Simulated)), but the LAFMT additionally recovered the scale offsets (correct LAFMT scale estimates: 79.2% (Oshawa), 66.6% (KN), 35.4% (Simulated)), which is required in the geospatial co-registration applications, targeted in this dissertation.

Moreover, the LAFMT assumes that the point clouds can be offset by any possible scale difference. Yet, within this framework, only the scale difference 3.2808 (and its inverse 0.3048) needs to be considered. This circumstance led to the idea of combining the PAFMT and LAFMT, pairing the high scoring rate of the LAFMT scale estimates with the PAFMT angle estimates, and adapting the LAFMT's scale output for geospatial applications. Thus, there are two ways in which a LAFMT with the reduced search space for scale is combined with a PAFMT:

1. The LAFMT is used to obtain both the scale and rotation in case the point clouds are offset by scale, and the PAFMT is performed only if both of the point clouds are referenced in the same units.

2. The LAFMT is used to obtain the scale offset only and the PAFMT is used to obtain the angular offset.

In the ensuing sections, the respective combinations will be referred to as “*Quasi-Log-Polar AMFT*” (QAFMT) algorithms. As will be shown in Section 6.2.1, the first variant, above, of the QAFMT intertwines the rotation estimate of the internal LAFMT and PAFMT more closely than the second variant, above. Therefore, it is referred to as “*Tightly-Coupled QAFMT*” (QAFMT_TC). In the second version of the QAFMT, the LAFMT and PAFMT are more capable of functioning independently. Thus, it is referred to as “*Loosely-Coupled QAFMT*” (QAFMT_LC).

6.2.1 Co-Registration using Tightly-Coupled QAFMT

6.2.1.1 The QAFMT_TC Algorithm

In the framework of this dissertation, only two scale offsets of 3.2808, and its inverse 0.3048, can occur, since the datasets can only be set apart by a difference in unit, i.e., reference in meter or feet. Thus, the QAFMT_TC started out by conducting a LAFMT where only the CPS peaks within a small buffer zone along the y-axis were considered in the search for the scale and angular offsets. After the LAFMT with the reduced search space settled on a CPS peak, the LAFMT output is combined with a PAFMT to obtain a refined angular offset. Building on “hypothesis 2” presented in Chapter 6.1.4.5, this angular offset should have the same or better accuracy than that provided by the PAFMT, in case the height maps possess a low SNR and are offset by scale. The PAFMT should, by design, estimate the angle offsets better than the LAFMT with the height maps at the same scale. Figure 6.12 outlines the QAFMT_TC algorithm.




Figure 6.12: The QAFMT_TC algorithm

In the tightly-coupled QAFMT, a LAFMT (steps A, B, C, D depicted in black in Figure 6.12) is undertaken first to decide whether the height maps are referenced in the same unit or not. If the LAFMT settles a scale offset “near” to 3.2808, then the scale offset is fixed to the nominal value of 3.2808, and the respective angular estimate from the same LAFMT is considered as the final estimate of the angular offset (since the previous LAFMT and PAFMT results showed that in presence of a scale offset the LAFMT angles tend to be more accurate than the respective PAFMT angles from the low SNR data, see Chapter 6.1.4.5 “hypothesis 2”). If the LAFMT results in a scale “near” to 1.0, it is assumed that the height maps were referenced in the same unit, and then the scale offset is fixed to the nominal value of 1.0, and the angular offset is derived from the PAFMT (steps F, G, H depicted in blue in Figure 6.12).

If a scale estimate by the LAFMT deviates strongly from the nominal value of 3.2808 or 1.0, the respective CPS peak is ignored and the next biggest CPS peak out of the 20 tallest CPS peaks will be analyzed (the same way it is done with the LAFMT and PAFMT). Furthermore, if the QAFMT_TC reaches the 20th peak in the scale search and locates an unplausible value far larger than 3.28, the scale is fixed to the nominal value of 3.2808 (this is unique to the QAFMT), as it is assumed that the QAFMT_TC did not converge even though a scale offset is present (otherwise it would have settled near to 1.0).

The QAFMT_TC requires to define a threshold for what values are considered “near” to the nominal scale offset of 3.2808 (and 0.3048) and 1.0. Thus, in the subsequent numerical evaluation of the QAFMT_TC, the output of the internal LAFMT was evaluated against the thresholds outlined in Table 5.1 (which were used earlier to evaluate the FMT, LAFMT, and PAFMT results) to decide if the datasets were offset by scale and/or rotation. Since Table 5.1 outlines thresholds for a 20% and 50% confidence interval, two sets of co-registration results were generated – one for each threshold. One set of results describes the QFMT_TC output where the scale estimates deviated less than 20% from the nominal values of 3.2808 and 1.0, while the second set of results describes the QAFMT_TC output gained with the respective 50% internal threshold for the scale estimates.

Figure 6.12 presents only results for the application of the scale offset 3.2808 and 1.00. As with the evaluation of the FMT, LAFMT, and PAFMT, the application of the scale offset of 0.3048 was circumvented and only the scale offsets of 3.2808 and 1.0 are applied in the QAFMT_TC. So, when the target point cloud is referenced in meters and the search point cloud is referenced in feet, and the LAFMT estimates a scale offset near 0.3, the scale can be fixed to 0.3048 and this scale number can be traced internally, and the PAFMT is executed with a native search point cloud and a 3.2808-up-scaled target point cloud instead of a 0.3048-downscaled search point cloud and native target point cloud. This tightly-coupled strategy has three advantages:

- Firstly, the loss of information from down sampling the search point cloud is avoided and the PAFMT can estimate the remaining rotation offset with a higher accuracy from a search- point cloud and a target point cloud that are sampled at a higher resolution.
- Secondly, it is challenging to retrieve the exact center location of CPS peaks that represent the scales <1.0 , since their peak is found in a sub-pixel location only within the first row of pixels in the CPS. While there are methods to accurately retrieve CPS peak locations at the sub-pixel accuracy (Wisetphanichkij and Dejhan, 2005), the suggested method of peak search in the CPS of inversely upscaled images is more robust since CPS locations across multiple rows of the CPS matrix represent scales of approximately 3.2808. Also, on the contrary to the sub-pixel CPS peak search, this strategy does not require additional computation.

- Thirdly, the QAFMT_TC does not require to rescale the imagery, only fixation of the scale estimate to 1.0 or 3.2808, unless the scale estimate was near to 0.3 (and then the reference image needs to be up-scaled).

Figure 6.13 illustrates the search in the QAFMT_TC for the peaks that represent a scale offset of 3.2808. In this example, the Oshawa laser point cloud served as the reference for the co-registration of the Oshawa tie point cloud that was offset by a scale of 3.2808.




Figure 6.13: Search of CPS peaks for co-registration of Oshawa laser- and tie point clouds (pix) in QAFMT_TC

Figure 6.13 illustrates the search of CPS peaks in the QAFMT_TC. The thumbnails in the first row indicate what the native padded laser point cloud and the 3.2808-upscaled tie point image looked like. The second row of thumbnails shows the same height maps in the log-polar (plot A1) and polar (plot A2) rendering – the form in which they entered the CPS of the QAFMT-internal LAFMT and PAFMT, respectively. The CPS plot in the plots (A1) and (A2) were obtained from

the LAFMT and PAFMT, respectively, in the QAFMT_TC with the application of the Chebyshev mask. As shown in Table 6.1, the LAFMT estimated an incorrect rotation offset of 90.9deg and a scale offset of 1.97 – the red arrow points to the respective CPS peak. Since this scale estimate deviated less than 50% from the nominal value of 3.2808, it was considered correct and replaced with the nominal value in the QAFMT_TC.

Evaluating the scale estimates against a confidence interval is equivalent to searching CPS diagram for the peaks that fall within a band (of 50% or 20% deviation) around the row of y-values that represents the scale offset of 3.2808. According to (6.8), this band situates at the rear of the CPS plot. In Figure 6.13 (A1) and the subsequent plots, the approximate location of this search space is indicated by an orange ellipse. Since the height maps in this example had the same orientation, the correct peak must situate along the first or last column (the frequency spectrum is infinitively repetitive) of x-values – in Figure 6.13 (A1), the CPS peak that represents the correct scale and rotation offset is marked with a green arrow.

With the LAFMT, only one peak within the band of acceptable y-values qualifies as the potential candidate in the CPS of log-polar images, while the PAFMT considers multiple peaks along the first and last column of x-values in the CPS of polar images as correct candidate. The respective areas are highlighted with yellow ellipses in Figure 6.13 (A2). As shown in Table 6.4, the PAFMT estimated a correct rotation offset of -10.3deg. The CPS peak that corresponds to this solution is located within the ellipses marked in orange on the right of the CPS graph in Figure 6.13 (A2).


6.2.1.2 Properties

Since the QAFMT_TC combines the LAFMT and PAFMT, it has their excellent characteristics:

- It is less susceptible to noise and delivers more robust registration results than the standard FMT, since it syndicates analytical AFMTs, – especially when the low-resolution DEMs with low elevation contrast (low SNR) are co-registered – because it utilizes the phase information of the height maps.

- The CPS of the LAFMT and PAFMT in the QAFMT_TC yield sharper peaks than the CPS in the classical FMT "since the spectral phase preserves the location of objects but is insensitive to the image energy" (Chen *et al.*, 1994). Therefore, the QAFMT_TC is able to retrieve the rotation and scale offsets within a larger range and/or at a higher precision and/or with an increased robustness than the FMT.
- It is free of “+”-artifacts because the QAFMT_TC encompasses only one DFT.
- Although window filters can be used to reduce the effects of image-tiling (see 5.1.2.2.2), - unlike in the FMT - the effects of image tiling are here small for height maps with small changes in elevation.

In Section 6.1.4.5, the “hypothesis 2” was made that a scale offset (CPS y- component) may have a favorable “pull” effect on the location of the corresponding x-component (rotation) in the log-polar CPS for data with low SNR. Thus, the presence of a scale offset may have a positive impact on the QAFMT_TC co-registration results when the log-polar rendering is used. In support of this theory, Figure 6.14 and Figure 6.15 illustrate that the log-polar CPS is slightly less noisy and has more distinct and pronounced peaks than the corresponding polar CPS, especially when a scale offset is prominent.






Figure 6.14: CPS of log-polar (left side) and polar (right side) sampled KN point clouds (pix), rotated by 90deg and scaled by 3.2808

Figure 6.14 illustrates the of CPS of the KN height profiles with different offsets. The thumbnails in the first row of Figure 6.14 (A1) and Figure 6.14 (A2) indicate what the native padded laser- and tie point cloud looked like. The second row of thumbnails show the height maps in the log-polar (plot A1) and polar (plot A2) rendering - the form in which they entered the CPS. The large CPS plot in the plots (A1) and (A2) were obtained from the LAFMT and PAFMT, respectively, in the QAFMT_TC with the application of the Chebyshev mask. As shown in Table 6.1, the LAFMT estimated a correct rotation offset of 4.7° and a correct scale offset of 1.00. However, as shown in Table 6.4, the PAFMT incorrectly estimated the rotation offset as 51.4° .

Figure 6.14 (B1) and Figure 6.14 (B2) illustrate the CPS after a 90deg rotation had been applied to the tie point cloud. Here, the LAFMT correctly estimated the rotation angle of 95.8deg and the scale offset of 1.00, however the PAFMT incorrectly estimated the rotation offset as 177.7deg .

The thumbnails in the first row of Figure 6.14 (C1) and Figure 6.14 (C2) indicate what the native padded laser point cloud and the 3.2808-scaled tie point cloud looked like, respectively. The second row of thumbnails show the height maps in the log-polar (Figure 6.14 (A1)) and polar (Figure 6.14 (A2)) rendering. The large CPS plot in Figure 6.14 (C1) and (C2) were obtained from the LAFMT (PAFMT) in the QAFMT_TC with the application of the Chebyshev mask, respectively. As shown in Table 6.1, the LAFMT estimated a correct rotation offset of 2.6° and a correct scale offset of 2.31. As shown in Table 6.4, the PAFMT estimated an incorrect rotation offset of 51.4deg .

Cells D1 and D2 illustrate the CPS after a 90° rotation had been applied to the tie point cloud. The log-polar CPS showed again more pronounced peaks and less noise than its polar counterpart. While the log-polar CPS yielded a correct rotation offset of 105.6° and a correct scale offset of 2.07, the polar CPS led to an incorrect angular offset of 177.7° .

To further support “hypothesis 2”, the respective CPSs with 90deg rotation and 3.2808 scale offset are shown in Figure 6.15 below.




Figure 6.15: CPS of log-polar (left side) and polar (right side) sampled Oshawa and simulated point clouds (pix), rotated by 90deg and scaled by 3.2808

Figure 6.15 A1 (A2) show the log-polar (polar) CPS for the Oshawa data, where the tie points were offset by a 90° rotation and 3.2808 scale. As before, the plots were obtained from the LAFMT (PAFMT) inside the QAFMT_TC with application of MATLAB's "Chebwin" Chebyshev mask (MathWorks, 2019), respectively. As shown in Table 6.1, the LAFMT estimated an incorrect rotation offset of 180.9° and a correct scale offset of 2.00. As shown in Table 6.4, the PAFMT estimated a correct rotation offset of 74.6° .

Figure 6.15(D1) and (D2) are the respective CPS plots from the simulated X5d4 and X5d5 data. Again, the log-polar CPS showed more pronounced peaks and lower noise than the polar version. the log-polar CPS resulted an incorrect angular offset of -89.9° and an incorrect scale of 1.55. The polar CPS resulted an incorrect angular offset of 141.2° , too.

6.2.1.3 Numerical Results

The QAFMT_TC requires to define a threshold for what to be considered as “near” to the nominal scale offsets of 3.2808 (and its inverse, 0.3048) and 1.0. Since Table 5.1 outlines thresholds for a 20% and 50% confidence interval, there are two sets of the co-registration results: One describes the QFMT_TC output where the scale estimates deviated less than 20% from the nominal values of 3.2808 and 1.0, while the second describes the QAFMT_TC output gained with the respective 50% internal threshold for the scale estimates.

Since the 20% threshold is tighter, the QAFMT_TC results obtained with the margin of 20% for the scale estimate contain more rotation estimates from the PAFMT (and fewer pairs of scale- and rotation estimates from the LAFMT) than those obtained with the internal threshold of 50%. Furthermore, since QAFMT_TC starts out with an internal LAFMT to identify the scale offset, it is expected that the scale and rotation estimates from the QAFMT_TC and from the standalone LAFMT should be identical in many cases. Unless the scale estimate deviates more than 20% or 50% from 1.0, the rotation estimate is determined by the PAFMT, and the rotation estimates by the QAFMT_TC are identical to those of the regular PAFMT. Supporting the earlier mentioned “hypothesis 2”, in presence of a scale offset that falls within the predefined margins, there may be instances where the LAFMT in the QAFMT_TC delivers better rotation estimates than what the PAMFT would deliver, especially when the data has low SNR.

In the preceding sections, the output of the QAFMT_LC was evaluated against the thresholds outlined in Table 5.1 to decide if the point clouds were offset by scale and/or rotation. The same color coding as in 6.1.3.3 above were applied. Instances where the LAFMT (results listed in Chapter 6.1.3) delivered an accurate rotation- or scale estimate and the QAFMT_TC failed were underlined with a black wave (“~~~~”). Similarly, numbers underlined with a blue wave (“~~~~”) indicate that QAFMT_TC delivered a better rotation- or scale estimate than the LAFMT. Likewise, a black underline (“__”) indicates that the PAFMT delivered an accurate rotation estimate where the QAFMT_TC failed, and a blue underline (“__”) indicates that QAFMT delivered an accurate rotation estimate where the PAFMT failed. Lastly, a dotted black line (“....”) indicates that both the standalone PAFMT and LAFMT resulted better rotation estimates than the QAFMT_TC, and a blue dotted line (“....”) indicates that the angle estimates of the QAFMT_TC superseded those of

the PAMFT and LAFMT. Whenever the internal LAFMT yielded a scale estimate larger than the 50% margin, the scale value was fixed at the nominal 3.28 and highlighted in purple. All other rotation and scale estimates were left in original format since here both the QAFMT_LC, LAFMT, and PAFMT delivered either identical or equivalent results. Lastly, since earlier tests showed that the interpolation method has little effect on the outcome, bilinear interpolation was used in all QAFMT_TC tests.

Table 6.9: Registration of Oshawa height maps using QAFMT_TC with the confidence interval of 20%

		Oshawa DEMs, Scale = 1.0, 20% confidence							
		0 ⁰		90 ⁰		180 ⁰		270 ⁰	
Laser – Laser	No filter	0.0	1.40	<u>-90.0</u>	1.00	-180.0	1.00	<u>270.0</u>	1.00
	Bartlett	<u>0.0</u>	1.00	90.0	1.00	<u>-182.6</u>	0.97	<u>0.0</u>	1.00
	Chebwin	0.0	1.00	<u>0.0</u>	1.00	-180.0	1.00	<u>0.0</u>	1.00
Tie - Tie	No filter	0.0	1.03	89.3	1.00	-180.0	1.00	270.7	1.00
	Bartlett	0.0	1.00	89.3	1.00	-177.0	0.97	273.7	1.00
	Chebwin	0.0	1.00	92.2	1.07	-177.0	0.97	273.7	1.03
Laser - Tie	No filter	352.3	0.97	<u>-180.0</u>	1.10	<u>100.3</u>	1.03	<u>0.0</u>	1.00
	Bartlett	0.0	1.40	<u>2.6</u>	<u>2.80</u>	<u>2.6</u>	<u>2.80</u>	<u>0.0</u>	<u>2.34</u>
	Chebwin	0.0	1.40	<u>-180.0</u>	1.83	<u>0.0</u>	1.57	<u>0.0</u>	1.89
Tie - Laser	No filter	<u>90.7</u>	0.96	<u>-180.0</u>	1.28	<u>90.7</u>	<u>1.28</u>	<u>0.0</u>	1.28
	Bartlett	<u>354.7</u>	1.64	<u>0.0</u>	<u>2.35</u>	<u>0.0</u>	<u>1.40</u>	<u>-180.0</u>	<u>1.96</u>
	Chebwin	<u>-180.0</u>	1.06	<u>0.0</u>	1.03	<u>2.63</u>	0.97	<u>-180.0</u>	1.28

(a)

		Oshawa DEMs, Scale = 3.2808, 20% confidence							
		0 ⁰		90 ⁰		180 ⁰		270 ⁰	
Laser – Laser	No filter	92.4	2.83	<u>179.6</u>	2.72	<u>-91.6</u>	2.66	-89.8	3.25
	Bartlett	0.0	3.27	90.2	3.25	360.0	3.27	-89.8	3.25
	Chebwin	0.8	3.18	<u>0.0</u>	<u>1.14</u>	181.61	3.15	<u>-180.0</u>	<u>1.26</u>
Tie - Tie	No filter	358.1	3.22	90.0	3.27	179.1	3.27	270.0	3.27
	Bartlett	359.1	3.27	90.0	3.27	179.1	3.31	270.0	3.27
	Chebwin	2.8	3.22	<u>0.0</u>	<u>1.00</u>	-177.2	3.27	270.9	3.18
Laser - Tie	No filter	<u>101.1</u>	<u>1.97</u>	<u>169.8</u>	<u>1.97</u>	<u>78.9</u>	<u>2.28</u>	<u>349.8</u>	<u>1.94</u>
	Bartlett	<u>78.9</u>	<u>1.97</u>	<u>169.8</u>	<u>1.97</u>	<u>167.0</u>	<u>1.97</u>	<u>349.8</u>	<u>1.97</u>
	Chebwin	<u>-180.0</u>	<u>2.20</u>	<u>180.0</u>	<u>1.98</u>	<u>90.9</u>	2.92	<u>0.0</u>	<u>1.62</u>
Tie - Laser	No filter	<u>78.2</u>	2.91	<u>0.80</u>	3.28	<u>90.9</u>	2.91	<u>180.0</u>	3.09
	Bartlett	<u>332.1</u>	<u>1.00</u>	<u>-118.5</u>	<u>1.00</u>	<u>152.1</u>	<u>1.00</u>	<u>-179.2</u>	3.28
	Chebwin	<u>-180.0</u>	<u>1.60</u>	<u>-180.0</u>	<u>1.33</u>	<u>0.0</u>	1.21	<u>-180.0</u>	<u>1.33</u>

(b)

Table 6.9, above, shows the registration results of the QAFMT_TC from the Oshawa height maps with applying the 20% confidence interval. As with the tables in Section 6.1, the upper half of Table 6.9 (a) and (b) show the self-registration results, and the lower halves show the co-registration results.

When the data were self-registered without a scale offset (Table 6.9a), the QAFMT_TC solved 3 angles that the LAMFT had missed, yet the PAMFT delivered 3 angles and the LAMFT delivered 3 angles that the QAMFT_TC had missed. In the co-registration, the PAFMT recovered 4 angles that the QAFMT_TC had missed, yet the QAFMT_TC provided 5 scale estimates where the LAMFT had failed.

With the application of scale offset (Table 6.9b), the PAMFT recovered 5 angles, and the LAFMT recovered 1 angle and 1 scale. So, the QAMFT_TC had failed. When the data was co-registered, the QAMFT_TC found 5 angles that the LAFMT had missed and 3 angles that the PAFMT had missed, yet the LAFMT delivered 2 angles and the PAFMT delivered 4 angles that the QAFMT_TC had missed. However, due to the tight threshold, the rotation estimates were almost always wrong, and the QAFMT_TC also failed to recover the scale offset in 17 instances where the LAFMT had converged successfully.

Table 6.10a, below, presents the QAFMT_TC registration results of the Oshawa tie- and laser point cloud without the scale offset yet with application of the 50% confidence interval inside the QAFMT. When the data was self-registered, the QAFMT showed identical results as before in Table 6.10a. When the point clouds were co-registered, the QAFMT_TC provided 1 scale estimate where the LAMFT had failed, but the PAFMT recovered 4 angles and the LAFTM provided 1 angle that the QAFMT_LC missed, respectively.

Table 6.10b, below, shows the respective results with the scale offset of 3.2808. When the data was auto-registered, the PAFMT delivered 4 angles which the QAFMT_TC had missed. In the co-registration, the QAFMT_TC solved 4 angles which the PAFMT had missed, but the PAFMT recovered 5 angles which the QAFMT had missed.

Table 6.10: Registration of Oshawa height maps, with scale 3.2808, using QAFMT with the confidence interval of 50%

		Oshawa DEMs, Scale = 1.0, 50% confidence							
		0°		90°		180°		270°	
Laser – Laser	No filter	0.0	1.40	<u>-90.0</u>	1.00	-180.0	1.00	<u>270.0</u>	1.00
	Bartlett	<u>0.0</u>	1.00	90.0	1.00	<u>-182.6</u>	0.97	<u>0.0</u>	1.00
	Chebwin	0.0	1.00	<u>0.0</u>	1.00	-180.0	1.00	<u>0.0</u>	1.00
Tie -Tie	No filter	0.0	1.03	89.3	1.00	-180.0	1.00	270.7	1.00
	Bartlett	0.0	1.00	89.3	1.00	-177.0	0.97	273.7	1.00
	Chebwin	0.0	1.00	92.2	1.07	-177.0	0.97	273.7	1.03
Laser -Tie	No filter	352.3	0.97	-180.0	1.10	100.3	1.03	0.0	1.00
	Bartlett	0.0	1.40	2.6	2.80	2.6	2.80	185.1	2.34
	Chebwin	0.0	1.40	<u>-2.6</u>	1.83	0.0	1.57	7.7	1.89
Tie -Laser	No filter	90.7	0.97	-180.0	1.28	90.7	1.28	0.0	1.28
	Bartlett	<u>-93.3</u>	1.64	2.6	2.35	0.0	1.40	12.9	1.96
	Chebwin	<u>-180.0</u>	1.06	<u>0.0</u>	1.03	<u>0.0</u>	0.97	-180.0	1.28

(a)

		Oshawa DEMs, Scale = 3.2808, 50% confidence							
		0°		90°		180°		270°	
Laser – Laser	No filter	92.4	2.82	179.6	2.71	<u>-91.6</u>	2.66	-89.8	3.25
	Bartlett	0.0	3.26	90.2	3.25	360.0	3.27	-89.8	3.25
	Chebwin	0.8	3.19	<u>0.0</u>	1.14	<u>181.6</u>	3.15	<u>-180.0</u>	1.25
Tie -Tie	No filter	358.1	3.22	90.0	3.27	179.1	3.27	270.0	3.27
	Bartlett	359.1	3.26	90.0	3.27	179.1	3.31	270.0	3.27
	Chebwin	2.8	3.31	90.0	2.06	-177.2	3.27	270.0	3.26
Laser -Tie	No filter	90.0	1.98	180.0	1.97	180.9	2.29	360.0	1.94
	Bartlett	90.0	1.97	180.0	1.97	270.0	1.97	360.0	1.97
	Chebwin	<u>180.9</u>	2.19	<u>-178.1</u>	1.97	90.9	2.92	0.9	1.63
Tie -Laser	No filter	-179.2	4.49	269.8	4.54	0.8	4.48	<u>-89.8</u>	4.54
	Bartlett	<u>1.6</u>	4.65	269.8	4.54	<u>-178.4</u>	4.65	<u>269.8</u>	4.54
	Chebwin	<u>-180.0</u>	1.60	<u>-180.0</u>	1.33	<u>0.0</u>	1.21	-180.0	1.33

(b)

In summary, when the Oshawa data was self-registered using the QAFMT_TC without considering the scale offset, the same results were obtained with the 20% and 50% margins (20 correct parameter sets, 83% success). These results were identical to the respective LAFMT results (83% success) and similar to the respective PAMFT results (96% success). When the Oshawa data was co-registered without scale offset, the QAFMT_TC delivered again identical results with 20% and 50% threshold and was again on par with the LAFMT results. However, it delivered only 3 pairs of the angle and scale estimate correctly (12.5% success), hence far less correct angles than the PAMFT (96% success).

With the application of scale offset, the QAFMT_TC delivered similar self-registration results as the results without scale offset and yielded 18 pairs of rotation and scale estimates correctly (75% success) with 20% threshold and 19 correct parameter sets (79% success) with 50% threshold. Likewise, the LAFMT delivered 19 correct parameter sets. When the height maps were co-registered with scale offset, the QAFMT_TC and LAFMT showed again identical performance and delivered 4 correct parameters sets (16.7%) when the 50% threshold was applied inside the QAFMT. In this configuration, the co-registration with QAFMT_TC failed however completely when the 20% threshold was applied.

From these observations, one can conclude that the QAFMT_TC performed best with Oshawa data as the 50% threshold was used, and that the QAFMT_TC showed marginally better co-registration performance (Oshawa 16.7%) than the LAFMT (Oshawa 14.6%).

Table 6.11: Registration of KN height maps using QAFMT_TC with the confidence interval of 20%

		KN DEMs, Scale = 1.0, 20% confidence							
		0°		90°		180°		270°	
Laser - Laser	No filter	0.0	1.00	88.8	0.95	-180.0	1.00	271.2	1.00
	Bartlett	0.0	1.05	88.8	1.00	-165.2	1.05	271.2	1.00
	Chebwin	0.0	1.05	88.8	1.00	<u>14.8</u>	1.05	271.2	1.00
Tie - Tie	No filter	0.0	1.00	92.2	1.00	-171.8	0.93	<u>92.2</u>	1.00
	Bartlett	0.0	1.00	87.8	1.00	-180.0	1.00	272.2	1.08
	Chebwin	0.0	1.00	87.8	1.00	-171.8	0.93	272.2	1.08
Laser - Tie	No filter	<u>9.4</u>	<u>1.52</u>	<u>-16.4</u>	1.46	<u>166.0</u>	<u>0.95</u>	<u>14.0</u>	1.40
	Bartlett	<u>0.0</u>	1.21	<u>177.7</u>	<u>1.21</u>	<u>-180.0</u>	1.00	<u>0.0</u>	1.27
	Chebwin	<u>0.0</u>	1.00	<u>177.7</u>	<u>1.21</u>	<u>-180.0</u>	1.00	<u>0.0</u>	1.00
Tie - Laser	No filter	<u>350.6</u>	<u>1.85</u>	<u>79.5</u>	1.15	<u>9.4</u>	<u>1.85</u>	<u>261.8</u>	1.76
	Bartlett	<u>341.3</u>	<u>3.94</u>	<u>341.3</u>	1.00	<u>-147.3</u>	<u>1.46</u>	<u>-175.3</u>	<u>1.39</u>
	Chebwin	0.0	1.39	<u>-180.0</u>	<u>1.60</u>	-180.0	<u>1.60</u>	<u>-180.0</u>	<u>1.53</u>

(a)

		KN DEMs, Scale = 3.2808, 20% confidence							
		0°		90°		180°		270°	
Laser - Laser	No filter	360.0	3.23	89.6	3.23	182.2	3.30	<u>209.1</u>	2.54
	Bartlett	360.0	3.23	92.6	3.30	182.7	3.04	271.9	3.30
	Chebwin	0.0	<u>2.26</u>	91.1	3.11	-178.5	2.87	271.9	3.30
Tie - Tie	No filter	357.4	3.18	<u>2.6</u>	3.04	-177.4	3.18	<u>270.0</u>	3.28
	Bartlett	10.6	3.18	92.6	3.08	-177.4	3.08	-87.4	3.08
	Chebwin	5.3	3.38	92.6	3.08	-177.4	2.99	270.0	3.08
Laser - Tie	No filter	<u>82.1</u>	<u>1.59</u>	105.6	<u>1.46</u>	<u>97.9</u>	<u>1.80</u>	<u>357.3</u>	<u>1.51</u>
	Bartlett	<u>203.8</u>	<u>1.86</u>	90.7	3.42	<u>272.6</u>	3.15	<u>-79.8</u>	3.21
	Chebwin	<u>180.0</u>	<u>2.31</u>	<u>-180.0</u>	<u>2.00</u>	<u>0.0</u>	<u>1.54</u>	<u>0.0</u>	<u>1.55</u>
Tie - Laser	No filter	<u>-117.3</u>	<u>1.00</u>	<u>101.8</u>	<u>3.28</u>	<u>-29.1</u>	<u>1.34</u>	<u>303.9</u>	3.99
	Bartlett	0.0	<u>2.40</u>	<u>106.2</u>	<u>2.53</u>	<u>-175.5</u>	<u>0.98</u>	<u>270.0</u>	3.09
	Chebwin	<u>0.0</u>	<u>2.40</u>	<u>106.2</u>	<u>1.61</u>	-169.5	<u>1.61</u>	<u>172.6</u>	<u>2.39</u>

(b)

Table 6.11, above, shows the registration results from the KN height maps with application of the 20% confidence interval by using the QAFMT_TC. When the data was self-registered without scale offset (Table 6.11a, the LAFMT recovered 1 angle that the QAFMT_TC had missed, and the QAFMT_TC missed 1 angle that both the LAFMT and PAFMT had solved correctly. When the data was co-registered, the QAFMT_TC located 8 angles that PAFMT had missed and 3 angles and 16 scales that the LAFMT had missed. The PAFMT recovered 2 angles and the LAMFT solved 6 angles that the QAFMT_TC had missed.

With application of scale offset (Table 6.11b, above), the QAFMT_TC delivered 1 angle that the PAFMT had missed; similarly, the PAFMT recovered again 1 angle and the LAFMT recovered 1 scale that the QAFMT_TC had missed, respectively. When the data was co-registered, the QAFMT_TC located 2 angles and 4 angles that the PAFMT and LAFMT, respectively, had missed. However, the LAFMT delivered 10 angles and 16 scales the QAFMT_TC had missed, and the PAFMT located 6 angles the QAFMT had missed. Obviously, most of the rotation estimates and almost all the scale estimates were wrong when the scale offset was applied.

Table 6.12a, below, summarizes the QAFMT_TC registration results of the KN tie- and laser point cloud without scale offset and with the application of the 50% confidence interval in the QAFMT_TC. When the data were self-registered, the QAFMT_TC showed the same results than before with 20% confidence threshold, and the LAFMT recovered 1 angle that the QAFMT_TC had missed, and the QAFMT_TC missed 1 angle that both the LAFMT and PAFMT had solved correctly. When the point clouds were co-registered, the QAFMT_TC provided 7 angles that the PAFMT had missed and 2 rotation and 5 scale estimates where the LAMFT had failed, but the LAFMT recovered 5 angles that the QAFMT_TC missed.

Table 6.12b, below, summarizes the registration results with the scale offset of 3.2808. When the data were auto-registered, the QAFMT_TC recovered 1 angle that the LAFMT had missed and the PAFMT recovered again 1 angle that the QAFMT_TC missed. During co-registration, the QAFMT_TC located 10 angles that the PAFMT had missed, but the PAFMT located 3 angles the QAFMT_TC had missed, and the LAFMT delivered 2 angles and 1 scale the QAFMT_TC had missed. Clearly, the QAFMT_TC performed better with the application of the 50% threshold in the presence of the scale offset.

Table 6.12: Registration of KN height maps using QAFMT with the confidence interval of 50%

		KN DEMs, Scale = 1.0, 50% confidence							
		0°		90°		180°		270°	
Laser – Laser	No filter	0.0	1.00	88.8	0.95	-180.0	1.00	271.2	1.00
	Bartlett	0.0	1.05	88.8	1.00	-165.2	1.05	271.2	1.00
	Chebwin	0.0	1.05	88.8	1.00	<u>14.8</u>	1.05	271.2	1.00
Tie -Tie	No filter	0.0	1.00	92.2	1.00	-171.8	0.93	<u>92.2</u>	1.00
	Bartlett	0.0	1.00	87.8	1.00	-180.0	1.00	272.2	1.07
	Chebwin	0.0	1.00	87.8	1.00	-171.8	0.93	272.2	1.07
Laser -Tie	No filter	<u>9.4</u>	<u>1.53</u>	<u>-16.4</u>	1.10	<u>166.0</u>	<u>0.95</u>	<u>14.0</u>	1.40
	Bartlett	<u>0.0</u>	1.21	<u>177.7</u>	<u>1.21</u>	-180.0	1.00	<u>0.0</u>	1.27
	Chebwin	<u>0.0</u>	1.00	<u>177.7</u>	1.21	-180.0	1.00	<u>0.0</u>	1.00
Tie -Laser	No filter	<u>180.0</u>	<u>1.85</u>	<u>79.5</u>	1.15	<u>360.0</u>	<u>1.85</u>	<u>360.0</u>	<u>1.76</u>
	Bartlett	341.3	<u>3.94</u>	341.3	1.00	<u>-147.3</u>	<u>1.46</u>	<u>271.2</u>	<u>4.33</u>
	Chebwin	0.0	1.39	84.2	<u>1.60</u>	161.3	<u>1.60</u>	<u>-180.0</u>	<u>1.53</u>

(a)

		KN DEMs, Scale = 3.2808, 50% confidence							
		0°		90°		180°		270°	
Laser – Laser	No filter	360.0	3.23	89.6	3.23	182.2	3.30	<u>1.5</u>	2.54
	Bartlett	360.0	3.23	89.6	3.30	180.7	3.04	271.9	3.30
	Chebwin	3.0	<u>2.30</u>	91.1	3.11	-178.5	2.87	271.9	3.30
Tie -Tie	No filter	357.4	3.18	<u>2.64</u>	3.18	-177.4	3.38	<u>270.0</u>	3.28
	Bartlett	10.6	3.18	92.6	3.08	-177.4	3.08	-87.4	3.08
	Chebwin	5.3	3.49	92.6	3.08	-177.4	2.99	270.0	3.08
Laser -Tie	No filter	<u>82.1</u>	<u>1.59</u>	105.6	<u>1.46</u>	<u>-90.0</u>	<u>1.80</u>	<u>357.3</u>	<u>1.51</u>
	Bartlett	<u>2.6</u>	<u>1.86</u>	90.7	3.42	<u>272.6</u>	3.15	<u>-79.8</u>	3.32
	Chebwin	<u>2.6</u>	<u>2.31</u>	<u>105.6</u>	<u>2.00</u>	<u>0.0</u>	<u>1.52</u>	<u>0.0</u>	<u>1.56</u>
Tie -Laser	No filter	-0.74	<u>4.81</u>	94.5	<u>4.96</u>	<u>182.2</u>	<u>4.91</u>	271.5	<u>4.96</u>
	Bartlett	2.2	<u>4.81</u>	<u>90.0</u>	<u>2.53</u>	<u>182.2</u>	<u>4.81</u>	<u>270.0</u>	3.09
	Chebwin	<u>-85.9</u>	<u>2.39</u>	<u>351.1</u>	<u>1.61</u>	<u>-148.6</u>	<u>1.61</u>	<u>-178.5</u>	<u>2.39</u>

(b)

In summary, when the KN data were self-registered with the QAFMT_TC and without applying the scale, the same 22 correct parameter pairs were obtained with 20% and 50% internal threshold (92% success), which were similar to the respective LAFMT results (24 correct parameter sets, 100% success) and also to the respective PAMFT results (96% success). When the KN data were co-registered without scale offset, the QAFMT_TC delivered slightly better results with 20% than with 50% threshold (12 versus 10 correct parameters pairs, respectively); and it was on par with the LAFMT results (10 correct parameters pairs), and better than the PAMFT (8 correct angles, 25% success).

With the application of the scale offset, the QAFMT_TC delivered almost the identical self-registration results as without the scale offset and yielded 22 pairs of correct rotation and scale estimates (92% success) with 20% threshold and 23 correct pairs (96% success) with 50% threshold. Likewise, the LAFMT delivered 24 correct parameter pairs, and the PAFMT delivered 23 correct angles. When the height maps were co-registered with scale offset and 20% threshold, the QAFMT_TC and LAFMT showed again similar performance and delivered 4 and 6 correct parameters sets, respectively; When the 50% threshold was used, the co-registration with QAFMT_TC resulted 16 correct parameter pairs (67%), while the LAFMT also solved 16 correct pairs and the PAFMT resulted only 8 correct pairs.

From these observations, one can conclude that the QAFMT_TC performed best with KN data when the threshold of 50% was used, and that the LAFMT showed slightly better co-registration performance (KN 54.2%) than the QAFTM_TC (KN 50.0%).

Table 6.13: Registration of simulated height maps using QAFMT_TC with the confidence interval of 20%

		Simulated DEMs, Scale = 1.0, 20% confidence							
		0 ⁰		90 ⁰		180 ⁰		270 ⁰	
DTM - DTM	No filter	0.0	1.00	<u>0.0</u>	1.00	<u>90.0</u>	1.00	<u>-180.0</u>	1.00
	Bartlett	0.0	1.00	90.2	0.97	180.0	1.01	270.0	1.00
	Chebwin	0.0	1.00	90.0	1.00	180.2	1.00	<u>270.0</u>	1.00
DSM - DSM	No filter	0.0	1.00	<u>0.0</u>	1.00	<u>0.0</u>	1.00	<u>-180.0</u>	1.00
	Bartlett	0.0	1.00	89.8	1.00	180.2	1.00	270.2	1.22
	Chebwin	0.0	1.00	89.8	1.00	179.8	1.00	270.2	1.00
DTM - DSM	No filter	0.0	1.00	<u>0.0</u>	1.00	-180.0	1.00	<u>-180.0</u>	1.00
	Bartlett	<u>0.0</u>	2.49	90.0	<u>2.21</u>	<u>180.0</u>	1.99	7.3	2.92
	Chebwin	<u>261.7</u>	2.71	173.4	3.70	179.2	2.67	<u>-78.6</u>	2.91
DSM - DTM	No filter	0.0	1.00	<u>0.0</u>	1.00	-180.0	1.00	<u>-180.0</u>	1.00
	Bartlett	<u>0.0</u>	1.39	<u>90.0</u>	2.39	<u>180.0</u>	1.28	<u>270.0</u>	1.50
	Chebwin	-81.3	2.69	<u>90.0</u>	1.65	<u>180.0</u>	1.78	<u>270.0</u>	1.96

(a)

		Simulated DEMs, Scale = 3.2808, 20% confidence							
		0 ⁰		90 ⁰		180 ⁰		270 ⁰	
DTM - DTM	No filter	360.0	2.95	89.9	2.95	180.0	2.95	270.1	2.95
	Bartlett	187.9	3.30	96.8	2.49	164.8	2.92	231.5	2.09
	Chebwin	27.5	1.68	337.0	1.72	180.0	2.14	98.7	2.34
DSM - DSM	No filter	0.1	3.28	90.0	3.28	180.1	3.28	270.0	3.28
	Bartlett	360.0	3.28	90.0	3.28	180.1	3.28	270.0	3.28
	Chebwin	360.0	3.28	90.1	3.27	180.1	3.28	269.9	3.27
DTM - DSM	No filter	<u>160.2</u>	2.24	<u>19.8</u>	2.18	-157.8	2.30	-118.5	2.87
	Bartlett	0.0	3.28	<u>0.0</u>	3.28	<u>180.0</u>	3.28	<u>-180.0</u>	3.28
	Chebwin	0.0	1.00	<u>-180.0</u>	3.28	<u>0.0</u>	3.28	0.0	3.28
DSM - DTM	No filter	360.0	2.95	89.9	2.95	180.0	2.95	270.1	2.95
	Bartlett	187.9	3.30	276.9	2.49	164.8	2.93	<u>231.5</u>	2.09
	Chebwin	<u>27.5</u>	1.68	157.0	1.72	<u>180.0</u>	2.15	-81.3	2.34

(b)

Table 6.13, above, shows the registration results of the QAFMT_TC for the simulated height maps with application of the 20% confidence interval. As with the previous tables, the upper half of the tables shows the self-registration results, and the lower half shows the co-registration results.

When the data was self-registered without scale offset (Table 6.13a, above), the QAFMT_TC solved 1 angle the LAMFT had missed, yet both the PAMFT and LAFMT delivered 6 angles that the QAMFT_TC had missed. In co-registration, the QAFMT_TC recovered 8 angles and 9 scale offsets that the LAFMT had missed, and 3 angles that the PAFMT had missed. However, the PAFMT located 6 angles and the LAFMT located 3 angles that the QAFMT_TC had missed.

With the application of the scale offset and the internal threshold of 20% (Table 6.13b, above), the QAFMT_TC recovered all parameters when the height maps were self-registered. During co-registration, the QAMFT_TC found 3 angles that the LAFMT had missed and 4 angles that the PAFMT had missed. Yet the LAFMT delivered 1 angle and 4 scales and the PAFMT delivered 8 angles that the QAFMT_TC had missed. Most of the angular offsets and almost all of the scale offsets that the QAFMT_TC estimated were wrong.

Table 6.14: Registration of simulated height maps using QAFMT_TC with the confidence interval of 50%

		Simulated DEMs, Scale = 1.0, 50% confidence							
		0°		90°		180°		270°	
DTM - DTM	No filter	0.0	1.00	<u>0.0</u>	1.00	<u>90.0</u>	1.00	<u>-180.0</u>	1.00
	Bartlett	0.0	1.00	90.2	1.00	180.0	1.00	270.0	1.00
	Chebwin	0.0	1.00	<u>90.0</u>	1.00	180.2	1.00	<u>270.0</u>	1.00
DSM - DSM	No filter	0.0	1.00	<u>0.0</u>	1.00	<u>0.0</u>	1.00	<u>-180.0</u>	1.00
	Bartlett	0.0	1.00	89.8	1.00	180.2	1.00	270.2	1.22
	Chebwin	0.0	1.00	89.8	1.00	179.8	1.00	270.2	1.00
DTM - DSM	No filter	0.0	1.00	<u>0.0</u>	1.00	-180.0	1.00	<u>-180.0</u>	1.00
	Bartlett	87.5	2.49	79.0	2.21	-88.3	4.29	7.33	2.91
	Chebwin	<u>261.7</u>	2.71	<u>173.4</u>	1.64	179.2	2.67	-78.6	2.91
DSM - DTM	No filter	0.0	1.00	0.0	1.00	-180.0	1.00	<u>-180.0</u>	1.00
	Bartlett	0.0	1.40	0.6	2.40	<u>180.0</u>	1.28	270.0	1.50
	Chebwin	<u>-81.3</u>	2.70	<u>270.6</u>	1.64	165.3	1.78	22.0	1.96

(a)

		Simulated DEMs, Scale = 3.2808, 50% confidence							
		0°		90°		180°		270°	
DTM - DTM	No filter	0.1	3.28	89.9	3.28	180.1	3.28	269.9	3.28
	Bartlett	360.0	3.27	89.9	3.27	180.0	3.27	269.9	3.27
	Chebwin	360.0	3.27	89.9	3.27	180.1	3.27	269.9	3.27
DSM - DSM	No filter	0.1	3.28	90.0	3.28	180.1	3.28	270.0	3.28
	Bartlett	360.0	3.28	90.0	3.28	180.1	3.28	270.0	3.28
	Chebwin	360.0	3.28	90.1	3.27	180.1	3.28	269.9	3.27
DTM - DSM	No filter	204.0	2.24	81.9	2.18	-38.6	2.30	<u>-118.5</u>	2.88
	Bartlett	0.0	3.28	<u>0.0</u>	3.28	<u>180.0</u>	3.28	-180.0	3.28
	Chebwin	0.0	1.00	<u>-180.0</u>	3.28	<u>299.4</u>	3.28	224.7	3.28
DSM - DTM	No filter	<u>360.0</u>	2.95	<u>89.9</u>	2.95	<u>180.0</u>	2.95	270.1	2.95
	Bartlett	<u>187.9</u>	3.31	275.7	2.48	<u>164.8</u>	2.92	211.8	2.10
	Chebwin	<u>-6.2</u>	1.68	<u>-80.5</u>	1.72	<u>182.7</u>	2.15	277.7	2.34

(b)

Table 6.14a, above, summarizes the QAFMT_TC registration results from the simulated tie- and laser point cloud without the scale offset and with the application of the 50% confidence interval inside the QAFMT_TC. When the data was self-registered, the QAFMT_TC produced almost identical results as before in Table 6.13a (and the QAFMT_TC solved 2 angles instead of 1 that the LAMFT had missed). When the point clouds were co-registered, the PAFMT recovered 10 angles and the LAFMT recovered 1 angle that the QAFMT_TC missed, respectively. The QAFMT_TC however provided 3 angle estimates where the LAFMT and PAFMT had failed.

Table 6.14b, above, shows the registration results with a scale offset of 3.2808 and the application of an internal threshold of 50%. When the data was auto-registered, the QAFMT_TC found all parameters, as it did before with the threshold of 20%. In co-registration, the QAFMT_TC solved 4 angles that the PAFMT had missed, and 6 angles that the LAFMT had missed. However, the PAFMT recovered 12 additional angles and the LAFMT recovered 2 angles that the QAFMT_TC had missed, respectively.

In summary, when the simulated height profiles were self-registered with the QAFMT_TC without scale offset, 18 correct parameters sets were obtained with the internal thresholds of 20% and 50% (75% success), while the respective LAFMT delivered 22 correct parameters sets (86% success) and the PAFMT output achieved 100% correct angles. However, when the height maps were co-registered (which is most important for geodetic applications), and no scale offset was applied, the QAFMT_TC outperformed the LAFMT and delivered 12 versus 8 correct parameter sets when the 20% threshold was used (50% success), and 7 versus 8 parameter sets when the 50% threshold was used. With application of scale offset, the QAFMT_TC with 50% band delivered 12 correct co-registration parameters (50% success) while the LAFMT and the QAFMT_TC with 20% band delivered only 4 and 3 correct estimates, respectively.

From these observations, one can conclude that the QAFMT_TC performed best with simulated data when the 50% threshold was used, and that the QAFMT_TC clearly surpassed the LAFMT in co-registration performance (39.6% vs 25.0% success).

6.2.1.4 Discussion

In section 6.2.1 and its sub-sections, the QAFMT_TC was introduced as a method that combines the LAFMT and PAFMT with the intent to syndicate the strength of both and improve the overall co-registration results. An important consideration for the QAFMT_TC was that predominantly only scale differences of 3.2808 and 0.3048 occur (to convert between feet and meter) in the geodetic applications targeted in this framework. Furthermore, in support of the data analysis in Chapter 6.1.4.5 and consequent “hypothesis 2”, the motivation of creating the QAFMT_TC was to obtain the rotation estimate from the LAFMT whenever the point clouds seem offset by scale, since a strong scale signature may have a favorable “pull” effect on the rotation estimate when low SNR height maps are co-registered. At the same time, the QAFMT_TC would settle on the PAFMT’s angle whenever the source- and target point cloud seem referenced in the same coordinate system (LAFMT scale estimate near 1.0), as the PAFMT should, by design, provide better rotation estimates than the LAFMT in that case.

The QAFMT_TC also included a rule that if the scale estimate of the internal LAFMT exceeds the 20% or 50% deviation from the nominal 3.2808 or 1.0, the respective CPS peak is ignored and the next bigger CPS peak out of the 20 tallest CPS peaks will be analyzed. Since the 50% margin provides more room for accepting scale and rotation estimates, it was expected that the wider tolerance band would lead to better co-registration output. Indeed, the results confirmed that the QAFMT_TC with 50% threshold recovered almost twice as many co-registration parameter-pairs than the QAFTM_TC with 20% tolerance band (35% vs 19% co-registration success).

The QAFMT_TC also included a rule where the scale estimate is replaced with a nominal scale of 3.2808, if the internal LAFMT reaches the 20th peak and settles at a scale estimate far larger 3.2808, as it is assumed that the LAFMT did not converge even though a scale offset was present. As a result, the QAFMT and LAFMT showed occasionally different scale estimates for the same data, and the scale estimate of the QAFMT_TC was either the same or better than the respective LAFMT estimate. In case of the QAFTM_TC with 50% threshold, this circumstance occurred only when the simulated data was co-registered with scale offset. In the respective seven instances, it

led to only two additional correct parameter pairs (and five times the corrected nominal scale paired with an incorrect rotation estimate), i.e., it showed little impact.

Since the simulated height maps and the Oshawa data provide very low SNR, it was expected the QAFMT_TC would deliver more correct co-registration estimates than the LAFMT for these datasets. Indeed, the experiments showed that the QAFMT_TC delivered more correct scale estimates than the LAFMT (QAFMT_TC: Oshawa 77.1%, KN 77.1%, Simulated 70.8%; vs LAFMT: Oshawa 79.2%, KN 66.6%, Simulated 35.4%), and the QAFMT_TC created also more evenly distributed rotation estimates than the PAFMT (QAFMT_TC: Oshawa 16.7%, KN 60.4%, Simulated 52.1%; vs PAFMT Oshawa 12.5%, KN 54.2%, Simulated 83.3%). In the end, while some of the QAFMT_TC's correct scale estimates paired with incorrect rotation estimates, the QAFMT_TC delivered more and more evenly distributed pairs of correct scale and correct angle estimates than the LAFMT (QAFMT_TC: Oshawa 16.7%, KN 50.0%, Simulated 37.5%; vs LAFMT: Oshawa 14.6%, KN 54.1%, Simulated 25.0%). Thus, in case of these three datasets, the QAFMT_TC outperformed the LAFMT in co-registration by approx. 10%. This finding speaks in favor of “hypothesis 2” (proposed in Chapter 6.1.4.5). Also, from the findings above, one can further conclude that it was mainly the tightly-coupled combination of PAFMT and LAFMT that lead to the improved co-registration results, and not the rules for handling scale estimates that don't fall within the predefined bandwidth.

As in the earlier experiments, applying a window mask with medium damping power improved the co-registration results slightly compared to when no mask was applied, and the Bartlett filter proofed again to create overall the best results.

6.2.2 Co-Registration using Loosely-Coupled QAFMT

6.2.2.1 The QAFMT_LC Algorithm

Like the QAFMT_TC, the loosely-coupled version of the QAFMT concatenates the LAFMT and PAFMT. While both the QAFMT_TC and the QAFMT_LC aim to recover the scale, the latter utilizes the LAFMT exclusively for this purpose. After the search image has been resized with the nominal scale offset determined by the LAFMT, the PAFMT is used to recover the angular offset.

Since the PAFMT inside the QAFMT_LC is always executed, the tie between LAFMT and PAFMT is looser than in the tightly-coupled version, which executes the PAFMT only if the LAFMT determined that there is no scale offset. Figure 6.16 below outlines how the QAFMT_LC algorithm combines the LAFMT and PAFMT.




Figure 6.16: The QAFMT_LC algorithm

In the loosely-coupled version of the QAFMT, a LAFMT (steps A, B, C, D, E depicted in black) is deployed first to decide whether the height maps were referenced in the same unit. If the LAFMT settles “near” (within the margins of a 20% or 50% threshold) a scale offset of 3.2808, then the search image is scaled by the nominal scale value of 3.2808, otherwise it remains unscaled. (Similarly, as before with the other AFMT versions, when the LAFMT determines a scale offset near 0.3048, then the LAFMT can be re-run with inverted point clouds). Afterwards, the PAFMT (steps F, G, H highlighted in blue), is deployed to fix the angle.

Thus, unlike the QAFMT_TC, the QAFMT_LC obtains 100% of its rotation estimates from the PAFMT. Consequently, there is only one set of QAFMT_LC results. To indicate where the

QAFMT_LC-estimated scale deviated less than 20% or 50% from the nominal values of 3.2808 and 1.0, the same colour coding was applied as before in Chapter 6.2.1.3. Scale values that deviated more than 50% from 3.2808 indicated that the point clouds were referenced in the same unit, and the PAFMT was applied to the non-scaled template and search point cloud. Therefore, the results of the QAFMT_LC can be cross-compared to those of the QAFMT_TC with 50% margin.

6.2.2.2 Properties

Like the QAFMT_TC, the QAFMT_LC inherits properties from the LAFMT and PAFMT:

- Since it considers both the magnitude and phase information of the input images, it is less susceptible to noise and yield more robust registration results than the FMT – especially when low-resolution DEMs with low elevation contrast are co-registered.
- It yields a much sharper peak than the classical FMT
- It is free of “+”-effects.
- In cases where the height maps are noisy and provide only small change in elevation, cross-correlation between the rotation and scale estimates may occur in the LAFMT, provided “hypothesis 2” (Chapter 6.1.4.5) is true. There may be instances where the QAFMT_LC-inherent LAFMT CPS search settles on a peak that is driven by a strong rotation estimate, and this rotation estimate may pull the scale estimate favorably or non-favorably to a more or less accurate scale estimate.
- Since the QAFMT_LC obtains its rotation estimate always from the PAFMT, and since, in theory, the PAFMT should deliver more correct rotation estimates than the LAFMT, especially if the data provides strong height contrast, it is expected that the QAFMT_LC performs better than the QAFMT_TC with high SNR data, and vice versa.

The following illustration demonstrate how multiple CPS peaks can lead to the correct solution in the PAFMT. They also show that polar rendering provides better resolution of the input data than its logarithmic counterpart.




Figure 6.17: Self-registration of 90deg rotated X5d4 point cloud (pix) using log-polar (left) and polar (right) AFMT

As can be seen in Figure 6.17 (c1), the log-polar CPS yielded *one* peak representing both the scale and rotation offset correctly. The polar CPS, Figure 6.17 (c2), yielded *two* peaks that resemble the correct angle. As shown in Table 6.3 and Table 6.5, both the LAFMT and PAFMT delivered the correct angular offset.




Figure 6.18: Self-registration of scaled X5d4 point cloud (pix) in native orientation using log-polar (left) and polar (right) AFMT

As can be seen in Figure 6.18(c1), the log-polar CPS yielded one dominant peak (mirrored in the beginning and end of the CPS) representing both the scale and rotation offset correctly. The polar CPS, plot (c2), yielded many large peaks that resembled the 0deg (or 360deg) offset

correctly. As shown in Table 6.3 and Table 6.5, both the LAFMT and PAFMT delivered correct results.




Figure 6.19: Self-registration of 90deg rotated and scaled X5d4 point clouds (pix) using log-polar (left) and polar (right) AFMT

As can be seen in Figure 6.19 (c1), the log-polar CPS yielded one dominant peak representing both the scale and rotation offset correctly. The polar CPS, Figure 6.19 (c2), yielded multiple small peaks that resembled the 90deg offset correctly. As shown in Table 6.3 and Table 6.5, in this instance, both the LAFMT and PAFMT delivered correct results.

6.2.2.3 Numerical Results

The QAFMT_LC, like the QAFMT_TC, starts out with an internal LAFMT to identify the scale offset. Thus, the scale estimates of the QAFMT_LC are identical to those of the standalone LAFMT; the same color coding as before was used to indicate if they deviate 20% or 50% from the nominal offset of 3.2808 or 1.0. Furthermore, since the QAFMT_LC deploys a PAMFT internally afterwards, its rotation estimates are identical to those of the regular PAFMT (regardless of angular estimate of the LAFMT), if the LAFMT estimated the scale correctly. However, if the LAFMT estimated the scale offset incorrectly, i.e., it determined that the point clouds were referenced in different units (or the same unit) even though they were not, the search image was scaled with the incorrect nominal scale, the angular estimate from the PAFMT can be disregarded regardless of whether it was correct or not, since only correct pairs of angle and scale offset really matter for successful co-registration.

At the same time, there can be instances where the LAFMT inside the QAFMT_TC delivers better rotation estimates than the consequent PAMFT, and the PAFMT inside the QAFMT_TC degrades the results.

In the preceding sections, the output of the QAFMT_LC was evaluated against the thresholds outlined in Table 5.1 to decide if the point clouds were offset by scale and/or rotation. In the following tables, the same formatting was applied as before, to highlight where the QAFMT_V2 outperformed or underperformed the PAFMT.

Table 6.15: Registration of Oshawa height maps using QAFMT_LC

		Oshawa laser and tie point cloud, Scale = 1.0							
		0°		90°		180°		270°	
Laser – Laser	No filter	0.0	1.40	90.0	1.00	-180.0	1.00	90.0	1.00
	Bartlett	0.0	1.00	90.0	1.00	-180.0	1.00	0.0	1.00
	Chebwin	0.0	1.00	0.0	1.00	-180.0	1.00	0.0	1.00
Tie -Tie	No filter	0.0	1.03	89.3	1.00	-180.0	1.00	270.7	1.03
	Bartlett	0.0	1.00	89.3	1.00	-180.0	1.00	273.7	1.00
	Chebwin	0.0	1.00	92.2	1.07	-180.0	1.00	273.7	1.03
Laser -Tie	No filter	352.3	1.03	-180.0	1.10	100.3	1.03	0.0	1.00
	Bartlett	0.0	1.40	-180.0	2.89	0.0	2.89	0.0	2.34
	Chebwin	0.0	1.44	90.0	1.83	-180.0	1.58	0.0	1.35
Tie -Laser	No filter	90.7	1.45	-180.0	1.40	90.7	1.45	0.0	1.27
	Bartlett	175.7	1.40	0.0	2.50	0.0	1.40	-180.0	2.50
	Chebwin	-180.0	1.03	0.0	1.13	0.0	1.03	-180.0	1.28

(a)

		Oshawa DEMs, Scale = 3.2808							
		0°		90°		180°		270°	
Laser – Laser	No filter	0.0	3.27	-89.8	2.78	-180.0	2.73	-89.8	3.29
	Bartlett	0.0	3.31	0.0	3.29	-180.0	3.27	0.0	3.29
	Chebwin	0.0	3.27	0.0	1.14	-180.0	3.27	0.0	1.35
Tie -Tie	No filter	0.0	3.27	-89.8	3.27	-180.0	3.31	-89.8	3.27
	Bartlett	0.0	3.31	0.0	3.27	-180.0	3.31	0.0	3.27
	Chebwin	0.0	3.22	0.0	2.06	-180.0	3.27	-180.0	3.18
Laser -Tie	No filter	-90.0	2.00	0.0	2.00	90.0	2.00	-180.0	1.97
	Bartlett	90.0	1.97	-180.0	2.00	-90.0	2.00	0.0	1.97
	Chebwin	0.0	1.97	90.0	2.00	-180.0	2.92	-90.0	1.63
Tie -Laser	No filter	90.0	4.60	-180.0	4.54	270.0	4.60	0.0	4.54
	Bartlett	270.0	4.71	0.0	4.54	90.0	4.71	-180.0	4.54
	Chebwin	-180.0	1.60	270.0	3.28	0.0	1.23	-90.0	3.28

(b)

Table 6.15a, above, shows the QAFMT_LC co-registration results of the Oshawa tie- and laser point clouds without scale offset. As with the tables in Chapter 6.2.1.3, the upper half of both tables show the self-registration results, and the lower halves show the co-registration results.

When the data was self-registered without scale offset (Table 6.15a), the QAFMT_LC solved 4 angles the LAMFT had missed, yet the LAMFT delivered 1 angle that the QAMFT_TC had missed. In co-registration, the PAFMT recovered 4 angles that the QAFMT_TC had missed, yet the QAFMT provided 4 scale estimates where the LAMFT had failed.

With application of the scale offset (Table 6.15b, above), the PAMFT recovered 5 angles, and the LAFMT recovered 1 angle and 1 scale, where the QAMFT_TC had failed, respectively. When the data was co-registered, the QAMFT_TC found 5 angles that the LAFMT had missed and 3 angles that the PAFMT had missed, and the LAFMT delivered 2 angles and the PAFMT delivered 4 angles that the QAFMT_TC had missed. The rotation estimates were however almost always wrong and the QAFMT_TC also failed to recover scales in 17 instances where the LAFMT had converged successfully.

In total, considering only angles that paired with a correct scale estimate, the LAFMT resulted 3 (1+2) angles correctly that the PAFMT estimated incorrectly, yet the PAFMT produced 8 (4+4) angles that were estimated incorrectly by the LAFMT. Thus, without scale offset, the PAFMT outperformed the LAFMT by estimating an additional 5 angles (28 angles versus 23 angles) correctly.

Table 6.16: Registration of KN height maps using QAFMT_LC

		KN DEMs (61x73, 40x37), Scale = 1.0							
		0°		90°		180°		270°	
Laser - Laser	No filter	0.0	1.00	91.2	1.00	-180.0	1.00	271.2	1.00
	Bartlett	0.0	1.05	88.8	1.00	-165.2	1.05	271.2	1.00
	Chebwin	0.0	1.05	88.8	1.00	-165.2	1.05	271.2	1.00
Tie - Tie	No filter	0.0	1.00	87.8	1.00	-180.0	1.33	92.2	1.00
	Bartlett	0.0	1.00	87.8	1.00	-180.0	1.00	272.2	1.08
	Chebwin	0.0	1.00	87.8	1.00	-180.0	1.00	272.2	1.08
Laser - Tie	No filter	-166.7	1.53	-16.4	1.05	168.2	1.53	14.0	1.40
	Bartlett	0.0	1.00	-180.0	1.70	-180.0	1.00	0.0	1.33
	Chebwin	0.0	1.00	-177.7	1.00	-180.0	1.00	0.0	1.00
Tie - Laser	No filter	-166.7	1.60	79.5	1.15	11.8	1.94	-10.3	1.85
	Bartlett	0.0	1.00	-97.4	3.94	180.0	1.53	0.0	4.33
	Chebwin	0.0	1.39	180.0	1.60	180.0	1.53	0.0	1.53

(a)

		KN DEMs, Scale = 3.2808							
		0°		90°		180°		270°	
Laser - Laser	No filter	0.0	3.30	90.0	3.30	-180.0	3.30	90.0	2.93
	Bartlett	0.0	3.30	88.6	3.30	180.0	3.30	90.0	3.30
	Chebwin	0.0	2.30	88.5	3.11	180.0	2.93	90.0	3.30
Tie - Tie	No filter	0.0	3.18	90.0	3.18	-180.0	3.18	270.0	3.28
	Bartlett	0.0	3.28	90.0	3.18	180.0	3.18	270.0	3.08
	Chebwin	0.0	3.49	90.0	3.08	0.0	2.99	270.0	3.08
Laser - Tie	No filter	9.3	1.59	11.8	1.66	166.0	1.50	14.0	1.51
	Bartlett	0.0	1.92	-180.0	3.42	-180.0	3.15	175.6	3.32
	Chebwin	0.0	2.31	0.0	2.07	180.0	1.92	180.0	1.88
Tie - Laser	No filter	-166.7	4.91	101.8	3.28	11.8	4.91	-10.3	3.28
	Bartlett	0.0	4.91	-97.4	2.53	180.0	4.91	0.0	3.09
	Chebwin	0.0	2.45	180.0	2.53	180.0	1.64	0.0	2.39

(b)

Table 6.16a, above, shows the LAFMT co-registration results of the Keyworth-Nottingham tie- and laser point cloud without scale offset. The upper half shows the self-registration results. Here, the LAFMT showed 100% success. The lower half shows the respective cross-correlation results. Here, the LAFMT converged correctly in 10 out of 24 cases, i.e., it showed a success rate of 42%. Considering only angles that paired with a correct scale estimate, the LAFMT estimated 8 (1+7) angles correctly, where the PAFMT failed. The polar AFMT performed better in only 1 instance. Thus, without scale offset, the LAFMT outperformed the PAFMT by estimating an additional 7 angles (38 angles versus 31 angles) correctly, that paired with a correct scale estimate, too.

Table 6.16b, above, shows the co-registration results for the KN tie- and laser point cloud with scale offset 3.2808. When the KN data was self-registered, the LAFMT failed only in 1 instance, i.e., it converged correctly in 96% of the test cases. When the KN data was co-registered, the LAFMT converged correctly in 16 cases, i.e., it showed a success rate of 67%. In total, the LAFMT resulted 12 (1+11) angles correctly (that paired with a correct scale estimate, too) that the PAFMT estimated incorrectly, and the PAFMT produced 3 (1+2) angles correctly that were estimated incorrectly by the LAFMT. Thus, with scale offset, the LAFMT outperformed the PAFMT by estimating an additional 9 angles (39 angles versus 30 angles) correctly.

The overall success rate of the co-registration of KN data with LAFMT was 55%. In total, the LAFMT outperformed the PAFMT and estimated an additional 9 (8) angles, that also paired with a correct scale estimate, when (no) scale offset was applied, respectively.

Table 6.17: Registration of simulated height maps using QAFMT_LC

		Simulated DEMs, Scale = 1.0							
		0°		90°		180°		270°	
DTM - DTM	No filter	0.0	1.00	0.0	1.00	90.0	1.00	180.0	1.00
	Bartlett	0.0	1.00	90.0	1.00	180.0	1.01	270.0	1.02
	Chebwin	0.0	1.00	90.0	1.00	180.2	1.01	270.0	1.01
DSM - DSM	No filter	0.0	1.30	0.0	1.00	180.0	1.00	-180.0	1.00
	Bartlett	0.0	1.00	89.8	1.00	180.2	1.00	271.2	1.00
	Chebwin	0.0	1.00	-90.2	1.00	-0.2	1.00	271.2	1.00
DTM - DSM	No filter	0.0	1.00	0.0	1.00	-180.0	1.00	-180.0	1.00
	Bartlett	0.0	2.49	0.0	2.21	180.0	4.29	-180.0	2.92
	Chebwin	0.0	2.71	-180.0	3.70	0.0	2.67	270.0	2.91
DSM - DTM	No filter	0.0	1.00	0.0	1.00	-180.0	1.00	-180.0	1.00
	Bartlett	0.0	1.40	-180.0	2.34	180.0	2.47	0.0	2.53
	Chebwin	0.0	1.02	0.0	1.65	180.0	2.49	0.0	2.43

(a)

		Simulated DEMs, Scale = 3.2808							
		0°		90°		180°		270°	
DTM - DTM	No filter	0.0	3.28	90.0	3.28	-180.0	3.28	270.0	3.28
	Bartlett	0.0	3.27	90.0	3.27	180.0	3.27	270.0	3.27
	Chebwin	0.0	3.27	90.0	3.27	-180.0	3.27	270.0	3.27
DSM - DSM	No filter	0.0	3.28	90.0	3.28	-180.0	3.28	-90.0	3.28
	Bartlett	0.0	3.28	88.8	3.28	180.2	3.28	271.2	3.28
	Chebwin	0.0	3.28	91.5	3.27	180.2	3.28	271.2	3.27
DTM - DSM	No filter	0.0	3.28	0.0	3.28	-180.0	3.28	-180.0	3.28
	Bartlett	0.0	3.28	0.0	3.28	180.0	3.28	-180.0	3.28
	Chebwin	0.0	3.28	-180.0	3.28	0.0	3.28	0.0	3.28
DSM - DTM	No filter	0.0	2.95	0.0	1.55	-180.0	2.95	-180.0	2.95
	Bartlett	0.0	2.87	-180	2.72	180.0	3.28	0.0	3.28
	Chebwin	0.0	1.68	0	3.28	180.0	2.15	0.0	3.28

(b)

Table 6.17a, above, shows the LAFMT co-registration results of the simulated tie- and laser point cloud without scale offset. When the data was self-registered, the LAFMT failed only in 2 instances and showed a 92% success rate. The lower half shows the respective cross-correlation results. Here, the LAFMT converged correctly in 8 out of 24 cases, i.e., it showed a success rate of 34%. Considering only angles that paired with a correct scale estimate, in total the PAFMT resulted 3 (2+1) angles correctly that the LAFMT estimated incorrectly. In total, the PAFMT outperformed the LAFMT and estimated an additional 13 angles (43 angles versus 30 angles) correctly.

Table 6.17b, above, shows the co-registration results of the simulated tie- and laser point cloud with scale offset 3.2808. When the data was self-registered, the LAFMT showed 100% success. When the simulated data was co-registered, the LAFMT converged correctly in only 4 cases, i.e., it showed a success rate of 17%. In total, the LAFMT resulted 2 angles correctly that the PAFMT estimated incorrectly, and the PAFMT produced 3 angles correctly that were estimated incorrectly by the LAFMT. Thus, with scale offset, the PAFMT outperformed the LAFMT angle estimates slightly.

The overall success rate of the co-registration of simulated data with the LAFMT was 26%. In total, the PAFMT outperformed the LAFMT' angles slightly and estimated an additional 3 (3) angles that paired with a correct scale estimate, when a (no) scale offset was applied, respectively.

6.2.2.4 Discussion

In section 6.2.2 and its sub-sections, the QAFMT_LC was introduced as a method that combines the LAFMT and PAFMT with the intend to syndicate the strength of both and improve the overall co-registration results. Like the QAFMT_TC, the QAFMT_LC was designed to fix the estimated scale offset at either 1.0 or 3.2808 (or its inverse, 0.3048), since this scale difference is common in geodetic coarse co-registration applications. Thus, the QAFMT_LC deploys the LAFMT first to fix the scale at a nominal value, too, but then, unlike the QAFMT_TC, it gets its rotation estimate always from the PAFMT.

As in the preceding experiments, bilinear interpolation was used in the QAFMT_TC tests, and the Bartlett filter proofed to create overall the best results.

While the LAFMT created many correct scale estimates (Oshawa 79.2%, KN 66.7%, Simulated 35.4%), a lot of them coupled with an incorrect rotation estimate (correct scale but wrong angle estimate: Oshawa 64.6%, KN 12.5%, Simulated 10.4%). Thus, the motivation of creating the QAFMT_LC was to combine the PAFMT's many correct estimates of the angular offset with the many correct scale estimates of the LAFMT, and perhaps result more instances where a correct LAFMT scale pairs with a correct PAFMT angle compared to the instances where a correct LAFMT scale pairs with a correct LAFMT angle estimate (disproving "hypothesis 2").

The experiments revealed that the QAFMT_LC created slightly fewer but more evenly distributed correct rotation estimates than the PAFMT (QAFMT_LC: 18.8% (Oshawa), 43.8% (KN), 47.9% (Simulated); vs PAFMT 12.5% (Oshawa), 54.2% (KN), 83.3% (Simulated)). The QAFMT_LC delivered also more correct scale estimates than the LAFMT (QAFMT_LC: 77.1% (Oshawa), 68.8% (KN), 64.6% (Simulated); vs LAFMT: 79.2% (Oshawa), 66.7% (KN), 35.4% (Simulated)), but many correct scale estimates paired with an incorrect rotation estimate.

In the end, the QAFMT_LC delivered more evenly distributed yet slightly fewer pairs of correct scale and correct angle estimates than the LAFMT (QAFMT_LC: 14.6% (Oshawa), 31.3% (KN), 35.4% (Simulated); vs LAFMT: 14.6% (Oshawa), 54.1% (KN), 25.0% (Simulated)). Thus, the QAFMT_LC did not turn out as clear, strong winner over the LAFMT; rather did the LAFMT win by a small margin since it delivered very similar co-registration output yet required less computational steps. This finding confirms the earlier findings of Section 6.2.1.4, supports "hypothesis 2" (Section 6.1.4.5), and rules in favor of the QAFMT_TC for these three datasets.

6.3 Summary

In the 2000s, the AFMT, - a modified FMT that maintains the phase information -, was proposed. Thus, repeating some of the earlier FMT experiments with the AFMT should provide more insight into the impact of the loss of phase information versus the impact of elevation profile on the co-registration results. Also, since the Fourier phase (that is preserved in the AFMT) carries

more information than the Fourier magnitude, it was expected that the AFMT delivers better results than the FMT.

Two forms of the AFMT have been examined: The PAFMT and the LAFMT. As expected, both the PAFMT and LAFMT showed a higher success rate than the FMT. The LAFMT delivered better rotation estimates than the PAMFT for the KN data, which could be attributed to the stronger phase signal of the KN data, and the PAFMT surpassed the LAFMT with the Oshawa and simulated data. Overall, the LAFMT proved to be the preferred method for co-registration since it delivers both rotation and scale and since its rotation estimates lagged the PAFMT angle estimates mostly only slightly.

Lastly, the EAFMT was introduced as an extended AFMT that can – like the FMT – recover not only rotation and scale, but also translation. But since the translations are recovered with Cartesian height images in an internal POMF, the range of translations that can be recovered with the EAFMT is smaller than that of the FMT. The most important drawback of the EAFMT, and also LAFMT and PAFMT, is however that they were not optimized for the presence of two units. Knowing that the laser and tie point clouds can only be either referenced in the same unit or offset by a scale of 3.2808, better optimized algorithms could be designed that account for this condition. Thus, Chapter 6.2 introduced the loosely- and tightly coupled QAFMT as novel co-registration algorithms along with the reasons why they should outperform the existing methods in global co-registration of airborne laser- and sparse tie- point clouds. Chapters 6.2.1.3, 6.2.1.4 and 6.2.2.3, 6.2.2.4 presented and discussed the respective co-registration results. In the ensuing paragraphs, these results have been compared against all co-registration results from the methods introduced earlier. Also, the results have been evaluated against the criteria outlined in Chapter 1.1 what optimal characteristics a good co-registration algorithm should have.

As with preceding analysis, this performance assessment was focused on the co-registration results (and not the self-registration results) since they necessarily imply practical geodetic application. While the count of individual correctly estimated co-registration parameters could be used to evaluate the co-registration success, the number of correctly estimated parameter pairs is a better performance metric, since only correct parameter pairs allow for successful co-registration. After this metric, Table 6.18 compares the QAFMT output against the earlier introduced frequency-based registration methods.

Table 6.18: Number of correct co-registration parameter pairs (angle- and respective scale estimate; min=0, max=4)

		LAFMT		PAFMT		QAFMT_TC 20% nth CPS		QAFMT_TC 50% nth CPS		QAFMT_LC		Total	
		Laser- Tie	Tie- Laser	Laser- Tie	Tie- Laser	Laser- Tie	Tie- Laser	Laser- Tie	Tie- Laser	Laser- Tie	Tie- Laser		
Oshawa	Scale=1.0	None	0	0	1	0	0	0	0	1	0	18	
		Bartlett	1	1	1	0	1	1	1	1	0		
		Chebwin	1	0	1	0	1	3	1	0	1		0
	Scale=3.3	None	0	1	0	0	0	0	1	1	0	0	17
		Bartlett	0	3	0	0	0	0	0	3	0	0	
		Chebwin	0	0	4	0	0	0	0	0	4	0	
KN	Scale=1.0	None	1	0	2	3	0	0	2	1	0	1	43
		Bartlett	3	1	2	2	1	1	1	2	2	1	
		Chebwin	4	1	2	2	1	0	2	2	2	1	
	Scale=3.3	None	1	2	1	1	0	1	0	4	0	1	52
		Bartlett	3	4	2	2	0	2	3	4	2	2	
		Chebwin	4	2	2	2	0	0	2	1	2	2	
Simulated	Scale=1.0	None	4	4	2	2	2	2	2	2	3	3	52
		Bartlett	0	0	4	4	0	3	0	3	1	1	
		Chebwin	0	0	4	4	0	0	0	0	1	1	
	Scale=3.3	None	0	1	2	2	1	1	2	4	2	2	47
		Bartlett	0	1	2	2	2	3	1	1	2	2	
		Chebwin	0	2	1	2	1	2	0	3	1	2	
Sums		22	23	33	28	10	16	18	32	25	19		
Total		45		30.5 (=61/2 ⁽⁴⁾)		26		50		44			

As shown in Table 6.18, the QAFMT_TC with the threshold of 50% (column highlighted in green) scored the highest and resulted 50 out of 144 possible correct co-registration parameter pairs. More, the QAFMT_TC with 50% threshold outperformed the LAFMT and the QAFMT_LC by 5 and 6 parameter pairs, respectively. Thus, it delivered a 7% better co-registration result than the rivaling algorithms.

Table 6.18 also shows that the QAFMT_TC with 20% threshold recovered only half as many parameter pairs than the QAFMT_TC with 50% threshold. While it was anticipated that the

⁴ Only angle estimates can be considered in the comparison to PAFMT output

QAFMT_TC would perform better with broader bandwidth, it was a surprise to see such a big difference in performance.

The table also reveals that the KN and simulated data allowed for twice as much co-registration success as the Oshawa data. This can be explained by the increased surface dynamic in the KN and simulated data compared to the Oshawa height maps. While the KN data performed better in presence of a scale offset (52 vs 43 correct pairs), the simulated height map worked better without scale offset (47 vs 52 correct pairs). More, the PAFMT performed slightly better in the absence of scale (33 vs 28 correct angles). These phenomena cannot be explained easily from this table alone and require more detailed information. Similarly, from this table it is difficult to clearly distinguish which filter worked best. Table 6.19 to Table 6.22 below break the above results up into more detail, which helped answer these questions.

Table 6.19: Number of correct parameter pairs (min=0, max=4), self-registration of laser point clouds

Self-registration of laser point clouds		Oshawa		KN		X5d4		Total
		Scale 1.0	Scale 3.3	Scale 1.0	Scale 3.3	Scale 1.0	Scale 3.3	
FMT (angle and scale only)	None	2	0	3	1	n/a	n/a	(6)
	Bartlett	4	2	2	1	n/a	n/a	(9)
	Chebwin	2	2	2	1	n/a	n/a	(7)
LAFMT (angle and scale)	None	3	2	4	3	4	4	20
	Bartlett	1	4	4	4	4	4	21
	Chebwin	4	1	4	4	2	4	19
QAFMT_TC with 20% threshold (angle and scale)	None	3	2	4	3	1	4	17
	Bartlett	3	4	4	4	4	3	22
	Chebwin	2	2	3	3	4	2	16
QAFMT_TC with 50% threshold (angle and scale)	None	3	2	4	4	1	4	18
	Bartlett	3	4	4	4	4	4	23
	Chebwin	2	1	3	4	4	4	18
QAFMT_LC (angle and scale)	None	3	2	4	4	4	4	21
	Bartlett	1	4	4	4	4	4	21
	Chebwin	4	1	4	4	2	4	19
PAFMT (angle only)	None	3	n/a	4	n/a	1	n/a	(16)
	Bartlett	3	n/a	4	n/a	4	n/a	(11)
	Chebwin	2	n/a	4	n/a	4	n/a	(10)
Total		30	29	46	45	38	45	

As can be seen from Table 6.19, the QAFMT_TC with the threshold of 50% delivered an almost perfect score for laser point clouds (the column highlighted in green shows 23 out of 24 possible correct self-registration parameter pairs). Similarly, the QAFMT_TC with the threshold of 20% resulted 22 and the QAFMT_LC and LAFMT resulted 21 correct pairs, respectively. All algorithms performed best when the Bartlett filter was applied.

The last row of Table 6.19 confirms what could already be observed from Table 6.18: The datasets with stronger undulation yield better results. However, unlike in co-registration (Table 6.18), self-registration of the simulated data performed better in presence of scale.

Table 6.19 presents also results from the instances where self-registration of the laser point clouds was performed with FMT. As described in Section 5.1.2 onwards, the FMT performed quite poorly in self-registration mode, and the respective co-registration experiments were therefore

skipped; since processing the simulated data with FMT proved to be time-intensive, only self-registration results from real-world data are listed above. Considering these FMT test cases, the FMT succeeded only in 22 out of 72 instances (30.56%), while the laser point clouds self-registered successfully in 235 out of 288 test cases (81.25%) when AFMTs (LAFMT, QAFMT_TC 20%, QAFMT_LC) were deployed, and in 59 out of 72 instances (81.95%) when the QAFMT_TC with the threshold of 50% was used.

For comparison, the respective self-registration results for the tie point clouds are displayed in Table 6.20 below.

Table 6.20: Number of correct parameters pairs (min=0, max=4), self-registration of tie point clouds

Self-registration of tie point clouds		Oshawa		KN		X5d4		Total
		Scale 1.0	Scale 3.3	Scale 1.0	Scale 3.3	Scale 1.0	Scale 3.3	
FMT (angle and scale only)	None	2	0	3	1	n/a	n/a	(7)
	Bartlett	4	2	2	1	n/a	n/a	(9)
	Chebwin	2	2	2	1	n/a	n/a	(7)
LAFMT (angle and scale)	None	4	4	4	4	4	4	24
	Bartlett	4	4	4	4	4	4	24
	Chebwin	4	4	4	4	4	4	24
QAFMT_TC with 20% threshold (angle and scale)	None	4	4	3	3	1	4	19
	Bartlett	4	4	4	4	4	4	24
	Chebwin	4	3	4	4	4	4	23
QAFMT_TC with 50% threshold (angle and scale)	None	4	4	3	3	4	4	22
	Bartlett	4	4	4	4	3	4	23
	Chebwin	4	4	4	4	2	4	22
QAFMT_LC (angle and scale)	None	4	4	4	4	4	4	24
	Bartlett	4	4	4	4	4	4	24
	Chebwin	4	4	4	4	4	4	24
PAFMT (angle only, *)	None	4	n/a	3	n/a	2	n/a	(9)
	Bartlett	4	n/a	4	n/a	4	n/a	(12)
	Chebwin	4	n/a	4	n/a	2	n/a	(10)
Total		48	47	46	46	42	48	

As can be seen from Table 6.20, almost all algorithms self-registered the laser point cloud perfectly when the Bartlett filter was applied: The QAFMT_TC with the threshold of 20%, the QAFMT_LC, the LAFMT, and the PAFMT delivered a perfect score. The QAFMT_TC with the threshold of 50% missed in one instance only, which is why it was highlighted with a light green color. The QAFMT_LC was highlighted in solid green to indicate that it delivered the overall best results here, considering also that it performed best in co-registration of the laser- and tie point clouds, as shown in Table 6.21 further below.

At first sight, the success rate of self-registration of the Oshawa tie point cloud (shown in the last row of Table 6.20) is surprisingly high, about twice as high than before for the Oshawa laser point cloud (Table 6.21). This is because a lot of elevation changes occur within the footprint of the tie point cloud, whereas the laser data stretches substantially beyond this area and covers a lot of ground with low elevation dynamics. The laser data also covers a very small area in the south-east corner with a strong slope (see APPENDIX C for details), however this slope acts like a rotary needle point, which makes the recovery of angular offsets difficult. Thus, the last row of Table 6.20 confirms what could already be observed from Table 6.18 and Table 6.19: *Datasets with stronger elevation dynamic are more likely to register successfully. This is also indicated by the fact that in presence of strong undulation, a scale offset hardly alters the co-registration results.*

As before with the self-registration of the laser point clouds, the FMT performed relative poorly with the real-world tie points, too, for the reasons mentioned in APPENDIX C. Considering the test cases where FMT was performed, the FMT succeeded only in 23 out of 72 instances (31.94%), while the tie point clouds self-registered successfully in 276 out of 288 test cases (95.83%) when AFMTs (excluding PAFMT) were deployed, and in 67 out of 72 instances (93.05%) when only the QAFMT_TC with the threshold of 50% was used.

Table 6.21: Number of correctly estimated co-registration parameter pairs of laser- and tie point cloud

Co-registration of Laser and tie point cloud		Oshawa		KN		X5d4		Total
		Scale 1.0	Scale 3.3	Scale 1.0	Scale 3.3	Scale 1.0	Scale 3.3	
FMT (angle and scale only)	None	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Bartlett	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Chebwin	n/a	n/a	n/a	n/a	n/a	n/a	n/a
LAFMT (angle and scale)	None	0	0	1	1	4	0	6
	Bartlett	1	0	3	3	0	0	7
	Chebwin	1	0	4	4	0	0	9
QAFMT_TC with 20% threshold (angle and scale)	None	1	0	2	0	2	2	7
	Bartlett	1	0	2	2	0	1	8
	Chebwin	1	0	2	0	0	0	3
QAFMT_TC with 50% threshold (angle and scale)	None	1	1	2	0	2	4	10
	Bartlett	1	0	2	3	3	1	10
	Chebwin	1	0	2	2	0	3	8
QAFMT_LC (angle and scale)	None	0	1	1	1	4	0	7
	Bartlett	1	1	3	3	0	0	8
	Chebwin	1	1	4	4	0	2	12
PAFMT (angle only)	None	1	n/a	2	n/a	2	n/a	(5)
	Bartlett	1	n/a	2	n/a	4	n/a	(7)
	Chebwin	1	n/a	2	n/a	4	n/a	(7)
Total		10	4	28	23	15	13(12)	

As can be seen from Table 6.21, the cross-correlation delivered less than half as many matches as the self-registration tests. Since the laser- and respective tie point clouds differ in size and surface characteristics, it was expected they yield a lower success rate in cross-correlation, however it was a surprise to see that complete co-registration occurred in only two instances (KN data with QAFMT_LC and LAFMT with application of the Chebyshev filter).

While the QAFMT_LC showed the best turnout (highlighted in solid green), the QAFMT_TC with 50% threshold was highlighted in light green to indicate that - even though its performance falls in second place for this test configuration - it performed well overall with all tests. As shown in the last row, the KN data performed much better than the Oshawa and simulated data. Unlike the Oshawa test site, the KN surface is relatively small and shows a lot of height difference over

short distance. Also, compared to the KN height maps, the simulated X5d5 DSM has a very noisy surface. Obviously, two smooth surfaces co-register better than a very noisy and a smooth surface (X5d4 DTM), which is why the co-registration of the KN point clouds performed much better than the co-registration of the simulated height models.

As described in 5.1.2, the FMT performed quite poorly in self-registration mode, and the respective co-registration experiments were therefore skipped (they also required considerably more processing time). Thus, there are no FMT co-registration results listed in Table 6.21. The laser- and tie point clouds co-registered successfully in 86 out of 288 test cases (29.86%) when AFMTs (excluding PAFMT) were deployed, and in 28 out of 72 instances (38.89%) when only the QAFMT_TC with the threshold of 50% was used.

Table 6.22: Number of correctly estimated co-registration parameter pairs (tie- and laser point cloud)

Co-registration of tie-laser point cloud		Oshawa		KN		X5d4		Total
		Scale 1.0	Scale 3.3	Scale 1.0	Scale 3.3	Scale 1.0	Scale 3.3	
FMT (angle and scale only)	None	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Bartlett	n/a	n/a	n/a	n/a	n/a	n/a	n/a
	Chebwin	n/a	n/a	n/a	n/a	n/a	n/a	n/a
LAFMT (angle and scale)	None	0	1	1	2	4	1	9
	Bartlett	1	3	1	4	0	1	10
	Chebwin	0	0	1	2	0	2	5
QAFMT_TC with 20% threshold (angle and scale)	None	0	0	3	1	2	4	10
	Bartlett	1	0	1	1	3	1	7
	Chebwin	0	0	2	0	0	0	2
QAFMT_TC with 50% threshold (angle and scale)	None	0	1	1	4	2	4	12
	Bartlett	0	2	2	4	3	1	12
	Chebwin	0	0	1	1	0	3	5
QAFMT_LC (angle and scale)	None	0	1	0	2	4	1	8
	Bartlett	1	3	1	4	0	0	9
	Chebwin	0	0	1	2	0	1	4
PAFMT (angle only)	None	0	n/a	3	n/a	2	n/a	5
	Bartlett	0	n/a	2	n/a	4	n/a	6
	Chebwin	0	n/a	2	n/a	4	n/a	6
Total		3	11	15	27(26)	18	19	

As presented by the green band in Table 6.22, the QAFMT_TC with the threshold of 50% and Bartlett filter outperformed the other co-registration algorithms again. As before in Table 6.21, the KN data delivered the highest score, as shown in the last column of Table 6.22. The tie- and laser point clouds co-registered successfully in 90 out of 288 test cases (31.25%) when AFMTs (excluding PAFMT) were deployed, and in 29 out of 72 instances (40.28%) when only the QAFMT_TC with the threshold of 50% was used.

As mentioned before, at first sight one might expect identical co-registration results, regardless of whether the laser- or the tie point cloud was used as reference point cloud. However, the extend of the zero-padded areas can vary slightly from case to case (depending on the dimensions of the original footprint and the applied scale- and angular offset), which can lead to slightly different effects of interpolation and application of window mask. If the data allows for strong correlation, the AFMT's CPS tends to have only few distinct spikes. However, if the data input is noisy or flat, the CPS surface tends to yield spikes of similar magnitude. In this case, the aforementioned minute differences in data input can invoke a slight change of the CPS and respective change of the spatial distribution of the 20 tallest CPS peaks.

Having this background in mind, one can conclude that the co-registration results in Table 6.21 and Table 6.22 are equivalent (40.28% and 38.89% co-registration success for the QAFMT_TC with wide bandwidth).

Further, the QAFMT_TC with 50% bandwidth succeeded approximately half as often in co-registration than in self-registration mode (overall 39.59% success rate in co-registration vs 87.50% in self-registration). The experiments also show that all algorithms performed best when the Bartlett filter was applied, and that datasets with stronger undulation yield better results. Furthermore, the presence of a scale offset did, on average, not alter the outcome, which confirmed that QAFMT is robust wrt scale even at low SNR.

Comparing these findings to the initial requirements for co-registration outlined in Chapter 1.1 (“a matching algorithm should be robust to low SNR, accurate even when the same scenery has different intensities or different orientations and sizes, and efficient numerically” (Kim and De Araújo, 2007), we can conclude that out of all algorithms that were researched in this thesis, the QAFMT_TC with 50% internal threshold and Bartlett filter meets these requirements best. It is superior to the other algorithms because it

- considers all data at once in search of a global maximum
- maintains the phase information
- is numerically efficient
- is optimized for geodetic scale offsets

and consequently, it showed the best performance, as it

- self-registered with 87.50% success
- co-registered with 39.95% success
- outperformed the standard FMT in self-registration by over 50%
- outperformed the LAFMT by approx. 7%.

Even though the QAFMT_TC with 50% bandwidth showed better results than all other methods, it still did not provide the level of co-registration success hoped for, i.e., something near 100%.

However, this lack of success can be attributed primarily to limitations of input data, such as

- tie- and laser points describing different surface elements
- lack in elevation differences in the height maps
- extreme levels of noise in the simulated data (in hindsight, the noise in X5d5 was too exaggerated compared to its point density and footprint)
- extremely sparse point density of the tie points compared to the laser points (approx. 1:10,000) and compared to the data footprint. (While the laser point clouds had different point densities, originated from different sensors, and were processed with different settings, the tie point clouds were generated with the same standard setting by the same tie point engine. So, while it would have been possible to alter the processing settings and create tie point clouds with higher density, this extra step was not considered during data preparation.)

rather than to shortcomings in the algorithm.

7.0 Conclusions and Recommendations

7.1 Conclusions

Co-registration is required for accurate data fusion and for cross-aided sensor calibration. As the vast volume of the published research over the past decades shows, the fully automatic co-registration of airborne laser scans and imagery is challenging. The increasing number of datasets from different sensor architectures should be supported with robust global co-registration methods no longer subject to any assumptions about the initial position and projection.

Many methods have been evolved, but none of them allows for a reliable fully automatic coarse co-registration without any metadata, and moreover, the existing fine registration schemes do not work for all large-format airborne applications (APPENDIX D). The existing registration methods are computationally intense, lack robustness, and/or need specific additional processing steps, which are not part of the standard camera and LiDAR data processing workflow (Chapter 2.3.1 and Chapter 2.3.2). Surprisingly, the use of the standard photogrammetric tie points as tie features for co-registration has been overlooked by academia and industry.

This dissertation offered a framework to innovate both the automatic coarse- and fine co-registration of large-area airborne LiDAR and camera data using the photogrammetric tie points as its tie features. Multiple space domain-based and time domain-based methods were proposed and examined. This dissertation delivered nine achievements:

- 1) The photogrammetric tie points, in their native form and in form of the height map they describe, have been utilized as the registration elements in the integration (coarse and fine co-registration with- and without simultaneous cross-aided sensor-driven calibration) of the airborne LiDAR and camera data. Since the respective research proposal was submitted in 2013 at York University, only two publications (Zheng *et al.*, 2013; Gneeniss *et al.*, 2014) have gotten involved in a similar idea, yet they relied on metadata and deployed dense tie points under less challenging conditions

(Chapter 2.3.2), which have been reviewed and compared to the proposed use of tie points in Chapter 2.0.

2) A systematic and thorough research was conducted as to how the airborne tie points from traditional large area surveys can be used for a fully automated co-registration (Chapter 4.0 and Chapter 5.0). At the time of writing this dissertation, this has been the first work that thoroughly researched the use of the extremely sparse photogrammetric tie points in the context of the co-registration and cross-aided sensor calibration. This has also generated a new knowledge about the use of non-geodetic tools for solving this problem.

3) The utilization of seven space domain-based algorithms, that are typically applied outside the geodetic domain, as the tools in the coarse co-registration was investigated (Chapter 4.2 and Chapter 4.3). The concepts of the differential surface analysis, fractal analysis, probabilistic analysis using graphical models, use of entropy, and other procedures were introduced and analyzed in terms of their suitability to transform the laser- and optical point clouds into a common coordinate system. This analysis showed that the laser points and photogrammetric tie points, - especially the image-based (non-key-point derived) tie points, - suffer from the extreme discrepancy in their point densities and the differences in their point locations. The discrepancy between their point densities can be in an order of thousands within the same surveying area, which proved to be the downfall for these space-domain based methods.

4) Two new methods for the fine-registration and cross-aided calibration of the LiDAR- and camera sensors and their data in space domain were proposed and systematically analyzed (Chapter 4.4.2 and Chapter 4.4.3). The concepts relied on using laser planes and photogrammetric tie points as tie elements. The analysis showed that more similarity in their point density and more compatible characteristics of their tie elements is required so that the impact of the missing parameters for the co-registration, and for the sensor- and system calibration can be overcome.

5) Three new methods for the fine registration (without the cross-aided sensor calibration) of LiDAR and camera data in the space domain were proposed and systematically analyzed (Chapter

4.4.4 and Chapter 4.4.5), and it was found that higher point densities and more similar point characteristics are required so that automatic fine-registration can succeed.

6) A systematic analysis was conducted as to what conditions the point clouds must meet so that the fine registration algorithms can converge successfully (Chapter 4.4.5.2). This research led to a quantitative description as to why the accurate fine registration with extremely sparse tie point clouds from traditional large area airborne surveys is impossible, and what conditions must be met, so that the registration can succeed using these tie elements (100 times higher point density and an approx. 10 times larger depth is required).

7) The use of frequency domain-based algorithms in the coarse co-registration was researched and, accordingly, the FMT algorithm was proposed as an efficient means to retrieve rotation, scale, and translational offsets (Chapter 5.1). Traditionally, this algorithm has been applied in medical- and computer vision. A detailed analysis of the FMT studied the reasons for its low self-registration performance with the airborne sparse point clouds and revealed the loss of phase information (and consecutive strong correlation of FMT-internal parameters) as its main root cause.

8) The AFMT with recursive POMF was studied as an improved version of the FMT (that maintains the phase information of the input signals) and a more efficient algorithm for the global co-registration in frequency domain (Chapter 6.1). The application of an expanded, recursive version of the AFMT (in form of the LAFMT and PAFMT) was applied to the laser- and tie point clouds and resulted in an improved co-registration performance and provenly a better suited algorithm to implement reliable, fully automated co-registration for this type of data.

9) Lastly, the QAFMT was introduced as a novel effective frequency domain algorithm for the coarse co-registration (Chapter 6.2). Three variants of the QAFMT were proposed: The QAFMT_TC with a small (20%) and large (50%) internal threshold setting, and the QAFMT_LC. All of them combine the LAFMT and PAFMT with the intend to improve the co-registration of airborne point clouds. As the tests with three datasets showed, the QAFMT_TC with 50% internal margin outperformed the LAFMT co-registration output by approx. 10%. Thus, amongst the algorithms researched in this dissertation, the QAFMT_LC is proven to be the most promising

algorithm for the coarse co-registration of the traditional large area airborne laser- and photogrammetric tie points.

To conclude along the remarks of Zitova and Flusser (2003): *“In the future, the idea of an ultimate registration method, able to recognize the type of given task and to decide by itself about the most appropriate solution, can motivate the development of expert systems. They will be based on the combination of various approaches, looking for consensus of particular results.”* As this thesis has demonstrated, the QAFMT_TC with use of photogrammetric tie points could be a useful addition to this database of co-registration algorithms.

7.2 Recommendations for future Work

Automatic co-registration remains an active topic in research and development. There is abundant room for continuous improvement of the co-registration methods. Whilst the problem of the fine registration has been solved for many applications, there is still room and need for improvement in special cases, such as the simultaneous co-registration and cross-aided airborne sensor calibration. Since photogrammetric tie points come with high precision and in large quantity, they possess potential to make good tie elements for the automatic integration of camera and LiDAR data, if the scenery provides strong height variations. Thus, the idea of using the tie points as co-registration elements could be revisited for the UAV and close-range mapping applications.

Similarly, the modern photogrammetric dense point engines can generate a dense grid of 3D points in the object space with the same accuracy as the tie points (Haala, 2013). While dense point matching is computationally intense and not part of the standard photogrammetric post-processing chain, the dense points assimilate better with the surface of the laser scan and they have great potential to be part of a highly efficient fine registration method: As outlined by Lindenthal *et al.* (2015), a single control point on a LiDAR roof polygon is sufficient to fully co-register the self-calibrated camera- and lidar data. Thus, the dense point matching could occur in a few small areas where tilted roof planes locate in the LiDAR data. Matching a few roof polygons from the laser point cloud and from the locally generated photogrammetric dense point cloud would allow for

the simultaneous cross-aided sensor calibration and co-registration with very little computational effort. This approach seems highly effective and efficient and should be explored in the future.

Lastly, the QAFMT algorithms are new and need to pursue more exploration. Within the scope of this dissertation, due to the time limit and other constraints, only datasets from three different sites were taken for evaluation, yet the algorithms could be optimized further when more data can be put to the test. Having more data could also help in the search of a specification parameter that can indicate the quality and reliability of the co-registration results. Further, as mentioned in Chapter 6.1.3.1, the QAFMT and the other AFMT versions could also be expanded to compensate for the effect of the swapped coordinate axes (X- and Y-axes are flipped in the moving and target point cloud). To conclude, future work could also be focused on customizing the AFMT derivatives for specific co-registration applications such as mapping of harbours, airports, large buildings, construction sites, and greenspaces.

8.0 References

- Abdelazeem, M., Elamin, A., Afifi, A. and El-Rabbany, A., 2021. Multi-sensor point cloud data fusion for precise 3D mapping. *The Egyptian Journal of Remote Sensing and Space Science*.
- Akca, D., 2007. Least Squares 3D surface matching (Vol. 92). *Dissertation*. ETH Zurich.
- Ackermann, F., 1992. Operational rules and accuracy models for GPS-aero triangulation. *Arch. ISPRS*, 1, pp.691-700.
- Alba, M., Barazzetti, L., Scaioni, M. and Remondino, F., 2011, September. Automatic registration of multiple laser scans using panoramic RGB and intensity images. In *Proceedings of the ISPRS Workshop Laser Scanning*.
- Albert Albiol, E., Arenas Sánchez, P., Genaim, S. and Puebla, G., 2015. A practical comparator of cost functions and its applications. *Science of Computer Programming 2015*.
- Al-Durgham, M., Datchev, I. and Habib, A., 2011. Analysis of two triangle-based multi-surface registration algorithms of irregular point clouds. *ISPRS-International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 3812, pp.61-66.
- Al-Durgham, K., 2014. Geometric Features Extraction and Automated Registration of Static Laser Scans Using Linear Features. *Dissertation*, University of Calgary, Department of Geomatics Engineering.
- Aldegalawy, M., Datchev, I.D. and Habib, A.F., 2008, April. Alternative procedures for the incorporation of LiDAR-derived linear and areal features for photogrammetric geo-referencing. In *Proc. Amer. Soc. Photogramm. Remote Sensing (ASPRS) Annu. Conf.*
- Al-Ruzouq, R., 2004. Semi-automatic registration of multi-source satellite imagery with varying geometric resolutions. *Dissertation*, University of Calgary, Department of Geomatics Engineering.
- Angelats, E., Blázquez, M., Colomina, I., de la Tecnología, P.M. and Gauss, A.C.F., 2012. Simultaneous orientation and calibration of images and laser point clouds with straight segments. *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 39, p.B1.
- Armenakis C. 1988. Displacement Monitoring by Integrating On-Line Photogrammetric Observations with Dynamic Information. *Dissertation*, Department of Surveying Engineering (now Geodesy and Geomatics Engineering), University of New Brunswick, 284p.; Also: Technical Report No 133, Department of Geodesy and Geomatics Engineering, University of New Brunswick, November 1987.

- Armenakis, C., 2000. Differential approach for map revision from new multi-resolution satellite imagery and existing topographic data. *International Archives of Photogrammetry and Remote Sensing*, 33(B4/1; Part 4), pp.99-104.
- Armenakis C., Gao, Y., and G. Sohn, 2010. 'Co-registration of LiDAR and photogrammetric data for updating building databases', *ISPRS Joint Workshop on Core Spatial Databases-Updating, Maintenance and Services-from Theory to Practice*, ISPRS Archives, Vol. XXXVIII, Part 4-8-2-W9.
- Armenakis, C., Gao, Y. and Sohn, G., 2013. Co-registration of aerial photogrammetric and LiDAR point clouds in urban environments using automatic plane correspondence. *Applied Geomatics*, 5(2), pp.155-166.
- Armenakis, C. and Nirupama, N., 2014. Flood risk mapping for the city of Toronto. *Procedia Economics and Finance*, 18, pp.320-326.
- Arun, P.V., 2013. A comparative analysis of different DEM interpolation methods. *The Egyptian Journal of Remote Sensing and Space Science*, 16(2), pp.133-139.
- Baltsavias, E.P., 1999. A comparison between photogrammetry and laser scanning. *ISPRS Journal of photogrammetry and Remote Sensing*, 54(2), pp.83-94.
- Bang, K. I. 2010a. Alternative Methodologies for LiDAR System Calibration. *Dissertation*. September 2010, UCGE Report 20317, University of Calgary, Canada
- Bang, K.I. 2010b. Alternative Methodologies for LiDAR System Calibration. *Presentation at the Department of Earth and Space Science and Engineering York University*, 5 November 2010
- Bay, H., Ess, A., Tuytelaars, T. and Van Gool, L., 2008. Speeded-up robust features (SURF). *Computer vision and image understanding*, 110(3), pp.346-359.
- Bayer, R. and McCreight, E., 2002. Organization and maintenance of large ordered indexes. In *Software pioneers* (pp. 245-262). *Book*. Springer, Berlin, Heidelberg.
- Bargoti, S., Mahajan, A., and Goktogan, A.H., 2012. Visual Odometry Based on the Fourier-Mellin Transform for an Experimental Mars Rover on Rough Terrain. *Proceedings from 12th Australian Space Science Conference*, 2012
- Benjemaa, R. and Schmitt, F., 1999. Fast global registration of 3D sampled surfaces using a multi-z-buffer technique. *Image and Vision Computing*, 17(2), pp.113-123.
- Bentley, J.L., 1975. Multidimensional binary search trees used for associative searching. *Communications of the ACM*, 18(9), pp.509-517
- Besl, P.J. and McKay, N.D., 1992. A method for registration of 3-D shapes. *IEEE Transactions on pattern analysis and machine intelligence*, 14(2), pp.239-256.

- Bisnath, S., Wang, J.-G., 2008. Advanced Satellite Positioning. *Lecture and scriptum*, York University, 2009
- Bisnath, S., and Y. Gao, 2009. Current state of precise point positioning and future prospects and limitations, *International Association of Geodesy Symposia*, Vol. 133, pp. 615-623.
- Blais, G. and Levine, M.D., 1995. Registering multi view range data to create 3D computer objects. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 17(8), pp.820-824.
- Boehm, J., 2002. Vorlesungsskript Nahbereichsphotogrammetrie. *Lecture notes* in close range photogrammetry. University of Stuttgart, Institute for Photogrammetry.
- Boeing, G. 2016. Visual Analysis of Nonlinear Dynamical Systems: Chaos, Fractals, Self-Similarity and the Limits of Prediction. *Systems*, 4 (4), 37. doi:10.3390/systems4040037
- Bombardier, A., 2013. Algorithmic of tie point implementation strategies. *Discussions* at Teledyne Optech, Toronto, 2013
- Bracewell, R.N., 1986. The Fourier transform and its applications (Vol. 31999). *Book*. New York: McGraw-Hill.
- Brooks, V.R. and Sitar, M. 2012. Accuracies of Orion and Gemini LiDAR systems. *Discussions* at Teledyne Optech, 2012
- Brown, L.G., 1992. A survey of image registration techniques. *ACM computing surveys (CSUR)*, 24(4), pp.325-376.
- Bszukalski, The ESRI Community Blog, 2017. Mostly Mapping Maps of the Day - Hurricanes Harvey and Irma. *Blog post* created by bszukalski-esristaff on Sep 12, 2017. <https://community.esri.com/blogs/bernszukalski/2017/09/13/mostly-mapping-maps-of-the-day-september-12-2017>. Accessed on Dec 10, 2017
- Budge, S.E., Badamkar, N.S. and Xie, X., 2015. Automatic registration of fused LiDAR/digital imagery (texel images) for three-dimensional image creation. *Optical Engineering*, 54(3), pp.031105-031105.
- CardinalSystems, 2018. <https://www.cardinalsystems1.net/>. Accessed June 2018
- Casasent, D. and Psaltis, D., 1976. Scale invariant optical transform. *Optical Engineering*, 15(3), p.153258.
- Canaz, S., 2012. Planar and linear feature-based registration of terrestrial laser scans with minimum overlap using photogrammetric data. *Dissertation*, University of Calgary, Department of Geomatics Engineering.
- Canny, J., 1986. A computational approach to edge detection. *IEEE Transactions on pattern analysis and machine intelligence*, (6), pp.679-698.

- Carpenter, J. and Hogarty, J., 2007. Multiple DEM Measured Accuracy. In *Proceedings of ASPRS 2007 Annual Conference*, Tampa, Florida.
- De Castro, E. and Morandi, C., 1987. Registration of translated and rotated images using finite Fourier transforms. *IEEE Transactions on pattern analysis and machine intelligence*, (5), pp.700-703.
- Censi, A., 2008, May. An ICP variant using a point-to-line metric. In *Robotics and Automation, 2008. ICRA 2008. IEEE International Conference on* (pp. 19-25). IEEE.
- Chen, Y. and Medioni, G., 1992. Object modelling by registration of multiple range images. *Image and vision computing*, 10(3), pp.145-155.
- Chen, Q.S., 1993. Image registration and its applications in medical imaging. *Dissertation*. Vrije Universiteit Brussel.
- Chen, Q.S., Defrise, M. and Deconinck, F., 1994. Symmetric phase-only matched filtering of Fourier-Mellin transforms for image registration and recognition. *IEEE Transactions on Pattern Analysis & Machine Intelligence*, (12), pp.1156-1168.
- Choi, K., Hong, J. and Lee, I., 2011. Precise geometric registration of aerial imagery and LiDAR data. *ETRI Journal*, 33(4), pp.506-516.
- CloudCompare, 2019. CloudCompare software v2, <https://www.danielgm.net/cc/> , accessed January 2019
- Congalton, R.G. and Green, K., 2008. Assessing the accuracy of remotely sensed data: principles and practices. *Article*. CRC press.
- Cooley, J.W. and Tukey, J.W., 1965. An algorithm for the machine calculation of complex Fourier series. *Mathematics of computation*, 19(90), pp.297-301.
- Corral-Soto, E.R. and Elder, J.H., 2014, September. Automatic single-view calibration and rectification from parallel planar curves. In *European Conference on Computer Vision* (pp. 813-827). Springer, Cham.
- Cramer, M., 1999. Introduction into photogrammetry. *Lecture*. University of Stuttgart, Institute for Photogrammetry.
- Cramer, M., 2001. Genauigkeitsuntersuchungen zur GPS/INS-Integration in der Aerophotogrammetrie. *Dissertation*. Deutsche Geodätische Kommission, Reihe C, Nr. 537, München 2001
- Cramer, M., Stallmann, D. and Haala, N., 2000. Direct georeferencing using GPS/inertial exterior orientations for photogrammetric applications. *International Archives of Photogrammetry and Remote Sensing*, 33(B3/1; PART 3), pp.198-205.

- Cramer, M. and Mostafa, M.A., 2004. A European network on digital camera calibration. *Photogrammetric Engineering & Remote Sensing*.
- Csanyi, N. and Toth, C.K., 2007. Improvement of LiDAR data accuracy using LiDAR-specific ground targets. *Photogrammetric Engineering & Remote Sensing*, 73(4), pp.385-396.
- D'Apuzzo, N., 2006. Overview of 3D surface digitization technologies in Europe. *Three-Dimensional Image Capture and Applications*, (6056).
- De Baar, J., 2020. Gradient-Enhanced Kriging. <https://www.mathworks.com/matlabcentral/file-exchange/60230-gradient-enhanced-kriging>, *MATLAB Central File Exchange*. Accessed April 2020
- De Castro, E. and Morandi, C., 1987. Registration of translated and rotated images using finite Fourier transforms. *IEEE Transactions on pattern analysis and machine intelligence*, (5), pp.700-703.
- Deng, F., Hu, M. and Guan, H., 2008. Automatic registration between LiDAR and digital images. *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, pp.487-490.
- Derpanis, K.G., 2010. Overview of the RANSAC Algorithm. *Image Rochester NY*, 4(1), pp.2-3.
- Derrode, S. and Ghorbel, F., 2001. Robust and efficient Fourier–Mellin transform approximations for gray-level image reconstruction and complete invariant description. *Computer vision and image understanding*, 83(1), pp.57-78.
- Dilley, M., 2005. Natural disaster hotspots: a global risk analysis (Vol. 5). *World Bank Publications*.
- Dos Santos, T.R., Seitel, A., Kilgus, T., Suwelack, S., Wekerle, A.L., Kenngott, H., Speidel, S., Schlemmer, H.P., Meinzer, H.P., Heimann, T. and Maier-Hein, L., 2014. Pose-independent surface matching for intra-operative soft-tissue marker-less registration. *Medical image analysis*, 18(7), pp.1101-1114.
- Douillard, B., Nourani-Vatani, N., Johnson-Roberson, M., Pizarro, O., Williams, S., Roman, C. and Vaughn, I., 2013. Frequency-based underwater terrain segmentation. *Autonomous Robots*, 35(4), pp.255-269.
- Duda, R.O. and Hart, P.E., 1972. Use of the Hough transformation to detect lines and curves in pictures. *Communications of the ACM*, 15(1), pp.11-15.
- Fan, S.K.S. and Chuang, Y.C., 2012. An entropy-based image registration method using image intensity difference on overlapped region. *Machine Vision and Applications*, 23(4), pp.791-804.
- Farrell, J., 2008. Aided navigation: GPS with high-rate sensors. *Book*. McGraw-Hill, Inc.
- Fernandez, J.C., Singhania, A., Caceres, J., Slatton, K.C., Starek, M. and Kumar, R., 2007. An

overview of LiDAR point cloud processing software. *GEM Center Report No. Rep_2007-12-001*, University of Florida.

Fischler, M.A. and Bolles, R.C., 1981. Random sample consensus: a paradigm for model fitting with applications to image analysis and automated cartography. *Communications of the ACM*, 24(6), pp.381-395.

Förstner, W. and Gülch, E., 1987, June. A fast operator for detection and precise location of distinct points, corners and centres of circular features. In *Proc. ISPRS intercommission conference on fast processing of photogrammetric data* (pp. 281-305).

Förstner, W., 2011. Probabilistic data analysis using graphical models. *Lecture scriptum*. Calgary, September 2011

Fourier, J., 1822. *Theorie Analytique de la Chaleur*. Firmin Didot, reissued by Cambridge University Press, 2009; ISBN 978-1-108-00180-9

Fonseca, L.M. and Manjunath, B.S., 1996. Registration techniques for multisensor remotely sensed imagery. *PE & RS- Photogrammetric Engineering & Remote Sensing*, 62(9), pp.1049-1056.

Friess, P., 2006. Towards a Rigorous Methodology for Airborne Laser Mapping, *Proceedings of the International Calibration and Validation Workshop EURO COW*, Castelldefels, Spain. Published as CD-ROM.

Friess, P., 2012. Use of 3D tie points for registration. *Discussions* at Teledyne Optech, 2011-2013

Friess, P., 2013. Brain storming, preparation of schematics, *discussions* about bundle adjustment calculus for integration of photogrammetric tie points with LiDAR plains. Teledyne Optech, 2013

Fritsch, D., 1999. Signalverarbeitung fuer Geodaeten. *Lecture*. University of Stuttgart, Institute for Photogrammetry.

Furukawa, Y., Curless, B., Seitz, S.M. and Szeliski, R., 2010. Towards internet-scale multi-view stereo. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2010 (pp. 1434-1441).

Gallant, J.C., and Austin, J.M., 2009. Stitching Fine Resolution DEMs, *18th World IMACS / MODSIM Congress, Cairns*, Australia 13-17 July 2009, <http://mssanz.org.au/modsim09>

Gao, Y., 2011. Coregistration of photogrammetric and lidar data in urban environments. *MscE Thesis*, York University 2011

Geiger, A., Lenz, P., Urtasun, R., 2012. Are we ready for autonomous driving? The Kitti Vision Benchmark Suite. *IEEE Conference on Computer Vision and Pattern Recognition (CVPR2012)*

Gelfand, N., Mitra, N.J., Guibas, L.J. and Pottmann, H., 2005, July. Robust global registration. In *Symposium on geometry processing* (Vol. 2, No. 3, p. 5).

GIP Aalen, 2018. *Email exchange* with Dr. Balazs Melykuti from GIP Aalen about comparison of ISAT and UltraMap generated tie points of the same flight project. May 2018

Glennie, C. L., 2007. Rigorous 3D error analysis of kinematic scanning LIDAR systems. *Journal of Applied Geodesy*, 1(3), pp. 147-157.

Gneeniss, A.S., Mills, J.P. and Miller, P.E., 2015. In-flight photogrammetric camera calibration and validation via complementary lidar. *ISPRS Journal of Photogrammetry and Remote Sensing*, 100, pp.3-13.

Gopaul, N., Wang, J.G., and Scherzinger, B., 2009. Simplified algorithms of variance component estimation for static and kinematic GPS single point positioning. *Journal of Global Positioning Systems*, 8(1), pp.43-52.

Grafarend, E.W., and Krumm, F., 2001. Adjustment Calculus 1&2&3 1999-2001. *Lecture series*. University of Stuttgart, Institute for Geodesy.

Grafarend, E.W., and Krumm, F. 2000. Map projections. *Lectures*. University of Stuttgart, Institute for Geodesy.

Grafarend, E.W., and Krumm, F. 2002. Physical Geodesy. *Lectures*. University of Stuttgart, Institute for Geodesy.

Grafarend, E.W., and Schwarz, F. 2001. Differential geometry. *Lectures*. University of Stuttgart, Institute for Photogrammetry.

Green, Jim, and Brooks, V.R., 2012. Accuracies of Orion and Gemini LiDAR systems. *Discussions at Teledyne Optech*, 2012

Grewal, M.S., Weill, L.R. and Andrews, A.P., 2007. Global positioning systems, inertial navigation, and integration. *Book*. John Wiley & Sons.

Gruber, M., 2017. *Conversations* about applications of Fourier Transform.

Gruber, M., Walcher, W., and Leberl, F., 2012. The Microsoft global ortho program. *ISPRS Ann Photogram Remote Sens Spat Inf Sci* I-4, pp.53-58.

Gruber, M. and Ladstätter, R., 2006, January. Geometric issues of the digital large format aerial camera UltraCamD. In *Proceedings of the International Calibration and Orientation Workshop EuroCOW* (pp. 25-27).

Gruber, M. and Walcher, W., 2014. Calibrating the new Ultracam osprey oblique aerial sensor. *The International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 40(3), p.47.

- Gruen, A., 1996. Least squares matching: A fundamental measurement algorithm, *Close Range Photogrammetry and Machine Vision* (K. Atkinson, editor), Whittles, pp. 217–255
- Gruen, A. and Akca, D., 2005. Least squares 3D surface and curve matching. *ISPRS Journal of Photogrammetry and Remote Sensing*, 59(3), pp.151-174.
- Guizar-Sicairos, M., Thurman, S.T. and Fienup, J.R., 2008. Efficient subpixel image registration algorithms. *Optics letters*, 33(2), pp.156-158.
- Haala, N., 2002. Digital Bilderverarbeitung 1&2. *Lecture series* in digital image processing by Prof. Dr.-Ing. Norbert Haala. University of Stuttgart, Institute for Photogrammetry, 2000-2002
- Haala, N., 2013, September. The landscape of dense image matching algorithms. In *Photogrammetric week* (Vol. 13, pp. 271-284).
- Habib, A. 2010. A Unified Conceptual Methodology for LiDAR and Camera System Calibration. University of Calgary, Digital Photogrammetry Research Group, *Presentation*
- Habib, A., Morgan, M. and Lee, Y.-R., 2002. Bundle adjustment with self-calibration using straight lines, *Photogrammetric Record*, 17(100), 2002, pp. 635-650.
- Habib, A., Ghanma, M., Michel, M., Al-Ruzouq, R., 2005. Photogrammetric and LiDAR data registration using linear features. *Photogramm. Eng. Remote Sens.* 71(6), pp. 699–707.
- Habib, A.F., Cheng, R.W.T., Kim, E.M., Mitishita, E.A., Frayne, R. and Ronsky, J.L., 2006. Automatic surface matching for the registration of LiDAR data and MR imagery. *ETRI journal*, 28(2), pp.162-174.
- Habib, A. and Rens, J., 2007. Quality assurance and quality control of LiDAR systems and derived data. In *Advanced LiDAR Workshop*, University of Northern Iowa.
- Habib, A. F., K. Bang, S.W. Shin, and E. Mitishita, 2007a: LiDAR system Self-calibration using Planar Patches from Photogrammetric Data. *The 5th International Symposium on Mobile Mapping Technology*, Padua, Italy, 28-31 May, Proceeding CD, 2007.
- Habib, A., Bang, K., Aldelgawy, M., and Shin, S., 2007b. Integration of Photogrammetric and LiDAR Data in a multi-primitive Triangulation Procedure, *Proceedings of the ASPRS 2007*, Tampa, Florida
- Habib, A., Bang, K. I., and Kersting, A. P., 2010a. LiDAR System Calibration: Impact on Plane Segmentation and photogrammetric Data Registration, *ASPRS 2010 Annual Conference*, San Diego, California, April 26 - May 30, 2010.
- Habib, A.F., Al-Durgham, M., Kersting, A.P. and Quackenbush, P., 2008, July. Error budget of LiDAR systems and quality control of the derived point cloud. In *Proceedings of the XXI ISPRS Congress, Commission I: Vol. XXXVII. Part B1* (pp. 203-209).
- Habib, A.F., 2013. *Hand out* about co-registration methods. Received in 2013

- Han, R.C. and Xiao, J.X., 2009. Deciding weighing by entropy value method is an error. In *Information and Computing Science, 2009. ICIC'09. Second International Conference* (Vol. 3, pp. 255-257). IEEE.
- Harris, C., and Stephens, M., 1988. A combined corner and edge detector. In: *Proc 4th Alvey Vis Conf. (AVC88)*. University of Manchester 1988. pp. 147-51.
- Hartley, R. and Zisserman, A., 2003. Multiple view geometry in computer vision. *Book. Cambridge university press*.
- Heipke, C., 1999. Automatic aerial triangulation: results of the OEEPE-ISPRS test and current developments. Wichmann, *Photogrammetric week*, pp.177-191.
- Heipke, C., Jacobsen, K. and Wegmann, H., 2002. Analysis of the results of the OEEPE test “Integrated Sensor Orientation. In *OEEPE Integrated Sensor Orientation Test Report and Workshop Proceedings*, Editors.
- Hexagon Intergraph, 2018. Intergraph Station Automatic Triangulation *Manual* (ISAT) 6.1
- Hexagon Leica, 2018. ADS100 and SH120 airborne sensor. <https://leica-geosystems.com/products/airborne-systems/imaging-sensors/leica-ads100-airborne-digital-sensor>. Accessed July 2018
- Hirschmuller, H., 2005, June. Accurate and efficient stereo processing by semi-global matching and mutual information. In *Computer Vision and Pattern Recognition, 2005. CVPR 2005. IEEE Computer Society Conference on* (Vol. 2, pp. 807-814). IEEE.
- Holmgren, J., Persson, Å. and Söderman, U., 2008. Species identification of individual trees by combining high resolution LiDAR data with multi-spectral images. *International Journal of Remote Sensing*, 29(5), pp. 1537–1552.
- Horn, B.K., 1987. Closed-form solution of absolute orientation using unit quaternions. *JOSA A*, 4(4), pp.629-642.
- Horner, J.L. and Gianino, P.D., 1984. Phase-only matched filtering. *Book. Applied optics*, 23(6), pp.812-816.
- Krystek, P., 1991. Fully automatic measurement of digital elevation models with MATCH-T. In *Proceedings of the 43rd Photogrammetric Week*, Stuttgart (pp. 203-214).
- Hu, L., He, Z., Liu, J. and Zheng, C., 2015. Method for Measuring the Information Content of Terrain from Digital Elevation Models. *Entropy*, 17(10), pp.7021-7051.
- Huang, J. and You, S., 2012. Point cloud matching based on 3D self-similarity. In *Computer Vision and Pattern Recognition Workshops (CVPRW)*, 2012 IEEE Computer Society Conference on (pp. 41-48). IEEE.

Huang, X., Mei, G., Zhang, J. and Abbas, R., 2021. A comprehensive survey on point cloud registration. *arXiv preprint arXiv:2103.02690*.

Hutton, J. and Mostafa, M.M., 2005. 10 years of direct georeferencing for airborne photogrammetry. *GIS Business (GeoBit)*, 11(1), pp.33-41.

IFP, 2017. Institute for Photogrammetry. Online link: http://www.ifp.uni-stuttgart.de/forschung/photo/test_site/. Accessed last in August 2017

IGI-Global (Idea Group Inc), 2017. Definition of the term “accuracy”. <https://www.igi-global.com/dictionary/geometric-accuracy/56382>. Accessed in December 2017

Ip, A.W.L., Mostafa, M.M.R. and El-Sheimy, N., 2004. Fast orthophoto production using the digital sensor system. In *Proceedings of the 7th Annual International Conference-Map India* (pp. 28-30).

Ismail, Z. and Jaafar, J., 2013, August. DEM derived from photogrammetric generated DSM using morphological filter. In *Control and System Graduate Research Colloquium (ICSGRC), 2013 IEEE 4th* (pp. 103-106).

Itti, L., Koch, C. and Niebur, E., 1998. A model of saliency-based visual attention for rapid scene analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 20(11), pp.1254-1259.

Jacobsen, K. 2003. System Calibration for Direct and Integrated Sensor Orientation, *ISPRS Working Group Working Group I/5 of the ISPRS Commision I*, Castlefels, Spain

Jähne, B., 2002. Digitale Bildverarbeitung. *Book*. Springer-Verlag.

Jung J., C. Armenakis, G. Sohn. 2010. Data driven alignment of 3D building models and digital aerial images. CD-ROM Proceedings *ISPRS Com VII Symposium*, Vienna.

Kager, H., 2004. Discrepancies between Overlapping Laser Scanning Strips- Simultaneous Fitting of Aerial Laser Scanner Strips. In: *The International Archives of the Photogrammetry, Remote Sensing and Spatial Information Science, Istanbul, Turkey, Vol. XXXV, Part B/1*, pp. 555 - 560.

Kahmen, H., 2006. Angewandte Geodaesie: Vermessungskunde, 20. *Book*. Voellig neu ueberarbeitete Auflage, de Gruyter, 700 pages.

Kashiwagi, N., 1993. On use of the Kalman filter for spatial smoothing. *Annals of the Institute of Statistical Mathematics*, 45(1), pp.21-34.

Kazik, T. and Göktoğan, A.H., 2011, April. Visual odometry based on the Fourier-Mellin transform for a rover using a monocular ground-facing camera. In *Mechatronics (ICM), 2011 IEEE International Conference on* (pp. 469-474). IEEE.

Keller, W., 2001. Satellite Geodesy 1&2. *Lecture*. University of Stuttgart, Institute for Geodesy.

- Keiner, J., Kunis, S. and Potts, D., 2007. Efficient reconstruction of functions on the sphere from scattered data. *Journal of Fourier Analysis and Applications*, 13(4), pp.435-458.
- Kekre, H.B. and Gharge, S.M., 2010. Image segmentation using extended edge operator for mammographic images. *International journal on computer science and Engineering*, 2(4), pp.1086-1091.
- Kersting, A.P.B., 2012. Quality Assurance of Multi-Sensor Systems. University of Calgary. Department of Geomatics, *Dissertation*. (URL: <http://www.geomatics.ucalgary.ca/graduatetheses>). December 2011
- Kim, H.Y. and De Araújo, S.A., 2007, December. Grayscale template-matching invariant to rotation, scale, translation, brightness and contrast. In *Pacific-Rim Symposium on Image and Video Technology* (pp. 100-113). Springer, Berlin, Heidelberg.
- King, A.D., 1998. Inertial navigation-forty years of evolution. *GEC review*, 13(3), pp.140-149.
- Kleusberg, A. and Schoeller, W. 2001-2002. Navigation 1&2. *Lecture*. University of Stuttgart, Institute for Navigation.
- Kok, M., Hol, J.D. and Schön, T.B., 2017. Using inertial sensors for position and orientation estimation. *arXiv preprint arXiv:1704.06053*.
- Khoshelham, K., 2009. Role of tie points in integrated sensor orientation for photogrammetric map compilation. *Photogrammetric Engineering & Remote Sensing*, 75(3), pp.305-311.
- Kacker, R.N., Lagergren, E.S. and Filliben, J.J., 1991. Taguchi's orthogonal arrays are classical designs of experiments. *Journal of research of the National Institute of Standards and Technology*, 96(5), p.577.
- Kraus, K., 2007. Photogrammetry: geometry from images and laser scans. *Book*. Walter de Gruyter.
- Kraus, K., 1993. Photogrammetry. Vol. 2. *Book*. Dümmler.
- Kruck, E., 1983. Lösung großer Gleichungssysteme für photogrammetrische Blockausgleichungen mit erweitertem funktionalem Modell. *Dissertation*. Geodät. Inst., Univ. Hannover
- Kruck, E., 2006. Combined IMU Sensor Calibration and Bundle Adjustment with BINGO-F, *OEEPE Official Publication No. 43*, "Integrated Sensor Orientation, Test Report and Workshop Proceedings".
- Kruck, E., Melikuti, B., 2013. *Discussions about the ATG module in Bingo and its capabilities of tying image blocks of multiple altitudes*
- Kuhlmann, H., 2002. Engineering Geodesy 2. *Lecture*. University of Stuttgart, Institute for

Applied Geodesy in Civil Engineering.

Kunis, S., and Potts, D., 2006. NFFT, software package, C subroutine library. <http://www.tu-chemnitz.de/~potts/nfft>, 2002 – 2006.

Kunis, S. 2006. Nonequispaced FFT. Generalization and Inversion. *Dissertation*. University of Lübeck, Germany

Kwak, E. and Habib, A., 2014. Automatic representation and reconstruction of DBM from LiDAR data using Recursive Minimum Bounding Rectangle. *ISPRS Journal of Photogrammetry and Remote Sensing*, 93, pp.171-191.

Kwak, E., 2013. Automatic 3D building model generation by integrating LiDAR and aerial images using a hybrid approach. *Dissertation*, University of Calgary.

Latulippe, M., Drouin, A., Giguere, P. and Laviolette, F., 2013, August. Accelerated Robust Point Cloud Registration in Natural Environments through Positive and Unlabeled Learning. In *IJCAI* (pp. 2480-2487).

Leberl, F., Irschara, A., Pock, T., Meixner, P., Gruber, M., Scholz, S. and Wiechert, A., 2010. Point clouds. *Photogrammetric Engineering & Remote Sensing*, 76(10), pp.1123-1134.

Lee, D.H., Lee, K.M., Lee, S.U., 2008. Fusion of LiDAR and imagery for reliable building extraction. *Photogrammetric Engineering and Remote Sensing* 74, 215-225.

Lee, W., Sim, D. and Oh, S.J., 2021. A CNN-Based High-Accuracy Registration for Remote Sensing Images. *Remote Sensing*, 13(8), p.1482.

Legat, K., Skaloud, J. and Schmidt, R., 2006. Reliability of Direct Georeferencing: A Case Study on Practical Problems and Solutions, Final Report on Phase 2. In *EuroSDR Official Publication No 51*, pp. 169–184.

Le Maistre Pratt, J.G.L.M., 2000. Application of the Fourier-Mellin transform to translation-, rotation-and scale-invariant plant leaf identification. *Aspen Bibliography. Paper 773*. http://digitalcommons.usu.edu/aspen_bib/773

Le Moigne, J., Netanyahu, N.S. and Eastman, R.D. eds., 2011. Image registration for remote sensing. Cambridge University Press.

Li, Z., Zhu, C. and Gold, C., 2004. Digital terrain modeling: principles and methodology. Book. CRC press.

Li, H., Zhong, C. and Huang, X., 2012. Reliable registration of LiDAR data and aerial images without orientation parameters. *Sensor Review*, 32(4), pp.280-287.

Li, C.C., Zhang, G.S., Lei, T.J. and Gong, A.D., 2011. Quick image-processing method of UAV

without control points data in earthquake disaster area. *Transactions of Nonferrous Metals Society of China*, 21, pp. s523-s528.

Lindenthal, S., 2003a. Automatische Texturextraktion in ArcGIS. *Erste Studienarbeit/First Practical*, University of Stuttgart, Institute for Photogrammetry, 2003

Lindenthal, S., 2003b. Automatische Texturextraktion und Bildentzerrung mit OpenGL. *Diplomarbeit/Master's Thesis*, University of Stuttgart, Institute for Photogrammetry, 2004

Lindenthal, S., 2004. Automatische Linienextraktion für ein Lane Departure Warning System / Lane Extraction and Tracking for a Lane Departure Warning System. *Zweite Studienarbeit/Second Practical*, Robert Bosch GmbH, Leonberg, department of Driver Assistance in Automobile Electronics, rear camera for Audi A8, in collaboration with University of Stuttgart, Institute for Photogrammetry, 2004

Lindenthal, S., 2005. Self-calibration of security cameras from by-passing objects. *Internship at Robert Bosch GmbH, Bosch-Blaupunkt, Hildesheim, department of CR/AEM5, Corporate Research/Automobile Electronics and Multimedia*, 2005

Lindenthal, S.M., Wang, J.G., Bisnath, S., and Canter, P., 2009. Performance analysis of GNSS network extrapolation supporting the moving reference station concept. *York University internal report*, generated in cooperation with Trimble Applanix, 2009

Lindenthal, S., Ussyshkin, V., Wang, J. and Pokorný, M., 2011. Airborne LIDAR: A fully automated self-calibration procedure. *ISPRS Achieves, ISPRS Calgary 2011 Workshop*, pp.29-31.

Lindenthal, S.M., Armenakis, C., Wang, J.G., Bisnath, S. and Goldman, H., 2015. Sensor-Driven Calibration and Co-Registration of Airborne Multisensor LiDAR Data. *Canadian Journal of Remote Sensing*, 41(3), pp.145-158.

Lingua, A., Marenchino, D. and Nex, F., 2009. Performance analysis of the SIFT operator for automatic feature extraction and matching in photogrammetric applications. *Sensors*, 9(5), pp.3745-3766.

Liu, W., Sun, J., Li, W., Hu, T. and Wang, P., 2019. Deep learning on point clouds and its application: A survey. *Sensors*, 19(19), p.4188.

Lowe, D.G., 2004. Distinctive image features from scale-invariant key points. *International journal of computer vision*, 60(2), pp.91-110.

Lu, W., Wan, G., Zhou, Y., Fu, X., Yuan, P. and Song, S., 2019. Deepvcv: An end-to-end deep neural network for point cloud registration. In *Proceedings of the IEEE/CVF International Conference on Computer Vision* (pp. 12-21).

Luhmann, T., Robson, S., Kyle, S. and Hartley, I. 2007. Close Range Photogrammetry. Principles,

Methods and Applications. *Book*. Whittles Publishing, 2007

Lyons, R. G., 2010, Understanding Digital Signal Processing, 3rd Edition, *Digital book*, published Nov 1, 2010, Prentice Hall.

Maas, H.G., 2003, September. Planimetric and height accuracy of airborne laser scanner data: User requirements and system performance. In *Proceedings* (Vol. 49, pp. 117-125).

Maes, F., Collignon, A., Vandermeulen, D., Marchal, G. and Suetens, P., 1997. Multimodality image registration by maximization of mutual information. *IEEE transactions on medical imaging*, 16(2), pp.187-198.

Maiseli, B., Gu, Y. and Gao, H., 2017. Recent developments and trends in point set registration methods. *Journal of Visual Communication and Image Representation*, 46, pp.95-106.

Makadia, A., Patterson, A. and Daniilidis, K., 2006, June. Fully automatic registration of 3D point clouds. In *Computer Vision and Pattern Recognition, 2006. IEEE Computer Society Conference on* (Vol. 1, pp. 1297-1304). IEEE.

Mandelbrot, B.B. and Pignoni, R., 1983. The fractal geometry of nature. *Macmillan*. ISBN 978-0-7167-1186-5.

Mastin, A., Kepner, J. and Fisher, J., 2009, June. Automatic registration of LIDAR and optical images of urban scenes. In *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on* (pp. 2639-2646). IEEE.

Masuda, T., Sakaue, K. and Yokoya, N., 1996, August. Registration and integration of multiple range images for 3-D model construction. In *Pattern Recognition, 1996. Proceedings of the 13th International Conference* (Vol. 1, pp. 879-883). IEEE.

Mattes. 2020. Mutual Information Based Registration of Medical Images. Washington State University, Course CSE577, Z5, *Course Notes*, Accessed April 2020

Maune, D. 2001. Digital Elevation model Technologies and Applications: The DEM Users Manual, American Society for Photogrammetry and Remote Sensing, Bethesda, MA, USA
MathWorks, 2018. https://www.mathworks.com/products.html?s_tid=gn_ps

McGuire, M., 1998. An image registration technique for recovering rotation, scale, and translation parameters. *NEC Res. Inst. Tech. Rep.*, TR, pp.98-018.

Meer, P. and Georgescu, B., 2001. Edge detection with embedded confidence. *IEEE Transactions on pattern analysis and machine intelligence*, 23(12), pp.1351-1365.

Merriam-Webster, 2017. Merriam-Webster's collegiate dictionary. Merriam-Webster.8 <https://www.merriam-webster.com/dictionary/salient>. Accessed last December 2017

- Mellado, N., Dellepiane, M. and Scopigno, R., 2015. Relative scale estimation and 3D registration of multi-modal geometry using Growing Least Squares. *IEEE transactions on visualization and computer graphics*, 22(9), pp.2160-2173.
- Mishra, R. K., and Zhang, Y., 2012. A review of optical imagery and airborne LiDAR data registration methods. *The Open Remote Sensing Journal*, 5(1).
- Mittal, A. and Sharma, M., 2008. Number Plate Recognition Using Analytical Fourier Mellin Transform. In *Novel Algorithms and Techniques in Telecommunications, Automation and Industrial Electronics* (pp. 37-42). Springer, Dordrecht.
- Moehlenbrink, W. and Glaeser, A., 2000. Engineering Geodesy 1. *Lecture*. University of Stuttgart, Institute for Applied Geodesy in Civil Engineering.
- Morgan, D. and Falkner, E., 2001. Aerial mapping: methods and applications. Second Edition. *Book*. CRC Press.
- Mostafa, M.M., 2001, June. Boresight calibration of integrated inertial/camera systems. In *Proc. International Symposium on Kinematic Systems in Geodesy, Geomatics and Navigation-KIS* (pp. 440-445).
- Moussa, W., 2014. Integration of digital photogrammetry and terrestrial laser scanning for cultural heritage data recording. *Dissertation*. University of Stuttgart, Germany
- Muick, M., 2018. *Conversation* about ideal number of tie points at Vexcel-Imaging GmbH, 2018
- National Geospatial-Intelligence Agency, NGA-mil, 2017. Accuracy. *Online publication*. <http://earth-info.nga.mil/GandG/publications/accuracy.pdf>. Accessed last 20th December 2017
- Niemeier, W., 2008. Ausgleichungsrechnung: Statistische Auswertemethoden. *Book*. Walter de Gruyter.
- Nittmann, J. and Stanley, H.E., 1987. Non-deterministic approach to anisotropic growth patterns with continuously tunable morphology: the fractal properties of some real snowflakes. *Journal of Physics A: Mathematical and General*, 20(17), p. L1185.
- Noble A., 1989. Descriptions of image surfaces. *Dissertation*, Oxford, UK, Oxford University.
- Nguyen et al., 2020. Coarse-to-fine registration of airborne lidar data and optical imagery on urban scenes. *IEEE Journal of Selected topics in applied earth observations and remote sensing*. Vol. 13, 2020
- Oppenheim, A.V. and Lim, J.S., 1981. The importance of phase in signals. *Proceedings of the IEEE*, 69(5), pp.529-541.
- Osten, W. 2006. Technische Optik 1&2, 2006. *Lectures in technical optics* by Prof. Dr. Wolfgang

Osten, University of Stuttgart, Institute of Technical Optics, 2005-2006

Osten, W. 2005. Optische Informationsverarbeitung, 2005. *Lecture in optical information processing* by Prof. Dr. Wolfgang Osten, University of Stuttgart, Institute of Technical Optics, 2005

Palenichka, R.M. and Zaremba, M.B., 2005, September. Object shape extraction based on the piecewise linear skeletal representation. In *International Conference Image Analysis and Recognition* (pp. 464-472). Springer, Berlin, Heidelberg.

Palenichka, R.M. and Zaremba, M.B., 2010. Automatic extraction of control points for the registration of optical satellite and LiDAR images. *IEEE Transactions on Geoscience and Remote sensing*, 48(7), pp.2864-2879.

Parmehr, E.G., Fraser, C.S., Zhang, C. and Leach, J., 2012, December. Automatic Registration of Aerial Images with 3D LiDAR Data Using a Hybrid Intensity-Based Method. In *Digital Image Computing Techniques and Applications (DICTA), 2012 International Conference on* (pp. 1-7). IEEE

Parmehr, E.G., Fraser, C.S., Zhang, C. and Leach, J., 2014. Automatic registration of optical imagery with 3D lidar data using statistical similarity, *ISPRS Journal of Photogrammetry and Remote Sensing*, vol. 88, pp. 28-40, 2014

Peckham, R.J. and Jordan, G., 2007. Digital terrain modelling. *Book*. Springer-Verlag Berlin Heidelberg.

Peng, F., Wu, Q., Fan, L., Zhang, J., You, Y., Lu, J. and Yang, J.Y., 2014, October. Street view cross-sourced point cloud matching and registration. In *2014 IEEE International Conference on Image Processing (ICIP)* (pp. 2026-2030). IEEE.

Peng, Z. and Li, C., 2019. Portable microwave radar systems for short-range localization and life tracking: A review. *Sensors*, 19(5), p.1136.

PennState University, 2016: GEOG481 – Topographic Mapping with LiDAR. *Online lecture notes*, https://www.e-education.psu.edu/geog481/11_p6.html. Accessed last on October 30th 2016.

Perez, P., 1998. Markov random fields and images (p. 31). *IRISA*.

Persad, R.A., and Armenakis, C., 2017. Automatic co-registration of 3D multi-sensor point clouds. *ISPRS Journal of Photogrammetry and Remote Sensing*, 130, pp.162-186.

Persad, Ravi Ancil, 2017. Automatic Alignment of 3D Multi-Sensor Point Clouds. *Dissertation*. York University, Toronto, Canada, 2017

Persson, Å.J.U.H., Holmgren, J., Söderman, U. and Olsson, H., 2004. Tree species classification of individual trees in Sweden by combining high resolution laser data with high resolution near-infrared digital images. *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 36(8), pp.204-207.

PhaseONE, 2018. CaptureOne software description. <https://www.phaseone.com/en/Products/Software/Capture-One-Pro/Highlights.aspx>. Accessed July 2018

Pluim, J.P., Maintz, J.A. and Viergever, M.A., 2003. Mutual-information-based registration of medical images: a survey. *IEEE transactions on medical imaging*, 22(8), pp.986-1004.

Pokorny, M., 2013. Horizontal accuracy assessment of LiDAR point clouds. *Discussion with Martin Pokorny, Lead LiDAR QA/QC, Teledyne Optech*, 2013

Puente, I., González-Jorge, H., Arias, P. and Armesto, J., 2011. Land-based mobile laser scanning systems: a review. *ISPRS Commission V, WG V/3*

Pulli, K. and Shapiro, L.G., 2000. Surface reconstruction and display from range and color data. *Graphical Models*, 62(3), pp.165-201.

Quantopian, 2018. Quantopian lecture series on Kalman Filtering. <https://www.quantopian.com/posts/kalman-filters-best-practices>. Accessed July 2018

Rapidlasso, 2018. <https://rapidlasso.com/category/dtm/>. Accessed September 2018

Ray, J.A. and Graham, L., 2008. New Horizontal Accuracy Assessment Tools and Techniques for LiDAR Data. In *Proceedings of the ASPRS Annual Conference*, Portland/Oregon, USA.

Reddy, B.S. and Chatterji, B.N., 1996. An FFT-based technique for translation, rotation, and scale-invariant image registration. *IEEE transactions on image processing*, 5(8), pp.1266-1271.

Renaudin, E., Habib, A. and Kersting, A.P., 2011. Feature-based registration of terrestrial laser scans with minimum overlap using photogrammetric data. *Etri Journal*, 33(4), pp.517-527.

Reynolds, D.A., 2009. Gaussian mixture models. *Encyclopedia of biometrics*, 741, pp.659-663.

Robert Bosch GmbH, 2008. Bosch Precision Wheel Alignment System. Robert Bosch GmbH, Schwieberdigen, 2008. *Video*. <https://www.youtube.com/watch?v=O0wcfOZpbzA>. Last accessed on Dec 20, 2017

Rönnholm, P. and Haggrén, H., 2012. Registration of laser scanning point clouds and aerial images using either artificial or natural tie features. *ISPRS Ann. Photogramm. Remote Sens. Spatial Inf. Sci.*, I-3, pp. 63-68, doi:10.5194/isprsannals-I-3-63-2012

Roscher, R., Waske, B. and Förstner, W., 2010, August. Kernel discriminative random fields for land cover classification. In *Pattern Recognition in Remote Sensing (PRRS)*, 2010 IAPR Workshop

on (pp. 1-5). IEEE.

Rothermel, M., Wenzel, K., Fritsch, D. and Haala, N., 2012, December. SURE: Photogrammetric surface reconstruction from imagery. In *Proceedings LC3D Workshop*, Berlin (Vol. 8).

Rublee, E., Rabaud, V., Konolige, K. and Bradski, G., 2011, November. ORB: An efficient alternative to SIFT or SURF. In *Computer Vision (ICCV), 2011 IEEE international conference on* (pp. 2564-2571). IEEE.

Rusinkiewicz, S. and Levoy, M., 2001. Efficient variants of the ICP algorithm. In *3-D Digital Imaging and Modeling, 2001. Proceedings. Third International Conference on* (pp. 145-152). IEEE.

Sabuncu, M.R., 2006. Entropy-based image registration. *Dissertation*. Princeton University.

Sakuraba, M., 2013. Course in Programming. Research Institute of Electrical Communication, Tohoku University. http://www5a.biglobe.ne.jp/~tenrou/programming/SnowCrystal1_WPF.htm. Accessed last January 2018

Schaer, P., Skaloud, J., Landtwig, S. and Legat, K., 2007. Accuracy estimation for laser point cloud including scanning geometry. *International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 36 (Part 5/C55), pp. 279-353

Shannon, C. and Weaver, W., 1949. The Mathematical Theory of Information. *Urbana*: University of Illinois Press.

Sharma, A., Tiwari, K.N., and Bhadoria, P.B.S., 2011. Determining the optimum cell size of digital elevation model for hydrologic application. *Journal of earth system science*, 120(4), pp.573-582.

Shorter, N. and Kasparis, T., 2008. Autonomous registration of LiDAR data to single aerial image. In *Geoscience and Remote Sensing Symposium, 2008. IGARSS 2008. IEEE International* (Vol. 5, pp. V-216). IEEE.

Scherzinger, B.M., 2000, September. Precise robust positioning with inertial/GPS RTK. In *Proceedings of the 13th International Technical Meeting for the Satellite Division of the Institute of Navigation (ION GPS)* (pp. 115-162).

Schneider, D., 2009. Geometrische und stochastische Modelle für die integrierte Auswertung terrestrischer Laserscannerdaten und photogrammetrischer Bilddaten. *Dissertation*. Verlag der Bayerischen Akademie der Wissenschaften.

Schwanghart, W., 2020. Ordinary Kriging. <https://www.mathworks.com/matlabcentral/fileexchange/29025-ordinary-kriging>. MATLAB Central File Exchange. Accessed April 2020.

Schwieger, V., 2003. Engineering Geodesy 3. *Lecture*. University of Stuttgart, Institute for

Applied Geodesy in Civil Engineering.

Sebe, N. and Lew, M.S., 2003. Comparing salient point detectors. *Pattern recognition letters*, 24(1), pp.89-96.

Sellami, M. and Ghorbel, F., 2012, August. An invariant similarity registration algorithm based on the analytical Fourier Mellin transform. In *Signal Processing Conference (EUSIPCO), 2012 Proceedings of the 20th European* (pp. 390-394). IEEE.

Sester, M., Hild, H. and Fritsch, D., 1998. Definition of ground-control features for image registration using GIS data. *International Archives of Photogrammetry and Remote Sensing*, 32, pp.538-543.

Shan, J. and Toth, C.K., 2010. Topographic laser ranging and scanning: principles and processing. *Book*. CRC press.

Shragai, Z., Barnea, S., Filin, S., Zalmanson, G. and Doytsher, Y., 2005, October. Automatic image sequence registration based on a linear solution and scale invariant keypoint matching. In *BenCOS-ISPRS Workshop in conjunction with ICCV* (pp. 5-11).

Shragai, Z., Even-Paz, A. and Klein, I., 2011, May. Automatic tie-point extraction using advanced approaches. In *Proceedings of the ASPRS 2011 Annual Conference*, Wisconsin.

Simon, D.A., 1996. Fast and accurate shape-based registration. *Dissertation*. Pittsburgh, Pennsylvania: Carnegie Mellon University.

Sitar, M. 2011. Internal Presentation on New Features in Optech LiDAR Mapping Suite LMS. *Optech Conference for Laser Scanning*, 2011.

Skaloud, J., 2011. Integrated Sensor Orientation. Introductory Text for ENGO699.55. *Lecture scriptum*, University of Calgary, July 2011

Soininen, A. and Burman, H., 2005. TerraMatch for MicroStation. *Terrasolid Ltd*, Finland.

Soininen, A. 2010a. *Presentation* on TerraPhoto software, including flight design requirements

Soininen, A., 2010b. New features in TerraScan. *Presentation at TerraScan User Event, Levi, Finland*. Available online: http://www.terrasolid.fi/en/presentations/latest_news_of_terrasolid_software/terrascan_new_features_03_09_02_10_arttu_soininen_levi_2010. Last date accessed 5 July 2012.

Spooren, N., Geelen, B., Tack, K., Lambrechts, A., Jayapala, M. Ginat, R. David, Y. Levi, E., and Grauer, Y., 2017. RGB-NIR Active Gated Imaging. *Electro-Optical and Infrared Systems: Technology and Applications XIII, SPIE*

Stone, H.S., Orchard, M.T., Chang, E.C. and Martucci, S.A., 2001. A fast, direct Fourier-based

algorithm for subpixel registration of images. *IEEE Transactions on geoscience and remote sensing*, 39(10), pp.2235-2243

Studholme, C., Hill, D.L. and Hawkes, D.J., 1999. An overlap invariant entropy measure of 3D medical image alignment. *Pattern recognition*, 32(1), pp.71-86.

Swoboda, D.M., 2014. A comprehensive characterization of the Asus Xtion pro depth sensor. *J. Sci. Technol*, 1, pp.16-20.

Tarsha-Kurdi, F., Landes, T. and Grussenmeyer, P., 2007, September. Hough-transform and extended ransac algorithms for automatic detection of 3d building roof planes from lidar data. In *ISPRS Workshop on Laser Scanning 2007 and SilviLaser 2007* (Vol. 36, pp. 407-412).

Teledyne Optech, 2012. Communication with Jeff Tripp, Principal Engineer at Teledyne Optech.

Teledyne Optech, 2013. <http://www.teledyneoptech.com/>. Accessed June 2018

Tempelmann, U., Börner, A., Chaplin, B., Hinsken, L., Mykhalevych, B., Miller, S., Recke, U., Reulke, R. and Uebbing, R., 2000. Photogrammetric software for the LH Systems ADS40 airborne digital sensor. *International Archives of Photogrammetry and Remote Sensing*, 33(B2; PART 2), pp.552-559.

Tempelmann, U., and Hinsken, L. 2003. ADS40 calibration and verification process. *ISPRS 2003*

Theiler, P.W. and Schindler, K., 2012. Automatic registration of terrestrial laser scanner point clouds using natural planar surfaces. *ISPRS Annals of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 3, pp.173-178.

Toth, C., Paska, E. and Brzezinska, D., 2008. Using road pavement markings as ground control for LiDAR data. *International Archives of Photogrammetry, Remote Sensing and Spatial Information Sciences*, 37, pp.189-195.

Trimble Applanix, 2017. POSAV Specifications. Online source. https://www.applanix.com/downloads/products/specs/posav_datasheet.pdf. Accessed Dec 2017

UCLA, 2017. Math lecture. www.math.ucla.edu/~cratsch/151b.2.05s/Lectures/Fourier/fourier.html. Accessed last January 2018

Ussyshkin, R.V. and Smith, B., 2006, May. Performance analysis of ALTM 3100EA: Instrument specifications and accuracy of LiDAR data. In *ISPRS Commission I Symposium, Paris; Proceedings, part B*.

Ussyshkin, R.V., Boba, M., Smith, B., Lane, T., Hlaing, D., 2008. Performance characterization of an airborne LiDAR system: Instrument specifications and attainable field accuracy. *ILMF Conference Proceeding, Denver, Colorado, U.S.A. 21-22 Feb. 2008*. Published as CD-ROM.

Vexcel Imaging GmbH, 2018. UltraCam Condor. <http://www.vexcel-imaging.com/ultracam-condor/>. Accessed July 2018.

Vieux, B.E., 1993. DEM aggregation and smoothing effects on surface runoff modeling. *Journal of Computing in Civil Engineering*, 7(3), pp.310-338.

Viola, P. and Wells III, W.M., 1997. Alignment by maximization of mutual information. *International journal of computer vision*, 24(2), pp.137-154.

Vollrath, A. 2010. The Nonequispaced Fast SO(3) Fourier Transform, Generalisations and Applications. *Dissertation*. University of Lübeck, Germany

Vosselman, G., 2000. Slope based filtering of laser altimetry data. *International Archives of Photogrammetry and Remote Sensing*, 33(B3/2; PART 3), pp.935-942.

Vosselman, G., 2008. Analysis of planimetric accuracy of airborne laser scanning surveys. *Bulletin N 22*, p.27.

Vosselman, G., Maas, H.G., 2010. Airborne and Terrestrial Laser Scanning, *Hardcover Book*, CRC Press, ISBN-13: 978-1439827987, March 5, 2010. Whittles Publishing. ISBN 978-1904445-87-6

Walter, V., 2003. GIS 2&3. *Lecture series in Geoinformation Systems* by Dr.-Ing. Volker Walter. University of Stuttgart, Institute of Photogrammetry, 2001-2003

Wang, J.G., 2008. Direct Readings. *Lecture*. York University, 2008

Weik, S., 1997, May. Registration of 3-D partial surface models using luminance and depth information. In *3-D Digital Imaging and Modeling, 1997*. Proceedings, International Conference on Recent Advances in (pp. 93-100). IEEE.

Wen, J., 2016. Kriging Toolbox. <https://www.mathworks.com/matlabcentral/fileexchange/59960-krigingtoolbox>, MATLAB Central File Exchange. Accessed April 2020

Wiechert, A., Gruber, M. and Karner, K., 2012. Ultramap: The all-in-one photogrammetric solution. *ISPRS-International Archives of the Photogrammetry, Remote Sensing and Spatial Information Sciences*, 39, p. B3.

Wilson, J.P. and Gallant, J.C., 2000. Digital terrain analyses : Principles and applications, *book*, 6(12), pp.1-27.

Wise, S., 2012. Information entropy as a measure of DEM quality. *Computers & geosciences*, 48, pp.102-110.

Wisetphanichkij, S. and Dejhan, K., 2005. Fast Fourier transform technique and affine transform estimation-based high precision image registration method. *GESTS Int'l Trans. Computer Science*

and *Engr*, 20(1), pp.179-191.

Wodzinski, M. and Skalski, A., 2021. Multistep, automatic and nonrigid image registration method for histology samples acquired using multiple stains. *Physics in Medicine & Biology*, 66(2), p.025006.

Wolberg, G. and Zokai, S., 2000. Robust image registration using log-polar transform. In *Image Processing, 2000. Proceedings. 2000 International Conference on* (Vol. 1, pp. 493-496). IEEE.

Wong A, and Orchard J., 2008. Efficient FFT-Accelerated Approach to Invariant Optical–LIDAR Registration. *IEEE Geoscience and Remote Sensing 2008*; 46(11): 3917-25

Wu, C.C., 2018. VisualSFM: A Visual Structure from Motion System. <http://ccwu.me/vsfm/>. Accessed June 2018

Wu, H., Li, Y., Li, J. and Gong, J., 2010. A two-step displacement correction algorithm for registration of lidar point clouds and aerial images without orientation parameters. *Photogrammetric Engineering & Remote Sensing*, 76(10), pp.1135-1145.

Wu, Z., Song, S., Khosla, A., Yu, F., Zhang, L., Tang, X., Xiao, J., 2015. 3d Shapenets: A Deep Representation for Volumetric Shapes. *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (VPR 2015)*, Boston, MA, USA, 7-1 June 2015

Wyawahare, M.V., Patil, P.M. and Abhyankar, H.K., 2009. Image registration techniques: an overview. *International Journal of Signal Processing, Image Processing and Pattern Recognition*, 2(3), pp.11-28.

Yanalak, M., 2003. Effect of gridding method on digital terrain model profile data based on scattered data. *Journal of Computing in Civil Engineering*, 17(1), pp.58-67.

Yilmaz, H.M., 2007. The effect of interpolation methods in surface definition: an experimental study. *Earth Surface Processes and Landforms: The Journal of the British Geomorphological Research Group*, 32(9), pp.1346-1361.

Yilmaz, H.M., Yakar, M., Mutluoglu, O. and Yildiz, F., 2004, July. Selection of the most suitable sizes of ground control points in the satellite images. In *Commission 4. Proceedings of the 20th ISPRS Congress* (pp. 738-742).

Yuan, X., Fu, J., Sun, H. and Toth, C., 2009. The application of GPS precise point positioning technology in aerial triangulation. *ISPRS Journal of Photogrammetry and Remote Sensing*, 64(6), pp.541-550.


Zeng, A., Song, S., Nießner, M., Fisher, M., Xiao, J. and Funkhouser, T., 2017. 3dmatch: Learning local geometric descriptors from rgb-d reconstructions. *IEEE conference on computer vision and pattern recognition* (pp. 1802-1811).

- Zhang, G., Lei, M. and Liu, X., 2009. Novel template matching method with sub-pixel accuracy based on correlation and Fourier-Mellin transform. *Optical Engineering*, 48(5), pp.057001-057001.
- Zhang, Z., Dai, Y. and Sun, J., 2020. Deep learning-based point cloud registration: an overview. *Virtual Reality & Intelligent Hardware*, 2(3), pp.222-246.
- Zheng, S., Huang, R. and Zhou, Y., 2013. Registration of optical images with LiDAR data and its accuracy assessment. *Photogrammetric Engineering & Remote Sensing*, 79(8), pp.731-741.
- Zitova, B. and Flusser, J., 2003. Image registration methods: a survey. *Image and vision computing*, 21(11), pp.977-1000.
- Zokai, S. and Wolberg, G., 2005. Image registration using log-polar mappings for recovery of large-scale similarity and projective transformations. *IEEE Transactions on Image Processing*, 14(10), pp.1422-1434.
- Zhu, H., Guo, B., Zou, K., Li, Y., Yuen, K.V., Mihaylova, L. and Leung, H., 2019. A review of point set registration: From pairwise registration to groupwise registration. *Sensors*, 19(5), p.1191.

APPENDIX A: Related AI Based Co-Registration Methods

Historically, machine learning started out with the registration of imagery from the same sensor about 15 years ago and was eventually expanded to register point clouds from the same sensor, i.e., point clouds with similar point density depicting similar scenery. During the last 5 years, new methods evolved that can co-register point clouds from different sensors, i.e., point clouds with different point densities, using their raw and irregular data structure.

This section overviews the most recent AI-based methods for registration and co-registration of point clouds and shows how these methods relate to the use of the traditional airborne laser- and photogrammetric tie points as tie features. There are multiple ways to classify registration methods. Most recent overviews of the optimization-based and deep-learning-based registration methods are from (Bello *et al.*, 2020; Huang *et al.*, 2021; Liu *et al.*, 2021; Zhang *et al.*, 2020). The following sections adapt the high-level classification scheme (Huang *et al.*, 2021) which grouped the registration methods into four classes:



Groups of same- and cross-source point cloud registration methods. Source: Huang *et al.* (2021)

- (A) Optimization-based methods: They rely on establishing the point-to-point correspondence, and estimating and refining the transformation parameters, iteratively. These methods do not require any training data, however according to Huang, “require sophisticated strategies to overcome variations of noise, outliers, density variations, [scale variations,] and partial overlap”. Refer to Abdellazeem *et al.* (2021) for a recent example.
- (B) Feature learning-based methods: They use “Neural Networks to learn the robust point features and respective correspondence search” (Huang *et al.*, 2021). Since deep learning is deployed as their feature extraction tool, such methods promise robust and accurate correspondences and respective registration results.
- (C) End-to-end learning-based methods: Here, the estimation of the transformation parameters is embedded in the neural network optimization, which optimizes for the entire registration task. Lu *et al.* (2019), for example, proposed to detect key points, and estimate the relative pose, simultaneously.
- (D) Cross-source learning-based methods: Like (C), this category covers the end-to-end learning, however Huang assigned this its own category since in comparison with registration (of same sensor data), the co-registration must overcome the relatively higher levels of noise, outliers, and differences in point density, and thus requires more complex optimization strategies.

RE (A): This category could be split up into (A.1) ICP-/Keypoint-based methods, (A.2) Feature-based methods (see for example in (Persad, 2017)), (A.3) graph-based methods (Förstner, 2011), and (A.4) Gaussian Mixture Model (GMM)-based methods (Reynolds, 2009), or similar categories. One of the early AI-based works in (A.1) was published by Peng *et al.* (2014), which applied deep learning embedded ICP in the deep learning to achieve coarse- and fine registration of point clouds from different sensors. For more information on other works in this area; refer to (Aoki *et al.*, 2019; Bello *et al.*, 2020; Huang *et al.*, 2021; Lee *et al.*, 2021; Liu *et al.*, 2021; Zhang *et al.*, 2020).

RE (B): This group could be divided further into (B.1) point-based methods and (B.2) tree-based methods. While the point-based methods extract the point agglomerates or voxels from the point clouds that describe the respective center points, the tree-based methods convert the point clouds into trees and ingest the tree structures into deep learning. The work in (Lee et al., 2021) is an example of B.1.

A major drawback of these methods is that the training process for the feature extraction requires a lot of data input; also, this process is separate from the estimation of correspondences and transformation parameters – a step that is typically non-iterative and involves RANSAC (Mellado *et al.*, 2015). Furthermore, these methods do not consider the information of the local structures (but rather how the clusters compare). Many of these methods require that the irregular point clouds be resampled in a structured 2D or 3D grid (Bello *et al.*, 2020), or mapped into a tree or graph structure, so that the traditional AI methods can be recycled and applied to either 2D renderings from multiple views, or on voxels.

PointNet (Charles *et al.*, 2017) was the first deep learning method/library that could be deployed on raw point clouds. Many others, such as PointNet++ and Point CNN, followed thereafter. A thorough review of the latest models for feature learning with raw point clouds can be found in (Bello *et al.*, 2021; Liu *et al.*, 2019; Zhang *et al.*, 2020).

In PointNet, the point clouds are resampled to reduce resolution, and group points are extracted based on neighborhood relation, and the nearest points are mapped into a feature vector that represents the local structure around the representative centroid point. The point clusters themselves are then either correlated locally in small regions, or multiple centroids are grouped, and their shape is compared against other patches or clusters.

The earlier mentioned graph-based methods can also be used for deep learning. It is common practice to build the graph structures with kd-trees. Then, the respective kd-network is used for deep learning (Liu *et al.*, 2021). Generally, in PointNet and all its successors, finding accurate correspondences between such elements is where deep learning occurs, as this allows for classification, segmentation, and registration.

RE (C): To distinguish between type (B) and (C), Huang *et al.* (2021) followed the same train of thought used in classifying Lidar- and camera calibration: Some methods separate sensor

calibration from the data calibration, while others calibrate the point cloud/imagery through the improved sensor- and mission parameters (Friess, 2006; Lindenthal *et al.*, 2014). Moreover, unlike the AI methods in type (B), the deep learning methods within the category (C) achieve the registration through combined feature learning and data transformation. Similar advantages apply: The method is optimized for the entire registration task at hand, not only for the feature extraction. In other words, the mathematical constraints used to minimize the point cloud offsets can be formulated such that the structure of the neural network-extracted features is maintained, and both the feature extraction and the mathematical optimization function are optimized in conjunction.

RE (D): The methods in this category are essentially the same as in category (C), which are just applied to the data from different sensors (co-registration, multi-modal data). Thus, the same core techniques are used, such as extraction and matching of key-points and respective clusters (using RANSAC for example), GMMs, trees, or graphs, in combination with outlier rejection and the optimization of the motion function based on the characteristics of the extracted features. Recent research work has been focusing on the raw point cloud registration. But the co-registration of the raw point clouds from different sensors in close range applications has been picked up only during the past few years. Huang *et al.* (2021) even claimed that “based on [our] knowledge, [we published] the first learning-based method to solve the cross-source point cloud registration [in 2020]”.

To this day, the AI methods have been optimized either for the registration of different types of imagery or for point clouds from close-range sensors only, such as scans or point clouds from stereo vision of indoor scenes. Refer to the examples: ModelNet40 (Wu *et al.* 2015), 3DMatch (Zeng *et al.*, 2017), or NYUDv2, SUN3D, S3DIS, SceneNN, ScanNet data listed in (Bello *et al.*, 2020), or 3D outdoor scenery such as the famous KITTI datasets (Geiger *et al.*, 2012), or data from robotics such as Robot 3D Scanning Repository (Grilli *et al.*, 2017), or autonomous driving applications. Liu (2020) and Bello (2020) provided a comprehensive overview of the most prominent point cloud datasets that were used for deep learning.

Typically, the point clouds with the high point density are obtained from a close-range LiDAR or from photogrammetric dense point matching (Haala, 2013), while the respective point clouds

with the lower density are obtained from a stereo-photogrammetric sensor such as Microsoft Kinect, or the Asus Xtion (Swoboda, 2014). In 3D body scanning applications (D'Apuzzo, N., 2006), they often stem from structure from motion software on tablets or cell phones; in automotive applications, they could be obtained from depth sensors like an ultra-high-resolution radar (Peng *et al.*, 2019), ultrasonic sensors, or the gated vision (Spooren *et al.*, 2017).

These point clouds have a lot of information about the depth and sample the same surface elements than their counterpart. Also, they differ in point density only by approx. 1:100. Also, often, the research community refers to the point clouds with the lower point density as “sparse point clouds”. Contrarily, the tie point clouds from traditional large- and medium-format airborne applications result point densities approx. 5,000-10,000 times lower than their counterpart.

As deep learning is a technique to implement machine learning and machine learning is an approach to achieve artificial intelligence, the “sparse” point clouds used in the AI community allow for feature extraction and learning of local point cloud characteristics. From the experiments given in Chapter 6, the “extremely sparse” tie point clouds targeted in this framework pick up different surface characteristics and do not describe the surface with sufficient detail to extract features that would be compatible to that extracted from the respective Lidar data in content of resolution and information. Consequently, the current AI methods cannot be used with the proposed point cloud intake. Likewise, the application of AI methods to imagery, for example, the application of Convolutional Neural Networks (CNN) to conducting image- or patch matching (Lee *et al.*, 2021), will not work with the respective elevation maps, either.

APPENDIX B: Characteristics of Image-based and Keypoint-based 3D Tie Points and Laser Points


This Appendix gives a reference to the knowledge related to tie point generation and their comparison to laser points, because the different characteristics of tie and laser points are critical and essential in the context of co-registration.

Photogrammetric Tie Points

Requirements

Automatic tie point generation must address the following difficulties:

- Low contrast scenery
- Changes of surface characteristics, such as loss of leaves and construction sites
- Moving objects, e.g., cars and trucks
- Change in sun angle, causing moving shadow points and image collection with different exposure settings
- Strong sun reflection, causing bright spots and image artifacts
- Strong variations in terrain height or flying height, causing differences in image resolution due to different heights and scales



Effects of elevation change on tie point search when (a) flying at different heights and (b) flying over mountainous terrain

In mountainous terrain, as in (b), or when images are acquired at multiple heights ((a) above), the image scale changes within a single image significantly. Therefore, tie point engines need to be capable of stitching mountainous and/or multi-altitude data.

For bundling large format imagery, a point coverage of approx. 400 points per frame is ideal (Muick, 2018). Whilst lower point densities may result in frail relative orientation, higher point densities tend to overpower the effects of the ground- and air control and do not add further improvements to the estimation of IOP.

2D Tie Point Extraction

As outlined in the previous section, the tie point engines must meet many challenges posed by image registration. The methodologies that evolved over the past decades can be classified in two categories:

- (1) The traditional feature-based or “image-based” (Shragai *et al.*, 2011) methods deploy a local corner search by correlating corner detection masks such as the Canny-, Kirsch-, (Haala, 2002; Deng *et al.*, 2008) or Edison (Meer and Georgescu, 2001) operator. The


search for matching points has been conducted on multi-scale image pyramids: The matching starts on images of low resolution, and then the algorithms refine the position of each potential tie point layer by layer, until its position has been located in the full-resolution imagery.

The “image-based” algorithms define the search area in a neighboring image by projecting the pixel location from one image onto the ground and from there back into the neighboring image (Bombardier, 2013). Thus, the input of an approximate terrain height and the approximate EO parameters is required. However, the image pyramids can be recycled and used to accelerate the image rendering during the stage of ground control pointing, as done in Intergraph’s ISAT software (Hexagon Intergraph, 2018), Trimble Applanix CalQC software (Mostafa, 2001), and the early versions of Vexcel’s UltraMap software (Wiechert *et al.*, 2012).

- (2) “Key-point based” registration algorithms are more robust, require less memory consumption and less computer power than the image-based algorithms (Lingua *et al.*, 2009). To extract the points of interest, their implementations typically rely on the SIFT, SURF (Bay *et al.*, 2008), or ORB descriptors (Rublee *et al.*, 2011). A widely used public library for generation of key-point based sparse tie points and dense points is VisualSFM (Wu, 2018). It integrates other software tools and libraries, such as SURE (Rothermel *et al.*, 2012) for photogrammetric SURface REconstruction and Microsoft’s PMVS/CMVS tool chain for multi-scale stereo and dense reconstruction (Furukawa *et al.*, 2010) and supported the aforementioned descriptors.

As will be shown in the next section, 7 or more evenly distributed tie points per image pair are sufficient for the standard photogrammetric post-processing process (where no camera self-calibration is required) because the denser coverage does not necessarily lead to more accurate registration results.

Whilst the conventional photogrammetric tie point engines may produce more tie point matches in urban areas, the key-point based algorithms generate tie points that are distributed more evenly across the imagery (figure below).



Tie points matched using (a) the Harris corner detector and (b) the SIFT operator.

Source: Li *et al.* (2012)

The key point-based tie point engines found their implementation in the newer tie point engines, such as Cardinal's VR software (Cardinal Systems, 2018), the camera module of Optech's LMS software (Teledyne Optech, 2018), and Vexcel's UltraMap software (Wiechert *et al.*, 2012). The table below summarizes the pros and cons of the image-based and key point-based image registration algorithms.


Tie Point Density

As outlined in the previous section, the tie point clouds of traditional large area mapping applications (acquired by large- or medium format cameras) are spaced irregularly and describe the uppermost layer of the surface. In aerial applications, Z always points upwards away from the ground. Since rural scenery provides less contrast than urban scenery, the tie point distribution in rural scenery is lower and also more irregular than that in urban scenery, especially if the image-based matching algorithms are used. What determines the shape of the tie point clouds is however

the software design, i.e., the intend of what the tie points are used for in the photogrammetric processing chain. The two figures below give an example where tie points of the same two images were generated by two software packages and used differently:

- The tie points in the first figure below were generated using the well-established ISAT software (Hexagon Intergraph, 2018), which uses tie points primarily for image orientation.
- The tie points in the second figure below were generated by UltraMap software (Wiechert *et al.*, 2012), which uses the tie points for image orientation, camera calibration, and generation of a course surface model.

In the figures below, a cardinal rotation of 90 degrees causes the images with different heading orientation. The images were collected with 60% overlap in flight direction and 30% overlap across flight direction, which is typical for regular survey missions. The image sensor's physical size is 104 mm x 68.4mm. And the units of the axis in the plots below are millimeter.








Image-based, image-based (ISAT-generated) tie points of two UltraCam Xp images






Key-point based (UltraMap-generated) tie points of two UltraCam Xp images


By comparing the figures above, the ISAT-generated tie points are obviously far less dense than those generated by the key-point-based UltraMap engine. However, the small number of ISAT-generated tie points is sufficient for stitching the images together to one coherent block with high accuracy, which is ISAT's primary purpose.

For camera calibration however, – particularly for the reliable estimation of lens distortion parameters, a more homogenous and denser point coverage is required. As can be seen from the figure below, the optical lens distortion shows small irregular geometric variations across the entire image space, which tells why the tie point observations are required all-across the image space.




Estimated image distortion residuals across image space

As illustrated in the figures above, for survey flights that are flown with little overlap, such as 60% overlap in line and 30% to 40% overlap across flight direction, the difference in point density can create a banding-artifact across the image block, i.e., there are more points in the areas where the images overlap with the neighboring flight lines. The figure below illustrates how the ISAT-generated tie points accumulate in 3 areas when combined graphically in one plot.



Distribution of ISAT-generated tie points accumulated across image space

The figures below shows another example, where imagery was collected with 60% (longitudinally) and 40% (laterally) overlap, resulting in a visible banding effect of the tie points.



Banding effect of tie points in a flight block with 20% and 45% image overlap

The above figures is an example of the tie point banding, where two flight lines with four images each were collected with 20% overlap in flight direction and approx. 40% overlap across flight direction with an UltraCam Xp camera. Yellow dots represent the 2-ray tie points, and the tie points in orange resemble the 3- and 4-ray tie points. In this example, the tie points were generated with key-point matching routines, resulting in 1852 2-ray points, 127 3-ray points, and 577 4-ray points, respectively, and a staggering 2536 tie points in total.

Whilst most of the large format digital cameras are typically flown with more overlap in flight direction (often 60% to 80%), banding can still occur when the side overlap is chosen to be smaller than 50%. Line-scanners like the ADS100 (Hexagon Leica, 2018) and the UltraCam Condor (Vexcel Imaging GmbH, 2018) are often flown with less than 40% side overlap, causing the band effects to be more prominent.

Whilst the sparse, banded tie point distribution shown in the last figure above is sufficient for accurate relative orientation of flight imagery, it is too irregular for a reliable estimation of the lens distortion parameters (GIP Aalen, 2018). To achieve a more homogenous point distribution and increase the number of the image rays per point with the older traditional photogrammetric software (such as ISAT), the flights for camera calibration are conducted with increased image overlap, typically 80-90% overlap in flight direction and 50-60% overlap across flight direction. The increased overlap results in a homogenous, band-free point distribution - see section 2.3.23.2 for illustrative examples. Without loss of generality of the foregoing, only the ISAT-generated tie points that are free of the banding effect will be used in this research for the co-registration with LiDAR data.


Laser Points

Laser Point Density

The ground sampling of an airborne laser scanner is irregular. Thus, unlike in aerial photogrammetry where the density of measurements is described as Ground Sampling Distance

(GSD) in meter or feet, the measurement density in LiDAR applications is described as point density in points/m² or points/ft². The point density varies also with the degree of overlap of neighboring scan lines: The areas of overlap show a higher density than the areas where no overlap of laser swaths occurs:-

The point density varies also with the angle of impact on the ground with respect to the detector, and the reflectivity of the scanned surface: Sloped roof tops have more surface compared to planar surface patches with the same aerial footprint and thus count more laser hits. Areas with vegetation tend to absorb some of the laser pulses, which can lead to a decrease in counts of returned signals. These effects can be seen in the figure below.




Point density of a LiDAR survey (Lindenthal *et al.*, 2015)

The figure above illustrates the point density of a LiDAR block acquired in Oshawa, a suburban area of Toronto. Eight lines were flown in North-South and two lines in East-West direction. Areas highlighted in green are the result of overlap from neighboring swaths, often covered with vegetation, whilst blue areas indicate areas with a high density of roof tops.

Comparison of LiDAR and Photogrammetric Tie Points

As illustrated in the previous section, an airborne LiDAR sensor actively scan the surface. Linear or planar surface elements are required to connect neighboring scan lines. Photogrammetric applications require that the ground is depicted at least from two different locations, whilst a single spot is sufficient for LiDAR to derive 3D information. The figure below highlights the sensor paradigms.



Archetype of laser scanning and stereo photogrammetry. Source: Kraus (2007)

The different nature of the sensors invokes certain contrasts in data characteristics and accuracies. Some of the synergic effects on LiDAR and photogrammetric data are (Hu *et al.*, 2011; Habib *et al.*, 2007):

- The horizontal (across line of sight) errors from photogrammetry are usually smaller than that from LiDAR whereas LiDAR can usually obtain higher vertical (in line of sight) accuracy (Baltsavias, 1999).
- The accuracy of a photogrammetrically obtained DSM is critically dependent on the flying height and the camera type while the accuracy of a LiDAR DSM depends on the point density and is almost independent of the flying height (see section 0).

- Due to the irregular nature of laser point clouds, it is harder to extract object boundaries, such as terrain breaks and contours of buildings, from LiDAR than from imagery. The extraction of building contours and treetops from LiDAR data is potentially affected by systematic errors in extraction (Matikainen *et al.*, 2009; Yu *et al.*, 2004).
- Laser scanning has the penetration capability with multiple signal echoes that can enable recovery of vertical structures like trees, while it is almost impossible to do so from imagery. Thus, DTMs are mostly derived from LiDAR data and not from imagery.
- Photogrammetry requires only color contrast so that tie points can be matched reliably. Laser pulses on the other hand get absorbed in water. Some metallic material, for instance copper, also do not return laser echoes.
- The location of a laser footprint is not directly visible. Most airborne laser scanners are therefore flown in tandem with a medium format camera, so that the laser points can be colorized, or the pixels be provided with their depth and can form texels (Budge *et al.*, 2015).

As will be shown in the upcoming sections, the most challenging discrepancies in point characteristics in co-registration are: a) the extremely sparse distribution of (ISAT-generated) tie points, b) the tie- and laser point clouds describe different details of the surface, and c) the tie points tend to locate in areas where there are insufficient laser points to accurately describe the geometry of the local surrounding. Typically, the disagreements between these two types of data are interpreted as errors in the land cover map derived from the remotely sensed data (Congalton and Green, 2008). However, for the reasons above, this interpretation is not valid when the conjugate laser- and tie point maps are compared.

APPENDIX C: Test Data

In this framework, three datasets were used for testing - one dataset of the simulated height maps was created and two were the real-world datasets. To address the question, to what degree these three datasets represent “typical and universal” airborne survey applications, the characteristics of a popular test site are highlighted, for comparison, first. As mentioned earlier, close range applications, such as terrestrial and UAV data collection, are not considered here. The ratio of the tie point spacing, and the captured elevation is vastly different for traditional large area mapping and UAV mapping applications (UAV data typically samples the surface with a tie point roughly all 2-3m while traditional airborne applications sample the ground with a tie point approx. all 100m), and thus, tie point clouds from UAV and terrestrial applications have very different characteristics.

Reference Dataset

One of the world’s most flown test sites for evaluating large- and medium format digital airborne sensors is the test field in Vaihingen and der Enz, Baden-Württemberg, Germany. While data from this area was not used, the terrain characteristics of this area can serve as a guideline for the simulated dataset and for comparison against the real-world datasets used in this thesis.

The Vaihingen and der Enz test site was established in 1995 by the University of Stuttgart for testing the geometrical performance acceptance of one of the first operational digital airborne line scanning systems (IFP, 2017). From 1996 onwards, the test site has been used by the OEEPE-ISPRS group to evaluate the performance of DGPS (Cramer, 2001), PPP (Heipke et al., 2002; Bisnath and Gao, 2009, Yuan *et al.*, 2009) and aerial triangulation with digital images (Heipke, 1999) etc. Numerous agencies and companies such as the DLR, Leica, Riegl, Vexcel Imaging, and Zeiss had their sensors flown over the test site for research and evaluation purposes.


The Vaihingen/Enz test site provides the following features:

- High density ground control - almost 200 control points have been spread across the area
- Mostly gently sloped terrain. Few areas with steep slopes
- A mixture of urban scenery and farmland
- Small areas with forest
- Small waterbodies, which typically create a challenge for LiDAR and photogrammetric tie point matching

The test site is enclosed with the red line in the figures below.



Aerial view of the OEEPE test site in Vaihingen/Enz, Germany. The dark green areas depict forest and water bodies. Source: IFP (2017)



Digital terrain model of the Vaihingen/Enz test site. The terrain variations are exaggerated by factor two. Source: IFP (2017)

The central area – enclosed with a purple line in the figure above – possesses dense distribution of GCPs and is intended for conducting high-resolution data acquisition and large-scale tests. This area is about 5050m (east-west) by 2750m (north-south) in size. The entire test area – enclosed with a red line in the figure above – is about 7400m (east-west) by 4700m (north-south) in size. The terrain elevations range between 193m and 356m above sea level. In Figure 3.2 above. The changes in elevation are exaggerated by factor two.

GCPs are painted targets mostly, which are intended to be used for block adjustment and to serve for accuracy verification of registration results. In total, 275 GCPs are available, of which 172 were painted 0.6m by 0.6m white squares and 103 were painted 0.3m by 0.3m white squares with black centre, and spread evenly across the test area (IFP, 2017). All points were surveyed with 2cm absolute accuracy (RMSE) in each of the three coordinates components.



Ground control points of the Vaihingen/Enz test site. Source: IFP (2017)

The Vaihingen/Enz test site served as a guideline for what the simulated test dataset should look like. It also confirmed that the real-world datasets prepared for testing the proposed co-registration algorithms well represent “typical” airborne survey data.


Real-World Datasets

This section introduces the two real-world datasets that have been used for testing some of the co-registration algorithms presented in this dissertation. First, it is shown how these datasets were

generated. The layout of the flight lines will be shown, followed by the footprints of LiDAR and camera data. Some screenshots of the fine-registered reference data will be presented to show the overlap, accuracy, and the noise level of the tie- and laser points.

Data Preparation

The process of data preparation is quite complex as six software packages were chained and intermitted data had to be input and output before the final product – the geometrically rectified and co-registered laser and tie point clouds – became available. The diagram below shall suffice for a high-level overview of the process.



Data preparation workflow of the real-world datasets

As in the figure above, the data preparation started with the flight planning of the data collects using FMS planner software (Teledyne Optech, 2013). During the data collection, DiOPS firmware (Teledyne Optech, 2013) was used to store the values of real-time parameters and data. After the data collection, the post-processing started by refining the GPS/IMU trajectory using

POSPac (Scherzinger, 2000; Hutton and Mostafa, 2010; Trimble Applanix, 2018), followed by developing the images in CaptureOne (PhaseONE, 2018), and conducting the camera calibration whilst generating the 3D tie points using ISAT (Hexagon Intergraph, 2018) and Bingo software (GIP Aalen, 2018). The LiDAR data was developed and calibrated through Optech's ALTM (Teledyne Optech, 2018) and LMS software (Friess, 2006; Lindenthal *et al.*, 2015). To generate the calibrated and fine aligned reference solutions, the 3D tie points, and laser points were cleaned from outliers and co-registered using TerraSolid software. In the figure above, the software modules from Optech were highlighted in red whilst the software components from third parties were highlighted in blue.

Oshawa Dataset

Flight Layout


Both the Oshawa data and Keyworth-Nottingham data showed flight lines instigating large degrees of data overlap (60% side lap), resulting in a practically band-free tie point cloud. This is untypical for regular survey missions, which would form much of the data input for coarse co-registration. However, these types of flights are common when data is collected for cross-aided sensor calibration. Here, accurate fine registration is mandatory. Also, co-registration approaches should be robust and be able to handle data with various geometric conditions. Thus, these datasets make great test material for the proposed co-registration methodologies.

The Oshawa data consisted of one “boresight block”, flown for deriving camera boresight parameters, and some additional flight lines acquired for LiDAR testing purposes. The boresight flight block was used for self-calibration of the LiDAR without ground control (Lindenthal *et al.*, 2015) and for cross-aided sensor calibration. The Oshawa flight divided into three sections:


- (1) Camera boresight collect: One altitude block for data processing in TerraPhoto software (Soininen and Burman, 2005) with 7 lines flown at 900m asl at 60% end lap:
 - 3 lines south-east to north-west, 60% side lap
 - 1 cross line south-west to north-east
 - 2 cross lines south-west to north-east acquired in test mode, excluded from processing

- 1 diagonal line
- (2) Additional lines used for LiDAR calibration purposes collected with 60% end lap:
- 2 perpendicular lines, 600m asl, included in processing
 - 2 perpendicular lines, 400m asl, excluded from processing

Only the data collected from the flight lines at an altitude of 850m and 600m asl were selected for co-registration. The flight lines flown with large overlap result in the higher laser and tie point densities than the areas with small overlap. The figures below show the layout of the selected flight from the Oshawa data. The resulting image block comprised 105 photos, 3641 tie points after the outlier removal, and 51 measured control points.




Oshawa flight vertical profile. Only data within the blue encircled times was used




A perspective view of the data selection (Oshawa flight)

Camera Data

As in the figure below, the Oshawa calibration site provided a very dense distribution of control points. All 51 ground control points within view were measured.




In the figure above, the tie points within the flight lines, also referred as the pass points, were depicted with the purple IDs. Their distribution is typically limited to the outskirts of the survey block. The tie points in-between flight lines were depicted with the point ID in turquoise. Since these points required the image overlap across the flight lines, a dense tie point distribution was expected to occur mostly in the block center. The homogenous and dense distribution in the combined display of the pass- and tie points (collectively called tie points) indicated good tie point matching results.




The combined display of the tie- and pass points (Oshawa test site)

As illustrated below, the area where multiple flight lines were flown close to each other in the North-South direction showed a particularly high tie point density. As can be seen from the figures below, the tie and control points clearly captured the changes in elevation of the terrain.



Oshawa tie and ground control distribution



The Oshawa tie and ground control points, colored by elevation


The tie points were generated using the Intergraph ISAT module (Hexagon Intergraph, 2018). After the blunder removal, 3641 tie points remained, equivalently about 35 tie points per image. The tie points covered an area of approximately 3.2km by 3.2km as the image footprint at the lower altitude was approximately 0.6km² (roughly 0.9km by 0.7km) leading to an approximate tie point density of 0.0006pts/m². A maximal GPS shift of 18cm and a minimal shift value of 3cm per line confirm good quality with the geo-referencing.

Conversion of the Tie Point Cloud to an Image


There are many ways to convert the point cloud into the regularly spaced data – see for example Fernandez *et al.* (2007) for more details. A surface grid was fit into the tie points using MATLAB software (MathWorks, 2019). MathWorks (2019) provided 7 interpolations methods for the 2.5D surface data: (1) Linear, (2) nearest neighbor, (3) cubic spline, (4) shape preserving, (5) biharmonic, (6) thin-plate spline, and (7) natural. Testing and visually comparisons of all these methods showed that the “natural” interpolation method delivered the best-fitting and most natural looking results for elevation data. Unfortunately, MathWorks (2019) still did not provide detailed information about this interpolation method, however, based on the works in (Yanalak, 2003; Yilmaz, 2007), it could be the Spline natural neighbor interpolation. The methods for the traditional Kriging (Arun, 2013; Wen, 2016; Schwanghart, 2020) and the gradient-enhanced and Kriging with incorporation of error information (De Baar, 2020) could also be found in the MATLAB user community (MathWorks, 2019). These methods may perhaps yield even flatter contours at the outskirts. But they surfaced years after the processing of these datasets was completed. Also, based on the user’s feedback in the MATLAB forum, their performance seemed questionable.

The figure below presents the tie point cloud overlaid with a natural-fit surface grid. The areas with the higher tie point density were captured with multiple flight lines and more image overlap than those areas with the lower point density. *For the average point density of 35 tie points per image, MATLAB considered a cell size of 40m by 40m to be optimal.* The second figure below

gives the respective grid, rendered to an image file. This image was used for testing the co-registration algorithms in this thesis.




The Oshawa tie point cloud overlaid with the MATLAB-generated surface mesh




The Oshawa tie point grid with 81x115 40m cells

LiDAR Data

Since the LiDAR and camera were flown at the same time, their data showed a similar layout. An elevation-coded plot of the laser points is given below. A closer look revealed that the raw LiDAR data was polluted by the atmospheric and other false returns. As can be seen, a large portion of the atmospheric outliers was situated in proximity of the sensor – also referred to as the cone of death (Teledyne Optech, 2012). Other outliers occurred more sporadic and the respective point locations were often so isolated that they could be seen from its appropriate enlargements: The outlier with the largest negative distance from the surface for example is 533m below the surface and invisible in the figure thereafter.




Top view on the full-resolution Oshawa laser point cloud (63Mio points) before outlier removal (outliers are marked in black)




Vertical cross section of full-resolution Oshawa laser point cloud before the outlier removal

The false returns were removed since they could have negative impact on the co-registration. As shown below, after cleaning the laser point cloud there was only one red line remain, which was caused by the reflection of a power line. Similarly, the small red quadratic spot in the south-east end of the block was not a remaining anomaly but stemmed from a high rise.



Oshawa, every 4th laser point (16Mio points), after the outlier removal, top view



Oshawa, every 4th laser point, after the outlier removal, vertical section

The full resolution point cloud had approximately 63Mio points. Covering an area of approximately 4.5 by 4.5km, the resulted average point density was 2.4 pts/m². Since the survey area did not cover the entire square-shaped grid that MALAB had fit, but rather a cross-shaped area roughly 2/3s the size of the full grid area, the actual average point density was roughly 4 pts/m².

Down-sampling Laser Points

During initial discussions, it seemed desirable to use the LiDAR at full resolution to ensure all information content is considered in the process of co-registration. However, handling millions of points (63M in case of the Oshawa dataset, 35M in case of the Keyworth-Oshawa dataset) in MATLAB would have been dreadful. A closer look revealed that handling the full-resolution data was not necessary:


- In case of the fine registration method presented in Chapter 4.4.2, roof polygons had already been extracted automatically from the full resolution point cloud and used as registration primitives; there was no need to repute the point cloud any further.
- In case of coarse co-registration, the registration algorithm can only make use of information content that is common to both datasets. All the extra information captured in the high-resolution laser point cloud had no counterpart in the tie point cloud. Thus, as long as the information content of the laser point cloud remained equivalent or slightly higher than that of the tie point cloud, cross-correlation could occur without restriction or loss of information and the co-registration results would be the same as those obtained using the full-resolution laser point cloud.

According to the Nyquist theorem (Fritsch, 1999), the laser point cloud should be sampled with the same or higher point density than that of the tie point cloud. Considering that the tie points were distributed very irregularly whilst the laser points, even when down-sampled, were distributed regularly in comparison, the laser points were down-sampled to 8775 points, i.e., a point density slightly higher than twice that of the tie point cloud.


The figure below displays the laser point cloud overlaid with a naturally-fit surface grid. Areas with higher laser point density were captured with multiple flight lines and more scan overlap than those areas with lower point density.

Conversion of the Laser Point Cloud to an Image

Once the LiDAR had been down sampled, both point clouds were gridded and rendered to images. The figure below illustrates the down-sampled point cloud overlaid with the grid. Figure 3.6 displays the respective grid, rendered to an image file. The color coding is the same as in Figure 3.3, however the large descent of the grid in the north-east corner causes the center of the grid to occur more yellow. The image file shown in Figure 3.6 was used for testing the co-registration algorithms described in Chapter 4.0, Chapter 4.0, and Chapter 5.0.




The Oshawa laser point cloud overlaid with the surface mesh



The Oshawa laser point grid with 117x138 40m cells

Co-Registered Data


The Oshawa tie and laser points showed a ratio of $3,700/63,000,000 = 1:17,000$, which confirmed the ratio of about $1/10,000$ that one would typically obtain from regular survey data. The co-registered point clouds are given below. The full-resolution laser data was so dense that it covered up the tie and ground control points. To demonstrate that the tie points show good fit and fall above and below the laser surface, the laser point cloud was thinned out by factor 10 so that the ground control and tie points became visible.




Overlay of every 10th Oshawa laser point (grey), GCPs (red), and the tie points (blue)

Originally, this research started out with the intend to only cover aspects of the fine registration methodology discussed in Chapter 4.4.2. Therefore, the point clouds were aligned as tightly as possible during data preparation; - the quality of fit is illustrated in more detail in APPENDIX G. Later, this highly accurate solution served as a reference not only for evaluating the fine-registration algorithms but for the coarse co-registration algorithms, too.

Figure 3.7 below displays the co-registered point clouds. In this graph, the tie points had been elevated by 10m to highlight the boundaries of the tie point area with respect to the laser points. As can be seen below, the negative curvature of the laser point cloud in the north-east end caused the mesh to extrapolate into negative height values. The laser points covered an area of approx. 4.5km by 4.5km whereas the tie points covered an area of approx. 3.2km by 3.2km.




Overlay of the Oshawa laser point cloud, 10m elevated tie point cloud, and surface mesh



Rotated rendering of Oshawa point clouds highlighting an extrapolated steep slope in the north-east corner

The high density of the laser points allowed for deriving an accurate DTM. If the photo and laser data have been aligned correctly, the difference of the DTM and the photogrammetric ground control points should average zero; Also, small vertical deviations should indicate good quality of the tie- and ground control points. The figure below shows the results of that experiment.




Vertical difference between Oshawa GCPs and LiDAR derived DTM in meters

As can be seen from the figure above, the surveyed ground points aligned with the laser DTM nearly perfectly, however the sporadic large deviations from zero indicated poor quality of the surveyed ground control.


As shown afore, the tie points rarely fall into the centre of the flat surfaces, but rather locate at corners of objects, usually away from the laser points. For example, the tie points tend to locate on top of trees or at corners of the rooftops whereas laser points tend to fall either onto bare ground near the respective tree or roof, or flat onto the roof surface. Thus, the tie points show occasionally a vertical difference of 10-15m from the surrounding laser points, which is roughly the height of trees and most buildings in sub-urban areas. The figure below confirms this observation and shows

that the distances of the tie points to the LiDAR key model points vary in the range of approx. $\pm 10\text{m}$.



Vertical difference between the Oshawa tie points and the LiDAR derived DTM

Since the tie points were not classified and were too random and loosely distributed to allow for reliable retrieval of information about the objects on which they located, the Z deviation from the ground model could be regarded as a vertical noise band. Thus, a 10m noise band was chosen in the DSM simulation (Chapter 5.6). The figure below shows this context at higher resolution for image 9, located at the center of the flight block. More details of image 9 are given in APPENDIX G.



Vertical difference between the Oshawa tie points and the key model ground points of image 9, in meters


Keyworth-Nottingham Dataset

Flight Layout


The Keyworth-Nottingham (KN) dataset consisted of a reduced “boresight block”, typically flown for calibrating camera boresight parameters, and additional flight lines for LiDAR testing purposes. The boresight flight block was also used for self-calibration of the LiDAR without ground control. The flight is divided into two sections:

- (1) Camera boresight data collect: One altitude block layout for data processing in TerraPhoto software (Soininen and Burman, 2005) with 9 lines flown at 880m asl at 60% end lap:
 - 4 lines south-east to north-west, 60% side lap
 - 4 cross lines south-west to north-east
 - 1 diagonal line
- (2) Additional lines used for LiDAR calibration and camera validation purposes were collected at 1250m asl with 60% end lap:
 - 1 line south-west to north-east, included in processing
 - 1 perpendicular line, included in processing

The resulting image block comprised 64 photos, 2705 tie points before and 1871 tie points after blunder removal, and 38 measured control points. Flight lines flown with high overlap resulted in higher laser and tie point densities than areas with small overlap. The two figures below illustrate the flight pattern. The camera footprint view of the survey block is displayed below as well.




KN vertical profile. Only the data within blue circles was used



A perspective view of the data selection (KN flight)


Camera Data

As illustrated below, the Keyworth-Nottingham calibration site provided a dense distribution of the control points. 38 ground control points had been surveyed.



KN block layout plotted in the ISAT footprint viewer (- the measured control points are in red)


In the ensuing figure, the tie points within flight lines, also referred as pass points, were depicted with their point IDs in purple. The distribution of the pass points was typically limited to the outskirts of the survey block. The tie points in-between flight lines were depicted with the turquoise point IDs. Since these points require specific image overlap across the flight lines, the tie point distribution was expected to occur mostly in the block center. The homogenous and dense distribution in the overlaying presentation of the pass- and tie points (collectively called tie points) in the figure below confirmed the good matching results of the tie points.




The combined display of the tie- and pass points of the (KN test site)

Conversion of the Tie Point Cloud to an Image

A surface grid was fit into the tie points using MATLAB's "natural" interpolation method (MathWorks, 2019). The ensuing figure displays the tie point cloud overlaid with the natural-fit surface grid. The areas with the higher tie point density were captured with multiple flight lines and the more image overlap than those areas with the lower point density. For the average point density of 30 tie points per image, MATLAB considered a 40m-by-40m cell size to be optimal. The second figure below plots the matching grid, rendered to an image file. This image was used for testing the co-registration algorithms throughout this thesis.




The KN tie point cloud overlaid with the MATLAB-generated surface mesh




The KN tie point grid with 36x35 40m cells

LiDAR Data


Since the LiDAR and camera were flown at the same time, their data showed similar layout. An elevation-coded plot of the laser points is shown below. Obviously, the LiDAR data was corrupted by the atmospheric and other false returns. A large portion of the outliers was situated along a vertical ray, visible in the figure below, caused by interference of the receptor and the pilot's radio signal. Other outliers occurred more sporadic and often the respective point locations were so isolated and could be brought out in the figure by significant enlargement. The following figures show the laser point cloud after the cleaning process.



The KN laser point cloud before the outlier removal



The KN laser point cloud after cleaning, side view




The KN laser point cloud after outlier removal (approx. 35M points), top view

Down-sampling the Laser Point Cloud


As conversed earlier, it was desirable that the laser point cloud, which serves as reference point cloud during co-registration, contains slightly more geometric information about the surface than the tie point cloud. This was achieved by down-sampling the laser point cloud to a point density slightly higher than that of the tie point cloud. The KN laser point cloud had approximately 35M points covering an area of 2.4km by 2.4km resulting in a point density of 5 pts/m² in its original configuration; the tie point cloud, on the other hand, covered an area of 1.5km by 1.5km with 1871 points, resulting in a point density of roughly 0.0008 pts/m². After down-sampling, the laser point cloud contained 2705 regularly distributed points.

Conversion of the Laser Point Cloud to an Image

Once the LiDAR had been down-sampled, both point clouds were gridded and rendered to images using the “natural” interpolation algorithm (MathWorks, 2019). The figures below illustrate the down-sampled point clouds overlaid with the grid, and the respective grid, rendered to an image file. The color coding was the same before, however the large descent of the grid in the north-east corner caused the center of the grid to occur more yellow.




The KN laser point cloud overlaid with the MATLAB-generated surface mesh



The KN laser point grid with 59x71 40m cells

Co-Registered Data


The co-registered point clouds are displayed below. The full-resolution laser data was so dense that it covered up the tie and control points. To demonstrate that the tie points showed a well-balanced fit and fell above and below the laser surface, the laser point cloud was thinned out by factor 10 so that the ground control and tie points became visible.



Overlay of every 10th KN laser point (grey) with GCPs (red), tie points (blue), and image footprints of photo 50 (purple) and 58 (green)


Originally, this research started out with the intend to only cover the fine registration methodology introduced in Chapter 4.4.2. Therefore, the point clouds were aligned as tightly as possible during data preparation. Later, this highly accurate solution served as a reference not only for evaluating the fine-registration algorithm but also the coarse co-registration algorithms.

In the figure below, images 50 and 58 show an example of the tight fit of the laser data and ground control – the control points fell mostly underneath the laser surface.




Overlay of every 10th KN laser point (grey) with ground control (red) and image footprints of photo 50 (purple) and 58 (green)

Displayed above are the co-registered point clouds. In this graph, the tie points had been elevated by 10m to highlight the boundaries of the tie point area with respect to the laser points. The laser points covered an area of approx. 2.4km by 2.4km and the tie points covered roughly 1.5km by 1.5km ground.



Overlay of the KN laser points, the 10m elevated tie points, and the surface mesh


The high laser point density allowed for deriving an accurate DTM. If the photo and laser data had been aligned correctly, the difference of the DTM and the photogrammetric ground control points should have averaged to zero; Then, small vertical deviations would indicate good quality of the photogrammetric ground control points. The figure below shows the results of that experiment.



Vertical difference between KN GCPs and LiDAR-derived DTM, in meters

As can be seen, the surveyed GCPs showed an offset of about -20cm from the laser DTM and the sporadic large deviations from zero confirmed poor quality of the surveyed ground control and/or laser DTM.

As illustrated in more detail in APPENDIX B, in suburban areas tie points tend to locate rather on top of trees or at corners of rooftops whereas laser points tend to fall either onto bare ground in proximity of the respective tree or roof, or flat onto the roof surfaces. Thus, tie points showed every so often a vertical difference of 10-15m from the surrounding laser points, which is roughly the height of trees and most buildings. Like the figure above, the figure below demonstrates the scattering of the tie points with respect to the laser derived DTM.



The vertical difference between the KN tie points and the LiDAR-derived DTM, in meters

Since the tie points were not classified and since they were too randomly and loosely distributed for reliable retrieval of information about the objects on which they located, the Z

deviation from the ground model could be regarded vertical noise. This finding lead to choosing a 10m noise band when simulating a DSM, as described in the next sections.

Simulated Dataset

One dataset was simulated by superimposing five 1D harmonics to a signal that mimicked the elevation magnitudes of the Vaihingen/Enz dataset (see section 3.2 above). To have a similar extend than the KN dataset, a signal length of 2km was chosen. The software code shows an excerpt of how the signal was composed in MATLAB:


```
% Magnitude in meters
a0 = 30;
a1 = 70;
a2 = 35;
a3 = 25;
a4 = 7;
a5 = 5;

% signal offset in meters
d0 = 300;

% ramp in meter
ramp = 90;

T = 1; % Sampling period, 1-meter intervals
L = 2000; % Length of signal in meter increments
t = (0:L-1)*T; % Sample vector
x_0 = a0*sin(2*pi*(t-p0)/(1*L))'+d0; % phase = 0, Amp= -/+0.5, F=25Hz
x_1 = a1*sin(2*pi*(t-p1)/(2*L))'+d0; % phase = pi/4, Amp= -/+1, F=25Hz
x_2 = a2*sin(2*pi*(t-p2)/(1/2*L))'+d0; % phase = pi/2, Amp= -/+2, F=50Hz
x_3 = a3*sin(2*pi*(t-p3)/(1/2*L))'+d0; % phase = 3*pi, Amp= -/+4, F=150Hz
x_4 = a4*sin(2*pi*(t-p4)/(1/3*L))'+d0; % phase = pi/8, Amp= -/+0.25, F=300Hz
```


The signal of the fifth harmonic component was chosen to have a magnitude of 5m, simulating finely resolved details on the ground. The resulting synthetic 1D elevation signal and its components are displayed below.



The 1D simulation of an elevation signal


To convert the assembled 1D signal into a 2D signal, its square root values were multiplied by themselves over a grid of 2000x2000 points. As can be seen above, the 1D signal had an absolute height offset of 250m. This height offset was chosen to mimic an average terrain height of 250m above sea level (asl), similar to the average elevations of Oshawa and KN.

A second version was generated where the minimum value of each 1D signal component was subtracted from the combined 1D signal, so that the border “edge” values of the assembled 2D dataset were close to zero. Subtracting the minimum value minimized elevation discrepancies at the borders such that the FFT (Cooley and Tukey, 1965) could observe little friction and create a small “+” effects (see Chapter 5.1.2.2.2 for more information). The signal was then shortened from 2000 pixels to a length of 1864 pixels since a few pixels were cut from both sides so that the border values approached zero values more closely. The resulting 1864x1864 DTM grid of 1m pixels is displayed in perspective- and top view below.



“Version 1” of the simulated DTM with the low-elevation border lines


The tilted version of this DTM, depicted below, was created to mimic slope – large portions of the Oshawa data for example and many alpine areas are imposed by a steadily increasing ramp-like change in elevation.



“Version 1” of the simulated DTM with tilt, mimicking sloped terrain

In the original non-tilted version, the DTM showed two peaks with the identical shape and height. By tilting the dataset (shown above) one of the peaks became higher than the other, since its absolute height increased.


The DTM was also modified to have the circular boundaries. The round boundaries imitate the effects of a window mask which – as described in Chapter 5.1 – could be advantageous in the process of the frequency-based co-registration. As shown in below, a small groove was cut out, too. This groove had two purposes: Firstly, it served as visual indicator of the rotation angle and allowed for fast visual inspection of the co-registration results. Secondly, it introduced asymmetry to the elevation profile, so that co-registration algorithms can retrieve rotational offsets without ambiguity.




“Version 1” of the simulated DTM with circular cut-off and trench

Both the sloped DTM and the rounded DTM were deployed during the early stage of FMT testing. Through experimenting with these data sets it became clear that the approximation of real-world slope effects through tilt as plotted above were ineffective for FMT and AFMT based co-registration (see Chapter 5.0, FMT and AFMT) since frequency analysis considers relative changes of the elevation profiles, which remained unchanged by the tilt. Further, while the tilt introduced the desired asymmetry of the absolute height values (as it increased the absolute height of one of the two mirror-symmetric peaks) the relative elevation changes still showed mirror-symmetry. The circular DTM dataset on the other hand proofed to invoked critical loss of information and resulted in degraded co-registration results when compared to the co-registration results of its rectangular counterpart.

Based on these learnings, rectangular borderlines were chosen instead of round boundaries and the elevation of one of the two symmetric hills was increased. Also, instead of keeping the pixel values close to zero at the border lines, it seemed more realistic to let the border pixels change magnitude. At the same time, it seemed desirable to keep the “+” effect small. To fulfill these requirements, the data was shifted horizontally. The result is displayed in the figure below. Here, the border values still transitioned seamlessly from one side to the opposite, keeping the magnitude of the “+” effect small.



“Version 2” of the simulated DTM with varying yet symmetric elevation borders




“Version 2” of the simulated DTM with symmetric borders, rear view

As can be seen from the figures above, the shifted “Version 2” DTM showed more realistic border values than the original “Version 1” yet allowed for seamless transition between opposite borders (that appear neighboring inside the FFT). After tilting the surface model back to the ground, the two second-highest peaks showed the same height and shape again, so one of the peaks had to change elevation to eliminate mirror symmetry. The change in elevation had to be strong enough to be clearly visible in the respective DSM dataset, too. Also, rather than turning the symmetric peaks into one pointy outstanding hilltop, the intend was to maintain the resemblance of real-world conditions as much as possible and keep elevation changes relatively small.


These requirements were met with “Version 2” of the simulated DTM, internally referred as “X5d2”, which is depicted in the left column of the figure below. Here, a Gaussian bell was fit into the elevation peak “B” to increase its height relative to peak “C” and to make the surface model non-symmetric. The second highest peak, highlighted with “B”, differed by roughly 20m from the highest peak, marked with “A”. Peak “B” was also shifted horizontally by approx. 120m, i.e., 3 pixels.

To convert a DTM into a realistic DSM can be challenging since in real-world conditions the height variations with respect to the bare earth model are not purely random, but depend on the scenery, sensor technology (active or passive data capture), and point density. For sparse unclassified tie point clouds, adding white noise to convert the DTM into a DSM was a practical solution. Without inhibiting major limitations on the general validity of the results, a 10m white noise band – mimicking the height of a three-story house – was added to the DTM. The second column of the figure below displays the respective DSM, internally referred as “X5d3”.



The simulated X5d2 DTM and the X5d3 DSM with increased elevation of the second peak


The “Version 3” DEMs were used for much of the initial FMT testing and for some of the initial AFTM experiments. A further fine-tuned “Version 4” DTM and “Version 5” DSM, illustrated as DTM X5d4 and DSM X5d5 in Figure 3.13 below, were also created. Here, the elevation of “peak B” was increased by approximately 10m and the Gaussian bell was shifted horizontally by a small extend to further decrease mirror-symmetry. A large percentage of the FMT and AFMT experiments were conducted using these simulated datasets.



The simulated X5d4 DTM and the X5d5 DSM with further modified second peak

The figure above illustrates side and top views of the X5d4 DTM and X5d5 DSM that were used for most of the AFMT testing. Out of all simulated data sets, the X5d4 has the most pronounced version of peak “B”.


The figure below summarizes the elevation differences of the “Version 3” and “Version 4” DTMs.



Comparison of the X5d2 and the X5d4 DTM

As can be seen from the figure above, DTM X5d4's peak "C" is almost twice as high than the peak of DTM X5d2. It was also shifted by a few pixels horizontally, while all other portions of the surface model remained untouched.

The figure below shows an example of how the final image-rendered point cloud looked like. As will be discussed in Chapter 5.1.2.2.2, there were multiple ways to pad the point cloud-rendered image – in the example below, the image was padded with zeros.



/


The X5d4 DTM image used for co-registration testing

APPENDIX D: Fine Registration using unclassified Tie Points and unclassified Laser Points using ICP

Methodology

From Section 4.4.2 to Section 4.4.5, the suitability of the point-to-patch based fine registration concepts were examined. Rather than involving the extraction of the roof planes, or the curved surface patches, or relying on the point cloud classification, this section explores the performance of the point-to-point based fine registration using ICP.

As discussed in Chapter 2.0, ICP triangulates the point clouds, establishes a point-to-point relationship, and then minimizes the point-to-point distance while estimating the (rigid body) transformation parameters. The graph below gives a high-level outline of the method.



Point-to-point based co-registration of unclassified tie- and laser points using ICP

As can be seen from the schematic above, ICP does not require any pre-processing or point cloud filtering. However, ICP does require that the point clouds have already been coarsely aligned so that the algorithm can converge to the correct solution. As with the previous figures, the (*) indicates that the spatial discrepancies could be minimized by optimizing camera- or LiDAR sensor calibration parameters or by fitting parameters of the rigid body transform.


At first sight it seems desirable to couple the tie point cloud with the full-resolution laser point cloud (instead of a down-sampled LiDAR) as this should provide the best possible co-registration results. However, the co-registration of two point clouds can only be as accurate as the level of geometric accuracy and information content of the lesser-accurate, lesser-dense of the two point clouds allows. In other words, whether the laser points are used in full resolution or down-sampled slightly higher than the tie points, the level of accuracy of the co-registration results should be practically-identical. Also, since ICP involves point cloud triangulation and rendering and since geodetic laser point clouds typically comprise millions of points, the speed of ICP degrades quickly when loading larger point clouds.

While ICPatch (Al-Durgham,2014) and LS3D (Akca, 2007; Gruen and Akca, 2005) may surpass ICP in co-registration performance for some airborne geodetic applications (see Chapter 2.3.2 for more details), ICP is accessible through commercial and publicly available software, such as CloudCompare (CloudCompare, 2019). The following sections examine the performance of ICP, using CloudCompare software, for co-registration of LiDAR and extremely sparse photogrammetric point clouds.


Results

Over the course of this thesis work, many co-registration tests with ICP were conducted using various point clouds and different versions of CloudCompare (CloudCompare, 2019). These tests showed that the precision of co-registration results varies with the data set, processing settings, and level of the coarse alignment. The following screenshots illustrate this finding:

The figure below shows the overlay of the Oshawa laser point cloud before and after co-registration. The point cloud was down-sampled to cover the ground evenly with one point per 40m x 40m cell, resulting in a total of 2739 points. CloudCompare was deployed with default settings and the random sampling limit was left at 50,000 points. The point clouds were set apart by a 90 rotation and the transformation function was set to allow for rotation in XYZ, translation in Tx, Ty, Tz, and no change in scale; in a second run the transformation function was set to allow for only for rotation in Z and translation in Tx, Ty and no scale.



(a)



(b)

Rotated KN laser points before (a) and after (b) ICP self-registration


The figure above shows a down-sampled LiDAR point cloud of the Oshawa area, where (a) displays the point cloud overlaid with its 90degree rotated counterpart. The plot in (b) shows the point clouds after registration. Obviously, the point clouds didn't fully coincide.

The table below highlights the estimated co-registration parameters which converged with a final RMS of 18.734m: /


Self-registration results, KN laser point cloud

R_{XYZ}			T_{XYZ}
0.941	-0.339	0.000	-110.735
0.339	0.941	0.000	-357.960
0.000	0.000	1.000	-2.904

The experiment was repeated, first again with the KN laser point cloud and then with the Oshawa laser point cloud, however instead of rotating the search point cloud 90 degree around its centre of gravity, this time the search point cloud was rotated around the point with the smallest coordinates, i.e., its south-east corner. In addition to the 90degree angle, this caused a large shift in West direction. The following figure and tables illustrate the results of these experiments:



(a)



(b)

Rotated (a) KN and (b) Oshawa laser point cloud after ICP self-registration

As the cell (a) above shows, the registration results degraded further (in comparison to the results obtain earlier), yet ICP was still able to converge with a final RMS of 67.150m. In case of the Oshawa data, cell (b) above, the ICP failed to move the point cloud and converged from the 8775 points with a final RMS 208.057m. The two tables below summarize the results:

Self-registration results, KN laser point cloud

R_{XYZ}			T_{XYZ}
0.990	-0.143	0.000	1314.314
-0.097	0.990	0.000	1206.999
0.000	0.000	1.000	54.556


Self-registration results, Oshawa laser point cloud

R_{XYZ}			T_{XYZ}
0.995	0.097	0.000	-219.243
-0.097	0.995	0.000	277.468
0.000	0.000	1.000	1.165

Lastly, the down-sampled KN laser point cloud was scaled by factor 3. In this self-registration test, ICP failed completely, and no parameter estimation occurred. However, - without going into details - ICP was able to recover translations of several hundred meters, i.e., several pixels for both the KN and Oshawa point clouds.

Since ICP has been designed primarily for fine registration, it was expected that ICP would fail in most of these experiments. To analyze the ICP's performance for fine registration of sparse point clouds, it was deployed on two tie point clouds that described an approximate area of 0.5km by 0.7km in the centre of the Oshawa flight area. These areas were located right next to each other

and showed relatively high point density – like what one would obtain from an image-based tie point engine –, which made them ideal candidates for comparison. The following figure shows the geometric fit after co-registration.



Co-registration results of two neighboring tie point clouds using ICP

In the figure above, cell (a) and (c) show colour-coded distances between tie points and the LiDAR ground model (see APPENDIX C) of the neighboring blocks. Cell (b) and (d) display the

respective distance graphs. The graphs were obtained with CloudCompare version 2 (CloudCompare, 2019) and default settings: Octree level of 9 and multi-threaded processing enabled.

In view of the high vertical accuracy of tie- and laser points (see APPENDIX B), vertical deviations of 20-60cm seemed quite large. Even more surprising was the change in distribution of the distance values in the histograms. Considering (1) the 20cm-level accuracy that was achieved when establishing the reference solution with TerraPhoto (see Chapter 3.3 and APPENDIX C for details) and considering (2) that the survey blocks located right next to each other and (3) that the blocks depicted similar scenery, one would expect these histograms to be more alike. This anomaly was also observed with other data tiles of similar resolution and led to the conclusion that – as already noted by Akra (2007) and Al-Durgham (2014) – ICP is not the optimal fine registration algorithm for this sparsely-sampled data.

Discussion

The ICP technique is one of the most widely-spread methods for fine registration. The experiments above isolated and benchmarked some of the conditions at which the ICP co-registration using the extremely-sparse point clouds is possible. Contrary to the aforementioned fine registration methodologies, it does not require a training phase nor does it rely on the classification or other forms of the data pre-processing. CloudCompare software (CloudCompare, 2019) was used to co-register a down-sampled version of the Oshawa laser point cloud.


The tests showed that point clouds must be freed from major scalar and rotational offsets and that translations should be in the range of few pixels so that successful fine registration can occur. The tests also showed that - while ICP was able to successfully complete the fine registration in most test cases - the geometric fit was worse than the expected 20cm (see APPENDIX D) and several decimeters away from photogrammetric ground control. This confirmed earlier conclusions that the point cloud was spaced too sparsely for accurate fine-registration and that the point clouds captured different parts of the same objects on the surface than the GCPs. Other methods, like LS3D (Akra, 2007) which were explicitly tested for co-registration of dense point clouds against coarse ground control, perform better and are less stringent on coarse pre-alignment.

APPENDIX E: Design of Experiments, Reduction of Design Space, and Refined Experiments with Simulated Data

The focus of the ensuing sections will be on reducing the design space, narrowing down choices for the number of correlation peaks, image cropping, window- and HF filtering, and the importance of undulation in the elevation profile from comparing the FMT performance of simulated DTM and DSM data.

Design of Experiments

As indicated in Chapter 5.1.2.1 and Figure 5.14, the standard FMT contains many parameters that can be adjusted. It is important to understand the function of each parameter and be aware of potential correlations, so that the FMT and its co-registration performance can be clearly characterized. An overview of all FMT parameters, arranged in form of a Pugh diagram (Kacker *et al.*, 1991), is shown in the figure below.



Pugh diagram of the standard FMT

The Pugh diagram gives an overview of the variables that can be adjusted in the standard FMT:


- Primary input to the FMT are height maps in image format, displayed in green on the left. Since FMT co-registration keeps one of the two input images static, as reference, while the other “floating” image is transformed onto the reference image, there are only two possible data combinations per survey site: In case of the real-world data, these are (1) static laser points with floating tie points, or (2) static tie points with floating laser points. During performance evaluation however, two additional data combinations can occur: (3) static laser points with moving laser points, or (4) static tie points with moving tie points.
- The output of the FMT, displayed in purple on the right, are the co-registration parameters. Most of the tests described in the ensuing paragraphs focused on the recovery of rotation and scale only. Therefore, the observation of translation offsets was greyed out.
- Encircled in orange are all FMT-internal parameters that can be adjusted. Displayed in grey are control parameters that were set to a constant value or reduced to a reduced sub-selection in the tests described in the ensuing discussions.
- The group of variables in the blue box describe environmental factors, or noise factors (Kacker *et al.*, 1991), that are typically out of the user’s control. In the context of co-registration, these

are the unknown co-registration parameters we are trying to recover. As before, the parameter values that were tested less frequently than others in the subsequent experiments were marked in grey.

- The “Symptoms” describe output attributes that can only be observed but cannot be controlled. Examples are robustness, SNR, performance characteristics like convergency behaviour (predicted vs actual, stability, artifacts related to the outcome like how many correct solution candidates stem from a certain amount of CPS peaks), and the cost of computational time.

Related to the symptoms is the recording, evaluation, and feedback of FMT results: From conducting experiments with different data input- and processing settings, the range of correctly solved scale-, translation-, and rotational offsets can be identified and looped back. Wilmer (2012) for example found that his FMT implementation could resolve a maximum scale factor of 5.0 when applied to images that show close-range scenery with people in them.

As the Pugh diagram shows, there are many noise factors and control factors that need to be considered when analyzing the FMT’s co-registration performance. This raises the question how these factors should be conditioned to explore the test space, or “design” space, efficiently and effectively. The fishbone diagram below relates the noise and control parameters and helps answer this question.



Fishbone diagram illustrating parameter combinations of the standard FMT

As the fishbone diagram above shows, the standard FMT covers a deep design space. Without assumptions and parameter restrictions (values in grey color), the design space would be even larger: Rotation angles could for example be selected in increments of 10degrees (rather than 90), and the HF and image cropping (see next section), and multiple methods of image padding could additionally be accounted for, too.


The following strategies for interrogating the design space are available:

1. Test all options (full factorial testing)
2. Run select options within the design space to create a “math model” for how the factors affect the response (response surface modeling, kriging)
3. Use optimization techniques (e.g., method of steepest ascent, genetic algorithm,)
4. Fractional factorial experimentation (Orthogonal Arrays (OA))

Full factorial testing (method 1) is required when it is expected that some of the parameters are correlated. In case of FMT testing, method 2 and 3 would infer unnecessary complexity. Since the

use of OA (method 4) promises a reduction in test cases (while isolating processing settings which achieve the best outcome), OAs will be explored further in the ensuing sections.

According to Kacker *et al.* (1991), the columns in an OA correspond to the noise factors (see Figure 5.20) and the entries in the columns correspond to the test levels of the factors, and the rows correspond to the test runs, i.e., the rows of an OA by $N(s^m)$ are a subset of the possible s^m test runs of a complete factorial plan in m factors each having s test levels. The figure shows an example of an 18-by-8 OA for the FMT.




N18-8 OA and respective fishbone diagram, with one variable (window filter) at two levels and two variables at three levels

The upper half in the above figure shows an 18-by-8 OA. Here, each column represents one set of noise conditions and each of the 18 rows represents a combination of control parameters that pairs with the 8 noise conditions. The table in the lower left lists the three control factors and their levels that go into the OA. The fishbone diagram on the right shows a visual presentation of the respective control and noise factors. This OA can be used to systematically test all combinations that are covered by this fishbone diagram.

The above OA greatly helps reduce the test combinations. As indicated in N18-OA, for each dataset there are 4 ways to combine the data itself (tie and laser point cloud), 3 different noise

factors (rotation, scale, translation) at 4, 3, and 2*5 levels, respectively, and each noise factor could be tested with 3 control factors at many levels (3, 6, and 8 levels), resulting in a total of $4 * (4+3+2*5) * (3+6+8) = 4*3*17*17 = 1156$ possible test combinations. In comparison, to only test the factors highlighted above in red colour, full factorial testing entails only $4*6*8 = 96$ tests.

OAs can also be used to design test cases for sensitivity analysis (SA). The OA illustrated below confirmed the systematic study the impact of the number for correlation peaks on the co-registration results.



L18-8* OA and respective fishbone diagram, with one variable (correlation peaks) at six levels and two variables at three levels

While OAs are a powerful tool to analyze FMT performance, they could not solve all challenges of the FMT test setup:

- (1) The real-world datasets contain a small set of datapoints; executing one FMT run with real-world data takes few minutes. The simulated datasets however contain many data points and running one test takes many hours. So even executing a reduced set of 96 tests would be very time consuming. If some of the test combinations could be eliminated with even less testing than implied by the OA above, it would be better.

- (2) OA assumes that all parameters in the design space are orthogonal, i.e., that there is no cross-correlation of control parameters. That however may not be applicable with the low SNR geodetic data on hand.

However, the design space can be reduced further by assessing the characteristics of control parameters in the context of co-registration and critically evaluating their impact on the test result. The following sections focus on isolating some of the control parameters that have very strong impact on the co-registration results and finding their optimal setting first, before evaluating and fine-tuning other parameters through an OA-like scheme.

Reduction of Design Space

As aforementioned, for co-registration of geospatial height maps, the full FMT parameter space would entail the large space indicated by the fishbone diagram above, and, additionally, it would accommodate different settings for the HF, image cropping, and characteristics of the simulated elevation profiles. Obviously, testing all combinations would be very painstaking.

Using the Toronto tie and laser points, initial tests were conducted where the number of evaluated correlation peaks, interpolation method, image filter, rotation angle, and scale were arbitrarily changed, which seemed particularly important factors. The goal of these experiments was to get a coarse understanding of the impact of each factor and/or perhaps isolate some parameter configurations that are critical for successful co-registration. The point clouds were auto-correlated and cross-correlated using Wilmer's FMT implementation (Wilmer, 2012).

While these tests gave a first impression what impact the FMT-internal parameters have on the co-registration results, less than 1/3 of the input noise parameters were recovered correctly during this initial test phase. *Indeed, the results did not allow to make a clear distinction between the impact of data characteristics and the impact of FMT-internal parameter settings.* These findings led to

1. Generating the simulated datasets introduced in Chapter 3.4, which allowed for systematic modification of size, resolution, elevation, and noise level of the input data. Having


simulated datasets allowed for testing more data combinations, evaluating the impact of the elevation profile, and cross-comparing co-registration results of higher resolution surface data.

2. Creating a test configuration where the simulated search point cloud were rotated systematically in 10deg increments with respect to each other (starting with 0deg, then 10deg, then 20deg, etc. up to 350deg) and the input angles were compared to the FMT-recovered angles. By additionally modifying one other parameter at a time, this setup was used to study the impact of (a) number of correlation peaks, (b) window mask, (c) HF, (d) interpolation method, (e) elevation signature of the point clouds.
3. Monitoring the co-registration results using real-world data whilst applying cardinal rotations and systematically altering one other FMT parameter at a time.

This second test phase revealed that it is especially important to decide upon the number of correlation peaks considered in one FMT run and image cropping, since these parameters showed strong impact on co-registration results. For studying the impact of image cropping, and the number of included correlation peaks, as well as the effect of window masks and interpolation methods, Wilmer's MATLAB implementation (Wilmer, 2012) was used. The respective findings will be presented in the following sections.

Number of Correlation Peaks evaluated in one FMT Run

As highlighted in Chapter 5.1.2.2.2, the CPS results one dominant peak (that corresponds to the correct set of co-registration parameters) if the point clouds allow for strong cross-correlation. This is typically the case when the height maps show strong change in magnitude, as for example surface scans of mountainous terrain. More, how the datasets are set apart determines the shape of the CPS graph, too. A cardinal offset of 180degrees for example yielded a more pronounced first CPS peak than other cardinal angles and non-cardinal offsets. The figure below illustrates an example:



Self-registration results of KN laser point cloud with a 180deg offset (left) and 90deg offset (right)

The left column of the figure above shows an example where the KN laser point cloud (which covers 69x81 pixels) was offset by 180deg. Here the CPS showed one dominant peak which corresponded to the correct set of co-registration parameters (a 180deg rotation offset and no offsets in scale and translation). Table 8.7 below lists the final set of co-registration parameters, too, for further analysis.

The right column shows co-registration parameters where the same point cloud was offset by 90deg. A couple of larger CPS peaks occurred and instead of peak 8 or 14, which represented the correct solution, the FMT settled for peak 20 which corresponded to a rotation of 271.4deg, a scale of 0.9332, and a translation of 1 and 11 in x and y, respectively. (This result was also recorded in row A6,2 of Table 8.7).

As this example showed, coarse geodetic elevation profiles may not always have enough signal strength to result in one isolated CPS peak, but rather can the CPS graph contain peaks of similar magnitude where each of these peaks reflects a potentially correct solution. Including more correlation peaks in the search for the correct set of co-registration parameters increases however computation efforts. Consequently, the following question had to be raised: How many CPS peaks must be evaluated in one FMT run to ensure the correct solution can be detected?

To answer this question, tests with real-world and simulated data were repeated whilst incorporating more correlation peaks with each run. Each time the number of considered CPS peaks was increased (first 10, then 20, then 30, and finally 40), the final co-registration estimates were compared. By doing so, a few rotation angles crystalized that always solved correctly (this observations will be investigated further in Chapter 5.1.2.2.4), and other rotation angles isolated that the FMT could not solve properly, even when 40 correlation peaks were considered.

In cases where the datasets where offset by multiples of 45 degrees ($\pi/4$), causing two “+”, it proofed sufficient to include only the 20 largest CPS peaks in the search of the correction set of co-registration parameters. Even when the point clouds were set apart by more challenging angles (that create three “+”), increasing the number of peaks from 20 to 30 helped recover the correct rotation angle only in two out of roughly 40 test cases, i.e., less than 5% of the tests.

When 20 CPS peaks were included in the search for the correct solution, one FMT run required a few seconds computation time when real-world data was used and 6-8hrs when simulated data was used. Consequently, generating a set of 10deg-incremented FMT solutions with simulated


data, required multiple days computation time. This was a practical indicator that 20 correlation peaks should be the maximum setting when evaluating DSM data.

The summary of these observations led to the conclusion that the number of CPS peaks considered in one FMT run should be fixed at 20. Thus, in all consecutive FMT and AFTM tests, only the 20 tallest CPS peaks were evaluated.

Effects of Image Cropping

The FMT co-registration algorithm can be injected with point clouds from different areas, with different resolution, different size, and perhaps expanded with a different padding method (see Chapter 5.1.2.2.2 for details). When the search image is larger in size than the template image, image cropping could be applied to cut off image excess. An example is the co-registration of an image and its 30deg rotated, non-scaled counterpart. Here, the rotated search image shows an increased footprint (compared to the reference frame) and cropping takes away scenery of the outstanding corners, so that both images end up having the same size.

The effect of image cropping was examined with both real-world data and simulated DSM data with Wilmer's FMT implementation (Wilmer, 2012). When the scale was set to 1.0, cropping showed no significant effects on the co-registration results, most likely because for the given datasets the areas with the largest change in elevation were located in the image center, and thus they remained unaffected by the trimming of the image corners. However, degraded co-registration results were observed in some instances with simulated data, and more frequently with real-world data, when scaling was applied. The following figure shows an example:




FMT auto-registered X5d4 DTM, scaled by factor 3, without (left) and with (right) image cropping

The figure above illustrates two sets of FMT parameter curves where the X5d4 DTM was autocorrelated and the scale factor was altered to 3.0 whilst the other registration parameters were left unchanged. When cropping was applied (right column), two of the CPS peaks resulted in two scale estimates of approx. 2.8 and 2.9, respectively. However, a closer look reveals that the two CPS peaks that retrieved the scale correctly paired with an incorrect rotation of 180deg and small translations of 7 and 15 pixels in X and Y, respectively. Without cropping (left column), only one of the 20 CPS peaks corresponded in a scale estimate near 3.0 (2.972). Here, the single CPS peak that re-produced the correct scale also corresponded to a correct rotation estimate of 0deg and translations of 0,0 in X and Y. The CPS peaks gained without cropping showed stronger magnitudes, too, which confirmed stronger correlation of the template and search point cloud.

While the above illustration showed an example where image cropping clearly degraded the FMT results, other tests showed only minor impact, especially when simulated DEM data was used. Generally, however, one can conclude: Since image cropping reduces data content, it should be avoided when co-registering geodetic data. *Instead of cropping the bigger image, padding, and resizing both input images to a common, larger footprint should be the preferred method.*

Impact of Window Filter

As mentioned in section 5.1.2.2.2, window filters can be used to blur the corners of the input imagery in space domain in order to minimize the “+” effect in the Fourier domain. To better understand the impact of different windows (and which window may lead to the best co-registration results), The figure below illustrates the visual effect of nine window filters:




X5d4 DTM after application of window masks

In the illustration above, the window filtered DTM images were sorted by the damping power of the window mask. Here, the filter strength increases from left to right and from top to bottom. Whilst the damping effect of most masks is clearly visible above, the differences between

Hamming and Hann window were very subtle and require a closer look; Also, in Wilmer's implementation (Wilmer, 2012) the visual effects of the Boxcar and MATLAB's Chebwin Chebyshev window filter (MathWorks, 2019), seemed to be identical. Out of these nine masks, the Hamming window showed the least amount of signal absorption and the MathWorks's "Chebwin" Chebyshev filter, showed the strongest damping effect. Thus, the Hamming and Chebyshev filter were used frequently in subsequent tests that examine the effect of window filters on co-registration results.


The effects of the Chebyshev window can be seen in the figure below. Here, the left side shows co-registration results where the search and template point clouds, both X5d5 DSM, were set apart only by rotation and the rotation angle was altered in 10deg increments. The right side shows results where the experiment was repeated whilst applying the Chebyshev mask. Both tests were executed with bicubic interpolation.




FMT derived registration parameters of the X5d5 DSM without (left) and with (right) application of a Chebyshev mask

As illustrated above, the registration parameters were estimated correctly in fourteen (thirteen) instances at 90deg, 180deg, 270deg +/- 10deg without (with) application of the Chebyshev filter. The scale and translation estimates showed a slightly more stability and symmetry amongst each other when the window mask was not applied.

The effect of MATLAB's "Chebwin" mask was analyzed further for the X5d4 DTM with a 20deg, 30deg, 60deg, and 80deg angular offset, i.e., angles that are difficult to recover since they yield a triple "+" that shadows the signature of the actual image content in the phase spectrum. The figure below illustrates co-registration results from a test where the X5d4 DTM was offset by a 60deg rotation, without application (left column) and with application (right column) of a Chebyshev filter. When this experiment was conducted for other non-cardinal angles, similar results were obtained. In the following, only the results gained with the 60deg and 20deg offset will be displayed as universal representation of non-cardinal angle cases. As before, the results shown in the graph below were obtained with bicubic interpolation, considering 20 CPS-peaks, and without image-cropping:






FMT self-registration results of 60deg rotated X5d4 DTM without (left) and with (right) application of a Chebyshev filter


In the graphs above, the left column displays results where no window mask was applied. The FMT identified the seventh-largest CPS peak correctly as solution and peak 13 proofed to be another correct solution candidate. An angular offset of 60.5degree, a scale of 1, and 25 pixels

offset in X and 43 pixels offset in Y were estimated. Considering the simulated DTM extends over a thousand pixels in each direction, 25- and 43-pixels translation can be considered a negligible, noise-like deviation from (0,0). The right column in the figure above shows registration results where the experiment was repeated whilst applying the Chebyshev filter. Peak ten was identified correctly as the only correct solution candidate, yielding a 60.0degree rotational offset, a scale of 1.0, and translational offsets of 2 pixels in X and Y.

As the figure above also shows, without window filter the scale was estimated at 1.0 twelve times yet only six times when the mask applied, so the window filter did degrade the FMT results slightly. However, since the variation of the scale estimates was small in magnitude (-20% to +5%), one can argue the scale was solved correctly for all 20 peaks in both cases.

The following figure displays self-registration results of a 20deg offset X5d4 DTM, i.e., another angle that is difficult to recover since it yields a triple “+” in the phase spectrum. As before, bicubic interpolation was deployed.





Co-registration results of 20deg rotated X5D4 DTM without (left) and with (right) application of Chebyshev mask

As before with the 60deg offset (previous figure), the resulting variation in scale was arguably small ($\pm 20\%$). Yet in this instance, the Chebyshev filter seemed to stabilize the scale estimates

slightly. The rotation estimates showed a lot of variation from peak to peak, with and without window mask, and overall, the solution appeared highly unstable.

In summary, out of the nine window masks that were introduced, the Chebyshev filter had the strongest dampening effect. The impact of the Chebyshev filter varied in the experiments and, overall, it seemed to degrade the registration results slightly. This preliminary observation was affirmed in Chapter 6.1 where the impact of window filters on the co-registration performance is analyzed with real-world data.

Impact of High-Pass Filter

As mentioned earlier in Chapter 5.1.2.2, in the FMT, the magnitude spectra is passed through an HF prior to sampling it into log-polar base to remove low frequency components before the second POMF inside the FMT. Obviously, the shape of the HF inside the first POMF is another degree of freedom that needs to be considered. During the early stage of performance testing, many of the earlier described experiments were repeated whilst gradually changing the shape of the HF, by increasing and decreasing factors 1.0 and 2.0 in equations (5.14) and (5.15). For analyzing the effects of the HF shape, an extended version of Laundon's MATLAB FMT implementation (Laundon, 2008) was deployed. Without going into details, the shape gained with default factors 1.0 and 2.0, as shown in cells C1,1 and C2,1 of Figure 5.15, and as suggest by Kazik and Göktoğan (2011) and Reddy and Chatterji (1996), gave the best co-registration results. Thus, this setting was used in all consecutive FMT tests.

Discussion

Through experimenting in a trial-and-error approach using real-world data and Wilmer's FMT implementation (Wilmer, 2012), first insights were gained into which control parameters have a strong impact on the co-registration results. It turned out that the number of correlation peaks considered in the search for the best co-registration solution is a very important factor for FMT

success. The experiments showed that evaluating 20 CPS peaks in one FMT run is an optimal setting. The experiments also showed that image cropping should be avoided.


Nine window masks and their effects were introduced. It was found that the Chebyshev window has the strongest damping power and that neither applying this window, nor applying no window at all, could consistently improve the co-registration results. Perhaps using a filter with low damping power, like the Bartlett window, would be a good compromise and benefit the co-registration results the most. Lastly, the impact of the HF was studied. Using the shape proposed by Reddy and Chatterji (1996) turned out the best setting for the HF.

Refined Experiments with Simulated Data

Based on the learnings from the initial test phase that focused on the impact of window masks, HF, and number of correlation peaks, another series of experiments was conducted to study the impact of magnitude and noise of the elevation profiles. For these experiments, simulated DEM data was used since it came with a known noise level. The point cloud images were left unpadding and co-registered in auto-correlation mode, and the point clouds were set apart only by rotation (whilst keeping the scale and translation unchanged). The rotation angle was altered in 10deg increments. To study the impact of the parameters of interest, the test series was repeated whilst changing only that respective parameter.

Impact of Elevation Dynamics

To analyze the importance of changes in the elevation magnitude, the X5d4 DTM and X5d1 DTM (which, as described in Chapter 3.4, has a slightly less pronounced peak “C” than X5d4) were self-registered. The figure below illustrates the autocorrelation results of the X5d4 DTM (left) and the X5d1 DTM (right):




FMT-derived co-registration parameters of self-registered X5d4 (left) and X5d1 DTM (right)

As can be seen from the figure above, the X5d4 DTM showed slightly better co-registration results than X5d1: All angular increments and scale offsets were retrieved correctly. The errors in translation were small and repeated with $\pi/2$ for the X5d4 DTM. The translational offsets in X

and Y were identical for all angular increments - the lines plotted in green and purple above were right on top of each other. The translation estimates showed occasionally abrupt changes, within small range, from one 10deg increment to the next, which – as discussed further in Chapter 6.1 – can be caused by effects of image rendering and is therefore unavoidable and acceptable.

Like X5d4, the self-registration of X5d1 showed accurate rotation and scale estimates. The translational estimates were small and repeated with $\pi/2$, too, however they were not quite as symmetrical, and contained a few outliers, which can be attributed to X5d1's smaller differences in elevation.

To study the impact of differences in the elevation profile, both the X5d4 and the X5d1 DTM were self-registered with a scale difference of 3.0 without rotation and without cropping. The plot below presents a graphical comparison of the outcome:




Self-registration results of X5d4 DTM (left) and X5d1 DTM (right), scaled by factor 3

The figure above shows two sets of FMT parameter curves where the X5d4 DTM and X5d1 DTM were auto-correlated and the scale factor was altered to 3.0 whilst the other registration parameters were left unchanged. Both tests produced only one CPS peak which yielded a correct scale of 3 (2.972) and a rotation offset of 0deg. The translations were estimated at 7,15 pixels and 15,7 pixels in X and Y, respectively. While X5d4 resulted 13 correct rotation estimates, there were only 11 instances with correct rotation estimates for X5d1.

As this experiment showed, an increased elevation contrast is critical for the stability of the co-registration results. This observation is affirmed in Chapter 5.1.2.2.4, where test results from real-world data are discussed and more insights into the importance of the elevation profile will be shared.

Impact of Noise in the Elevation Map

To study the FTM's sensitivity to noise, the experiments above were repeated using the X5d5 DSM, which (as described in Chapter 3.4) is the X5d4 DTM plus a 10m noise band. Final FMT parameters for the X5d5 map were already shown above, and the following experiment shows the FMT-internal results for one rotational offset. Here, a 30deg rotational offset, which is particularly difficult to solve as it yields a triple “+” in Fourier domain, was selected. The results are displayed below.



FMT parameters of 30deg rotated X5d4 DTM (left) and X5d5 DSM (right)


The figure above displays self-registration results where the simulated DTM and DSM were offset by a 30 deg angle. In both tests a Chebyshev window was applied. In case of the X5d5 DSM, the CPS peak with the ninth highest magnitude, corresponding to a 100deg offset, was identified incorrectly as solution (right column above) even though the first CPS peak would have provided the correct set of parameters. The graphs shown in the left column were obtained using the X5d4 data. Here, the FMT identified the fifth largest CPS peak correctly as solution. When the experiment was repeated with other window masks, the results remained identical.

As can also be observed, the registration parameters were represented correctly in four instances for the X5d4 data and but only by one peak for the X5d5 data. Obviously, the scale and translation estimates of the 20 CPS largest CPS peaks showed more consistency and stability when the elevation signal was free of noise. This finding – that the FMT is very sensitive to noise in the elevation signal – will be re-affirmed in Chapter 6.1, where similar tests with DSM data will be preformed.

Effects of Image Rendering

As described in Chapter 5.1.2.2, rotation offsets that are uneven multiples of $\pi/4$ create two “+”, and rotation offsets that are even multiples of $\pi/4$ create one “+” effect in the phase spectrum. The experiments described thus far showed that the FMT could recover such angular offsets when simulated data was used. The experiments with simulated data also revealed that for angular offsets that are 10deg away from even multiples of $\pi/4$, i.e., the triplets {350deg, 0deg, 10deg} and {80deg, 90deg, 100deg}, the FMT solved the rotation and scale offsets correctly and showed small translation estimates, too.

This observation led to a closer investigation of these areas with a resolution of 5deg increments. Here, the experiment described above (which compared co-registration values of the X5d5 and X5d1 DEMs) was repeated for the first quadrant using 5deg angular increments, instead of 10deg increments. The results are plotted below:





FMT-derived rotations of the simulated X5d5 DSM dataset using 10deg (left) and 5deg (right) rotational increments

The left side in the figure above is identical to what was shown few figures earlier, where the X5d5 DSM was auto-correlated using a Chebyshev mask. The right side shows the results of the first quadrant at a higher resolution. As can be seen, registration parameters related to any CPS

peak and its neighboring (next biggest) CPS peak did not form a linear trend: Whilst the scale values fluctuated little, the rotation and translation estimates showed drastic changes from one CPS peak to the next, for both 5degree and 10degree rotation increments. At 5deg increment resolution, the FMT self-registration resulted in two additional correct parameter sets in the first quadrant - at 45deg and 55deg rotation. At 40deg, 50deg, and 60deg rotation, the FMT failed however (both with and without application of a window filter). The graphs above also showed cross-correlation between the co-registration parameters: Wherever the translation estimates met their expected value, the respective rotation and scale estimates converged correctly, too.

Similar observations were made when the above experiment was repeated, whilst again applying MATLAB's Chebwin mask (MathWorks, 2019) and replacing the X5d1 dataset with X5d2. Since the X5d2 DTM had a slightly smoother elevation signature (see Chapter 3.4) it was expected that is caused slightly less leakage and less “+” effect. Since the results of the previous experiment showed periodicity with $\pi/2$, this test was also conducted for the first quadrant only, using 5deg rotation increments. The rotation estimates turned out identical and the scale estimates showed similar variations as before. Of special interest were the translation estimates – they showed small variations that followed symmetric patterns, as shown below.





FMT-derived translations of self-registered DTMs X5d4 (left) and X5d2 (right) without and with application of Chebyshev mask

As can be seen from the figure above, the translation estimates for the X5d4 DTM (left column) were identical for all 10deg increments to the values of the previous experiment (left column in previous figure) within the first quadrant. Since the x- and y- component of the estimates were always identical, the plots appeared in a single color. Obviously, the translation estimates that always solved correctly were for rotations of 0 and 90 degrees, and 45deg. Here, only one and two “+” occur, respectively. More importantly, for the angles the border lines between padded areas and the elevation map in the search and template image have the same length everywhere and form triangles with two equal sides. However, for rotation angles that cause three “+”, these border lines vary in length and the padded areas form non-symmetric triangles.

A closer look at the MATLAB-rotated, zero-padded plots revealed that MATLAB’s “imrotate” function caused indeed small shifts of the height maps that varied with the applied rotation angle, i.e., the exact location of the image centre and the number of zero-padded pixels on the left and right side of the DTM plot altered by a few pixel from one rotation increment to the next. These small shifts happened whenever rotations with more than two respective “+” occurred. It seemed that the tooth-saw (zigzag) effect of the pixel rendering at the image borders appeared more pronounced for some rotation angles than for others and contributed to these small translational artifacts. Thus, depending on the rotation angle, MATLAB’s image rendering itself added uncertainty and inaccuracy to the co-registration results and the translation (and respective rotation

and scale) estimates were truly supposed to change by a small quantity when the DEM images were rotated from one 5degree increment to the next.

Convergency Behavior

Since the simulated and real-world airborne laser and tie point clouds show relatively small changes in magnitude, the magnitudes of their Fourier phase spectra are small, too. Thus, when put in auto-correlation, the energy of the CPS tends to spread across multiple peaks of similar magnitude. Moreover, for angular offsets that are not a multiple of $\pi/4$ this phenomenon is enhanced by the “+” effect.


During FMT analysis, many CPS graphs were generated. The analysis of these plots revealed that the CPS graphs can be categorized in three types:

- (1) The CPS displays a dominant first peak followed by significantly smaller peaks
- (2) The CPS shows multiple dominant peaks of similar magnitude (typically located along the axis representing rotation), i.e., with only subtle decrease in magnitude from one peak to the next
- (3) The CPS graph is filled with many smaller peaks and looks very noisy


These observations lead to the presumption that a clear dominant first peak followed by a clear falloff could perhaps serve as indicator that the first peak corresponds to the correct solution, or that a certain magnitude or ratio of a peak and its neighboring peak could serve as an indicator that this peak represents the correct set of registration parameters. More generally, it was expected that some sort of reliability indicator could be derived from the CPS graph and/or the graphs of the co-registration parameters that correspond to the 20 tallest CPS peaks.

To study the significance of the magnitude and ratio of the CPS peaks (as potential indicator of a correct solution), it turned out helpful to illustrate the magnitude of the first n largest peaks (where n is a user selected number) in the type of 2D graph. Here, the peak curve represents a non-logarithmic equivalent of the twenty largest CPS peaks. This graph proved to simplify the analysis especially in cases where the CPS graph tends to be noisy and difficult to interpret, as for example for a 30deg, 60deg, 80deg rotation (and multiples thereof). The graph below shows some examples.

Below: CPS of self-registered X5d2 with 30deg offset




(a1)




(a2)

Below: CPS of self-registered X5d2 with 45deg offset




(b1)




(b2)

Below: CPS of self-registered X5d4 with 60deg offset



(c1)



(c2)

FMT-CPS peaks of X5D2 dataset for different rotation angles

As can be seen from cell (a) above, in case of 30deg rotation where the padding areas caused a triple “+” effect, the CPS energy scattered across many peaks, and the fall-off in magnitude from one CPS peak to the next was small and gradual. In case of the 45deg rotation, cell (b), the CPS graph was dominated by the impact of only two “+” lobes and the peak plot produced four peaks with strong magnitude. The graph in cell (c) looked similar to that in cell (a) yet resulted in slightly more noise and slightly steeper decline in magnitude of the peaks.

Further examination of CPS curves showed that – against expectations – a strong fall-off in CPS magnitude from one peak to the next peak does not imply that the dominant peak corresponds to the correct solution. In fact, neither the shape of the CPS graph nor the magnitude of the peaks allowed to make conclusions about what peak corresponds to the correct solution. In the previous figure for example, it was peak eleven in (a2), peak one in (b2), and peak eight in (c2) that corresponded to the correct results.

Likewise, the experiments showed no clear correlation between the shape of the parameter curves (that correspond to the CPS graph) and which parameter value resembles the correct one. The two last figures above, for example, showed strong variation from one parameter estimate to the next, yet settled on the only estimate that reflected the correct solution. In the same way, a stable series of rotation and scale estimates that solved correctly for the majority of the 20 CPS peaks turned out to not be an indicator that the FMT converged successfully. This phenomenon can also be observed from the figure above where all scale estimates, except for one varied around 1.0, yet only one out of the 20 estimates reflected the correct, final solution.

APPENDIX F: Interpretation of FMT Experiments with Real-World Data

The experiments described in APPENDIX E showed that the FMT retrieved rotational offsets correctly for multiples of $\pi/2$ for the simulated DSM and DTM. Applying the Chebyshev mask caused no significant difference in the results. Compared to the self-registration of the simulated DSMs, the self-registration of the DTMs showed better results. In both cases, the parameter curves showed slightly more symmetric patterns when no window mask was applied, probably because more of the original signal came through to outweigh the degrading effect of the “+” lobes.


The tests in APPENDIX E also exposed that the FMT is sensitive to noise in the elevation map and that the auto-registration of the simulated DSM is difficult especially for angles that are not a multiple of $\pi/4$. Here, the padding areas have sides of different length which cause the CPS energy to spread across multiple “+” lobes, which – as explained in Chapter 5.1.2.2 – can overpower the frequency signature of the actual elevation signal of the DSM.

While in APPENDIX E the focus was on reducing the design space and narrowing down choices for the number of correlation peaks, image cropping, window- and HF filtering, in the next section, the effects of interpolation and presence of scale offset will be explored further using real-world data.

The results presented in Chapter 5.1.2.2.4 provided insight how some of the FMT-internal factors influence the FMT output:

- 1) Out of the three interpolation methods – (1) nearest neighbour, (2) bilinear, and (3) cubic – the cubic interpolation should, in theory, deliver the best results, since it maintains most of the signal information. However, changing the interpolation method showed very little impact on the co-registration results.
- 2) The FMT, while robust to the interpolation methods, proved quite susceptible to noise parameters: It performed correctly more often when no angular offset was applied compared to when 90deg, 180deg, and 270deg offsets were applied.

- 3) The best co-registration results were achieved in absence of scale:
- The FMT almost always failed when the search point cloud was scaled by 0.3048. This configuration mimicked the case where the search point cloud is referenced in feet whilst the target point cloud is referenced in meters. Here the search point cloud covered the same area on the ground with 3x3 times as many pixels than the target point cloud. To recreate this, the search point cloud was down sampled by factor 0.3048. As can be seen from below, this led to a significant loss of surface details and the CPS ended up correlating a search image the same size than the target image but with reduced information content. This loss of information is magnified inside the FMT where according to equation (5.18) the scale appears in log-polar representation.




KN laser data in (a) original resolution and (b) scaled by 0.3048

As the figure above shows, the pixel size increased when the point cloud image was scaled by 0.3408. In Wilmer's FMT implementation (Wilmer, 2012), the pixel size was increased by $1/0.3408$ and the number of pixels was reduced by factor 0.3408 so that the footprint remained the same as that of the original.

- When the search point cloud was scaled by factor 3.2808, the FMT solved correctly less frequently than without scale, but clearly more often than with scale 0.3048. As can be seen from below, more information content was maintained at scale 3.2808,

and the search image appeared as an interpolated, smoothed version of the target image with approx. 3x3 times as many pixels.



KN laser data in (a) original resolution and (b) scaled by 3.2808

As the figure above shows, most information of the original point cloud image was maintained when it was scaled by 3.2808. The pixel size was reduced by $1/3.2808$ and the number of pixels was increased by factor 3.2808 to maintain the footprint of the original point cloud image. To create this rendering, Wilmer's FMT implementation (Wilmer, 2012) was used, and cubic interpolation was chosen as sampling method.

- 4) The FMT performed better when the point clouds were set apart by either rotation or scale, and it performed worse, when the point clouds were offset by both rotation and scale at the same time. This suggests that there may be some cross-correlation of scale and to rotation inside the FMT. As discussed in 5.1.2.2.3, this correlation could be an indicator that the point clouds did not provide enough elevation difference for the FMT to fully discern between rotation- and scale offsets.
- 5) Table 5.2 to Table 5.7 also showed that use of Bartlett and Chebyshev filter created better results than FTM runs without filter. While the Chebyshev mask was only tested for

comparison with the Bartlett filter in one test configuration, the data from these tests showed that there was only a small difference in outcome, and that the Bartlett filter seemed to slightly outperform the use of the Chebyshev mask.

In summary, from all these findings gained with real-world data, one can conclude that the FMT performed best without rotation and without scale offset whilst the Bartlett filter was applied. As well, the co-registration failed in all test combinations when the scalar offset was 0.3048 and almost always when the offset in scale was 3.2808. The test combinations also showed that the interpolation method has little effect on the co-registration outcome. Overall, the Oshawa data resulted in better self-registration outcome than the KN data, most likely because the Oshawa laser point cloud shows stronger differences in elevation magnitude and more asymmetric distribution of elevation changes than the Keyworth-Nottingham point clouds.





APPENDIX G: Precise Alignment of Point Clouds to build a Reference Data





This section illustrates, by the example of two images, how the alignment of the datasets was checked and what levels of accuracies were achieved.



Example: Oshawa data, Image 9

The table below illustrates effects of laser point thinning, tie point locations in the image and in the laser data. It also demonstrates the tight fit of the point clouds for a few selected scenes of image 9 of the Oshawa collect.




	
Ground control distribution	Ground control and tie point distribution


	
LiDAR footprint of image 9 colored by line and intensity. A total of 6,853,812 points	LiDAR footprint of image 9 colored by intensity. Full resolution of LiDAR points resulting in a point density between $2/\text{m}^2$ to $21/\text{m}^2$
	
LiDAR points reduced to every 4th point resulting in a total of 1,713,453 points and a point density between $0.5/\text{m}^2$ to $5/\text{m}^2$	LiDAR points reduced to every 16th point resulting in a total of 428,363 points and a point density between $0.2/\text{m}^2$ to $1.5/\text{m}^2$

	
Control point 261 and tie points 30083407, 30083412, 30083416, 30084210, 30084205, 30084253, 30084199	Area around control point 261 in a las-viewer, colored by line and intensity. Average point density using the full resolution data is approx. $10/\text{m}^2$ to $26/\text{m}^2$
	
Area around control point 261. The average point density using every 4th point is approx. $2.9/\text{m}^2$ to $7/\text{m}^2$	Area around control point 261. Average point density using every 16th point is approx. $0.3/\text{m}^2$ to $1.6/\text{m}^2$

	
Control point 261 and tie points line up horizontally with the laser point cloud within <10cm (<1GSD)	Vertical sections line up vertically with the laser point cloud within <10cm (<1GSD)
	
Control point 37 and tie points 30299987, 30300191, 30260786	Area around control point 37 colored by line and intensity. Average point density full of the resolution data is approx. 12/m²






	
Vertical profiles line up both horizontally and vertically with the laser point cloud within <10cm (<1GSD), where tie points 30082880, 30081966, 30081940 are encircled in orange and tie points 30082880, 30082898 are encircled in green	
	
Control point 25 and tie points 30079842, 30080473, 30081060	Area around control point 25 in .las viewer, colored by line and intensity. Average point density using full resolution data on the right side and south of the intersection is approx. 6/m², and approx. 14/m² north of the intersection respectively

	
Control point 25 and tie points line up horizontally with the laser point cloud within <10cm (<1GSD)	Vertical sections line up vertically with the laser point cloud within <10cm (<1GSD)


Example: Oshawa data, Image 45

The figures below illustrate effects of laser point thinning, tie point locations in the image and in the laser data. It also demonstrates the tight fit of the point clouds for a few selected scenes of image 9 of the Oshawa collect.





Control ID	Easting	Northing	Known Z	Laser Z	DZ
25	667799	4862519	82.217	82.22	0.003
26	667522.9	4862885	87.059	87.14	0.081
37	667177.7	4862772	81.893	81.96	0.067
84	667769.2	4862874	81.691	81.56	-0.131
88	667354.1	4862542	93.673	93.75	0.077
92	667317.9	4863024	90.322	90.38	0.058
261	668011.2	4862863	83.383	83.45	0.067
Average dz					0.032
RMS					0.077




Vertical difference between tie points and LiDAR-derived key model ground points


Tie point statistics, 390 tie points, difference to full resolution laser point cloud in meters

Average dz	-0.091
Minimum dz	-10.44
Maximum dz	13.09
Avg magnitude	0.533
RMS	1.454
Std	1.453


APPENDIX H: Change of Data Footprint and Entropy when Estimating Data Shifts using POMF




Autocorrelation of KN elevation maps - expected vs recovered translations



Cross correlation of KN elevation maps – change of size and entropies



Autocorrelation of X5d4 and X5d5 – change of size and entropies



Cross correlation of X5d4 and X5d5 – change of size and entropies