

# Biocomotion Detection in Videos

Soo Min Kang

A Dissertation submitted to  
the Faculty of Graduate Studies  
in Partial Fulfillment of the Requirements  
for the Degree of  
Doctor of Philosophy

Graduate Program in  
Electrical Engineering and Computer Science  
York University  
Toronto, Ontario

August 2019

© Soo Min Kang 2019

## Abstract

Animals locomote for various reasons: to search for food, to find suitable habitat, to pursue prey, to escape from predators, or to seek a mate. The grand scale of biodiversity contributes to the great locomotory design and mode diversity. In this dissertation, the locomotion of general biological species is referred to as **biolocomotion**. The goal of this dissertation is to develop a computational approach to detect biolocomotion in any unprocessed video.

The ways biological entities locomote through an environment are extremely diverse. Various creatures make use of legs, wings, fins, and other means to move through the world. Significantly, the motion exhibited by the body parts to navigate through an environment can be modelled by a combination of an overall positional advance with an overlaid asymmetric oscillatory pattern, a distinctive signature that tends to be absent in non-biological objects in locomotion. In this dissertation, this key trait of positional advance with asymmetric oscillation along with differences in an object's common motion (extrinsic motion) and localized motion of its parts (intrinsic motion) is exploited to detect biolocomotion. In particular, a computational algorithm is developed to measure the presence of these traits in tracked objects to determine if they correspond to a biological entity in locomotion. An alternative algorithm, based on generic handcrafted features combined with learning is assembled out of components from allied areas of investigation, also is presented as a basis of

comparison to the main proposed algorithm.

A novel biolocomotion dataset encompassing a wide range of moving biological and non-biological objects in natural settings is provided. Additionally, biolocomotion annotations to an extant camouflage animals dataset also is provided. Quantitative results indicate that the proposed algorithm considerably outperforms the alternative approach, supporting the hypothesis that biolocomotion can be detected reliably based on its distinct signature of positional advance with asymmetric oscillation and extrinsic/intrinsic motion dissimilarity.

## Acknowledgements

I would like to express my sincere and utmost gratitude towards my supervisor, Professor Richard Wildes for his guidance throughout my studies, his extreme patience when I was slow to comprehend topics, implement ideas, and write reviews, and his continued support, especially when I had my doubts. It has truly been a blessing to be supervised by Rick.

I would like to thank my supervisory committee members, Professor Michael Brown and Professor Marcus Brubaker, for their insightful comments during our meetings, their continued support, encouragement, and advice as I continued to pursue this research topic. I would also like to extend my gratitude towards my internal member, Professor Nikolaus Troje, and external examiner, Professor Dorita Chang, for their suggestions that stemmed from their work in psychophysics. I would also like to thank Professor Robert Allison for agreeing to serve as the chair on my examining committee. Many of my discussions with my examining committee members contributed to the final output of this dissertation in various forms, whether it be a component in the algorithm or comments that triggered intellectual thinking.

I would like to thank the Technical Staff in the Department of Electrical Engineering and Computer Science (EECS) for all their help since my arrival at York University. As they helped me troubleshoot various technical problems to help me



make progress in my research, they also walked through the problems with me to teach me how to solve them myself.

I would like to thank my former and current peers in the Computer Vision Lab and the Department of EECS. I was lucky to be surrounded by such studious and ambitious individuals filled with curiosity who were also very kind and patient. Not only did it motivate me to work similarly, to set similar goals, and to explore to fulfil my curiosities, but I also was able to leverage their experiences in research and life.

Most importantly, I would like to thank my family for their unconditional love and support. It was through my parents' sacrifices and hard work that I had this uncommon opportunity to pursue education abroad for so many years. I would also like to thank my sister for always making me feel like a superstar.

I am indebted to everyone that have been part of this journey with me. Thank you.

# Table of Contents

<b>Abstract</b>	<b>ii</b>
<b>Acknowledgements</b>	<b>iv</b>
<b>List of Tables</b>	<b>x</b>
<b>List of Figures</b>	<b>xi</b>
<b>1 Introduction</b>	<b>1</b>
1.1 Motivation . . . . .	1
1.2 Challenges . . . . .	3
1.3 Contributions . . . . .	4
1.4 Dissertation Outline . . . . .	5
<b>2 Related Research</b>	<b>6</b>
2.1 Introduction . . . . .	6
2.2 Biomechanics . . . . .	7
2.3 Psychophysics . . . . .	12

2.4	Computational Vision . . . . .	14
2.5	Summary . . . . .	18
<b>3</b>	<b>Technical Approach</b>	<b>20</b>
3.1	Introduction . . . . .	20
3.2	Feature Extraction . . . . .	23
3.2.1	Point Trajectories . . . . .	23
3.2.2	Trajectory Post-Processing . . . . .	27
3.3	Biocomotion Detector . . . . .	31
3.3.1	Measure of Locomotion, $\mathcal{M}_\Lambda$ . . . . .	34
3.3.2	Measure of Oscillation, $\mathcal{M}_\Omega$ . . . . .	37
3.3.3	Measure of Asymmetry, $\mathcal{M}_\Sigma$ . . . . .	40
3.3.4	Measure of Extrinsic Motion Dissimilarity, $\mathcal{M}_E$ . . . . .	44
3.3.5	Biocomotion Detector, $S_B$ . . . . .	48
3.4	Adaptive Sliding Temporal Window . . . . .	51
3.5	Algorithmic Summary . . . . .	52
<b>4</b>	<b>Evaluation</b>	<b>54</b>
4.1	Introduction . . . . .	54
4.2	Datasets . . . . .	55
4.2.1	Biological Objects in Locomotion Detection (BOLD) Dataset .	56
4.2.2	Camouflage Animals Dataset (CAD) . . . . .	59
4.2.3	Biocomotion Annotations . . . . .	59
4.2.4	Summary . . . . .	62

4.3	Evaluation Metrics . . . . .	64
4.4	Algorithm Parameter Values . . . . .	67
4.5	Alternative Algorithm . . . . .	68
4.6	Results . . . . .	70
4.6.1	Proposed Algorithm vs. Alternative Algorithm . . . . .	71
4.6.2	Component Analysis . . . . .	72
4.6.3	Qualitative Results . . . . .	77
4.7	Summary . . . . .	79
<b>5</b>	<b>Conclusion</b>	<b>80</b>
5.1	Summary . . . . .	80
5.2	Future Work . . . . .	81
	<b>Appendices</b>	<b>85</b>
<b>A</b>	<b>Trajectory Processing</b>	<b>85</b>
A.1	Pruning Trajectories . . . . .	86
A.1.1	Removing Static Trajectories . . . . .	86
A.1.2	Removing Random Trajectories . . . . .	88
A.2	Clustering Trajectories . . . . .	89
A.3	Stabilizing Trajectories . . . . .	98
A.4	Elongating Trajectories . . . . .	106
<b>B</b>	<b>Alternative Oscillation Measures</b>	<b>109</b>
<b>C</b>	<b>Additional Empirical Results</b>	<b>112</b>



# List of Tables

4.1	Breakdown of motion and objects in BOLD . . . . .	57
4.2	Summary of the datasets. . . . .	63
4.3	Parameters to convert raw measurements into confidence scores . . . . .	67
4.4	AP scores for BOLD and CAD . . . . .	71
C.1	Precision values at specified threshold $\tau_B$ . . . . .	114

# List of Figures

1.1	Variations in modes of locomotion and species . . . . .	3
2.1	Illustration of motion patterns exhibited by locomoting biological species	11
2.2	Illustration of a point-light display . . . . .	13
3.1	General framework of the biolocomotion detection algorithm . . . . .	21
3.2	Illustration of characteristic motion patterns exhibited by biological and non-biological objects . . . . .	22
3.3	Detrending a trajectory . . . . .	27
3.4	Camera motion compensation via stabilization . . . . .	30
3.5	Developed biolocomotion detector component composition . . . . .	32
3.6	A simple real-world illustration of biolocomotion analysis . . . . .	33
3.7	Centroid of points at frame $t$ . . . . .	34
3.8	Calculation of the magnitude of asymmetry . . . . .	42
3.9	Calculation of extrinsic and intrinsic motion dissimilarity . . . . .	47
3.10	Computing displacement vectors for various temporal windows . . . . .	50
4.1	Venn diagram of A2D and BOLD . . . . .	57

4.2	Illustration of videos in BOLD and its annotations . . . . .	58
4.3	Illustration of videos from CAD and its annotations . . . . .	59
4.4	Biocomotion annotation illustrations . . . . .	61
4.5	General framework of the alternative algorithm . . . . .	69
4.6	PR curve for the proposed algorithm . . . . .	74
4.7	PR curve for the alternative algorithm . . . . .	76
4.8	Qualitative results from the proposed biocomotion algorithm . . . .	78
5.1	Illustration of perception of animacy . . . . .	84
A.1	Comparison of single object per frame vs. multiple objects per frame	89
A.2	Silhouettes . . . . .	95
A.3	Illustration of trajectories in presence of linear expansion or contraction	100
A.4	Scale factor . . . . .	104
B.1	Alternative measures of amplitude . . . . .	111
C.1	PR curve for the proposed algorithm . . . . .	113



# Chapter 1

## Introduction

### 1.1 Motivation

Videos have become a vital component of our lives as they contain important information about the world. This information has served humans in various domains from security to robotics to entertainment and many more. Not only are videos important sources of information, their sheer quantity is becoming overwhelming. Given the potential quality of information available from videos and their vast quantity, automated systems for processing and analyzing videos are of great importance.


A wide range of video analysis tasks have been considered in computer vision (e.g., segmentation [40, 127], tracking [20, 39, 66], and action recognition [48, 59, 64]). Curiously, however, a task that appears to not have yet been addressed is the detection of biological entities as they locomote through their environment. In this disserta-

tion, general biological objects in locomotion will be referred to as **biolocomotion** and its detection will refer to determining its spatiotemporal loci in the video.

Be it a natural or artificial intelligent system, the ability to detect biological entities as they locomote would provide the system with powerful information about what subsequent actions to take. In the realm of humans, the detection of biolocomotion can be used in monitoring systems to focus on regions of interest, assistive robots to adapt its behaviour to assist its person of interest [70], sports analysis for broadcast, coaching, or training [81, 110], and autonomous vehicle technology for safe navigation. Similarly, its application can be extended to other animals to monitor wildlife (e.g., to help preserve biodiversity). In addition, a concrete algorithm for biolocomotion detection could provide the basis for a model of how natural systems perform this task.

Biological evidence suggests that natural systems have the ability to detect biolocomotion on the basis of very limited visual data [57, 115]. This ability is especially striking given the wide range of species and modes of locomotion; see Figure 1.1. Perhaps the very wide range of intra-class variations has discouraged previous researchers in artificial systems from delving into biolocomotion detection. Nevertheless, the ability of biological systems to detect biolocomotion from limited data raises the possibility that there may be a distinctive signature to biolocomotion independent of species type and mode of locomotion that could be leveraged by a computer vision system. Given that all biological systems are governed by biomechanical prin-

ciples, a possible basis for such a signature comes from biomechanics. Fortunately, there is a rich literature on the biomechanics of animal locomotion from which to draw [3, 9].



(a) Humans can locomote in various forms (e.g., climb, swim, roll, and walk).



(b) Variations in biological species (e.g., snake, turtle, terrestrial quadruped, and bird).

Figure 1.1: Variations in modes of locomotion and species. (a) Humans can locomote in various ways (e.g., climb, swim, roll, and walk) and (b) biodiversity encompasses a wide range of species (e.g., snake, turtle, cat, and bird).

The apparent ability of natural systems to detect biolocomotion from visual data combined with its potential usefulness provides a strong motivation for studying biolocomotion from a computational perspective.

## 1.2 Challenges

Like any image- or video-based recognition or detection task, biolocomotion detection must be robust to variable acquisition scenarios (e.g., illumination, clutter, intrinsic and extrinsic camera parameters). Beyond these usual concerns, two additional out-

standing challenges must be considered in biolocomotion detection in videos. First, there is extreme diversity in how biological entities move through the environment (e.g., mammals use legs, birds use wings, and fish use fins); even within species, wide variations are present (e.g., humans can walk, run, skip, or swim). Thus, a general biolocomotion detection must be robust to such variations. Second, prior to the work described in this dissertation, there were no extant datasets suitable for developing and evaluating biolocomotion detection algorithms. Correspondingly, it was necessary to construct a novel dataset for the task.

## 1.3 Contributions

The contributions of this dissertation are as follows.

- Biolocomotion detection in videos is introduced as a new research topic in computer vision. Despite the strong motivation for this area of study, it appears that no previous computer vision research has addressed this topic.
- A novel algorithm capable of spatiotemporally detecting biolocomotion in videos is proposed. The proposed algorithm is motivated by biomechanical properties of animals in locomotion and psychophysical studies on biological motion perception; thus, it benefits from not being required to learn the within class variations of biological and non-biological objects in motion.
- An alternative algorithm for biolocomotion detection in videos is presented to provide the main proposed algorithm with a basis for comparison in evaluation.

This algorithm is assembled out of components from allied areas of investigation (action proposals [118] and recognition [121]) and, unlike the main proposed algorithm, relies on training rather than biomechanical modelling.

- A novel biolocomotion video dataset is introduced. The dataset is extremely diverse in capturing terrestrial, aquatic, and aerial biological entities as well as non-biological objects moving in various ways. Biolocomotion groundtruth labels are provided for this dataset as well as an extant camouflage animals dataset from allied area of motion segmentation [8]. These datasets are used to evaluate both of the developed algorithms.

## 1.4 Dissertation Outline

This dissertation unfolds in five chapters. This initial chapter has served to motivate the importance of the proposed problem, biolocomotion detection in videos. Chapter 2 covers related work from various fields ranging across biomechanics, psychophysics, and computer vision. Chapter 3 provides a unified computational algorithm inspired by biomechanics and psychophysics to detect biolocomotion in videos. Chapter 4 provides empirical evaluation of the approach, including introduction of a novel biolocomotion dataset as well as an alternative baseline algorithm. Finally, Chapter 5 provides an overall summary of the presented research, as well as suggestions for future research.

# Chapter 2

## Related Research

### 2.1 Introduction

Animal locomotion has been studied extensively in ethology for some time. Animals locomote for various reasons: to search for food, to find suitable habitat, to pursue prey, to escape from predators, or to seek a mate. The grand scale of biodiversity (ranging from mammals, lizards, birds, fish, insects, and many more) contributes to the great locomotory design and mode diversity. Fortunately, there are common principles that underlie most of these components. Thus, understanding these physical principles would provide a general understanding of why certain biological structures evolved for movement. In complement, biological motion can be detected with very limited visual data by natural visual systems [57, 115]. The goal of this dissertation is to use common traits found in biomechanics of moving biological species and biological motion perception in psychophysics to build a unified computational

approach to detect arbitrary locomoting biological species in videos.

This chapter unfolds in five sections. This first section has served to identify related fields for the development of a biolocomotion detection algorithm in videos. Section 2.2 covers the biomechanical properties that underlie locomotion in biological species. Section 2.3 describes characteristic motions that induce the perception of biological motion. Section 2.4 describes computational work developed for biological motion analysis. Finally, Section 2.5 provides an overall summary of the work from various fields that are necessary to build a unified approach for the detection of biolocomotion in videos. Note that extensive reviews of animal locomotion using biomechanical properties can be found in [3, 9] and surveys on the perception of biological motion in psychophysics can be found in [113, 114].

## 2.2 Biomechanics

While the means by which animals traverse their environment is extremely varied, common locomotory mechanisms have emerged as a result of biomechanical constraints. In particular, animal locomotion is typically accompanied by the overall positional advance with an asymmetric oscillatory trace of the body parts to provide a propulsive force, as will be detailed in this section.

Land, air, and water constitute the type of environments animals move through. The properties of these media, such as density and viscosity, can influence the loco-

motory mechanisms evolved by the animals [9]. Air has lower density than water. Thus, aerial animals must exert sufficient forces to support their weight in air, while most aquatic animals are neutrally buoyant since their body density is nearly the same as water [9]. Air also has lower viscosity than water. This lower viscosity imposes less, but non-absent, drag forces on flying and terrestrial animals compared to aquatic animals. In essence: aerial animals must generate enough force to lift their bodies as well as thrust to overcome the aerodynamic drag forces associated with moving forward; aquatic animals must swim strategically to reduce drag forces induced by high viscosity of water; and terrestrial animals must overcome gravitational forces as they move.

For each medium, animals can locomote in various forms. Movement in air can be achieved by gliding or flapping, where flight by *glide* generates lift by keeping the wing fixed and exploits the airflow for movement and *flapping* generates lift and thrust forces simultaneously by continuous wing oscillation [3, 19]. Movement in water can be achieved by lift-powered swimming, undulation, drag-powered swimming, or jet propulsion. *Lift-powered swimming* involves flapping fins or tail to propel forward, *undulation* refers to the oscillation of the entire body, *drag-powered swimming* pushes water backwards by using fins or limbs as oars that move back and forth, and *jet-propulsion* involves sequential ingestion and expulsion of finite mass of water [3, 19]. Movement on land can be achieved by crawling, walking, running, hopping, or jumping [3, 19]. A common trait across these different means of powered locomotion is the activation of muscles to lengthen and shorten at constantly changing speeds



to accelerate and decelerate moving body parts. An exception is glide, as glide is an unpowered flight as it generates no mechanical power with its flight muscles [9, 19]. These powered actions apply forces in an oscillatory manner such that a structure with mass (body and/or its part) oscillates in its environment [3, 9].

A generalized model to understand the dynamics of legged terrestrial locomotion (e.g., trot, run, and hop) often builds on the **bouncing spring-mass model** [13, 78]. The spring-mass model consists of a massless spring attached to a point mass, where the leg is represented by the spring and the body of an animal is represented by the point mass. A point-mass spring has been used to model uniformly a wide range of species (e.g., humans, dogs, kangaroos, land birds, crabs, and cockroaches) across different locomotory designs (e.g., number, length, shape, position, and skeleton type of legs), as the relative vertical ground-reaction force and the relative compression of the leg are the essential, yet common, components required to move an animal's centre of mass [12]. Indeed, several walking [27, 76], running, and hopping [94] robots and computational models for tracking [18] have emerged through the understanding of legged locomotion via the bouncing spring-mass model [19, 76]. Overall, the bouncing spring-mass model further underlines that oscillation can be an important component of biological locomotion.

The use of rotating systems as a means of transport has brought tremendous efficiency in artificial locomotory devices (e.g., wheels and propellers). However, very few biological species have adapted rotating systems as a means for locomotion;

notable exceptions include bacterial flagellum, rolling spiders, caterpillar-tred stomatopods, and pangolins [67]. The lack of rotating systems in natural systems can be attributed to environmental constraints or the level of efficiency. That is, consider wheel-based transportation on terrestrial surfaces. These are efficient modes on flat rigid terrains (e.g., asphalt and concrete), but far less on irregular terrains - a very common characteristic of natural terrains. In addition, wheel-based transportation require corners to be wide enough and not too sharp for manoeuvrability, which would be a severe disadvantage for species in cluttered terrains. As a result, natural selection favours the evolution of limbs capable of travelling on irregular surfaces and manoeuvring around obstacles [9, 67]. As another example, thrusting in aerial or aquatic media by propellers (as done by artificial locomotory devices) is far less energetically efficient than oscillating flexible foils (as in caudal fin of fish and bird wings) [61]. Consequently, natural systems favour the oscillation of the body and/or its appendages for its energetic savings [67].

Moving efficiently is a very important aspect of biological motion. In terrestrial locomotion, specifically in walking, electromyography (EMG) data revealed that once the leg muscles are activated to set the foot into motion during the stance phase, its muscles are almost inactive during the swing phase, such that the foot moves entirely under the influence of gravity [29]. In aerial locomotion, the wings rotate to move down and in front of its body during the downstroke and move up and slightly backwards during the upstroke to generate lift and thrust forces simultaneously [7, 9, 129]. In aquatic locomotion, the body and/or appendages accelerate to

produce a propulsive force, then decelerate before initiating its subsequent propulsive stroke to maintain steady speed in its viscous environment [9, 42]. In general, the effective means of locomotion in various media results in an asymmetric path traced by body parts of animals as they traverse their environments. This asymmetry results because a steeper slope is observed during the lift/stance phase of a walk compared to the swing phase, downstroke compared to upstroke during flight, and initiation compared to completion of a propulsive stroke in swimming; see Figure 2.1.




Figure 2.1: Illustration of characteristic motion patterns exhibited by locomoting biological objects in various media (human on land, bird in air, and fish in water). As objects move across time (indicated by varying grey levels, with lighter-to-darker moving forward in time), the means that biological objects use to locomote (legs for humans, wings for birds, and body/fins for fish) exhibit an asymmetric sinusoidal pattern (marked in red).

Overall, consideration of the biomechanics of animal locomotion shows that regardless of the wide range of locomotory designs of different animals in various media, animal locomotion involves the use of its body and/or appendages to generate propulsion. Specifically, the acceleration and deceleration in an oscillatory fashion of the body and/or its parts to move in its environment is commonly observed. Moreover, the oscillation typically unfolds in an asymmetric fashion across time. Provided such patterns can be visually observed, there is potential for the development of a principled approach for biolocomotion detection in videos.

## 2.3 Psychophysics

Psychophysical studies have shown that human observers are able to perceive a set of dynamic dots as a coherent figure representative of a person or other animals in motion, provided the dots are located near major joints and their motions are consistent with their representative figures [55, 75]. These visual stimuli are referred to as **point-light displays**; see Figure 2.2. Studies on point-light displays show that biological motion can be perceived in the absence of any relevant appearance information (e.g., body silhouette, texture, or colour). It also has been shown that other animals can perceive biological motion in such displays (e.g., cats [11], pigeons [35], and chicks [117]). Indeed, not only can gross motion patterns be discriminated (e.g., walk, run, and stair climb), but more subtle differences can be perceived (e.g., gender [4, 65, 112] and emotion [36, 91]). It is argued that in making such inferences, biological visual systems decompose motion in terms of common motion, referred to as **extrinsic motion**, and relative motion between parts, referred to as **intrinsic motion** [55]. Interestingly, neuroimaging studies have been able to localize biological motion processing in the brain [22].

The **direction discrimination task** is a task that asks the observers to indicate *which* direction (left or right) the figures depicted in point-light displays are facing. Such experiments have shown that direction discrimination accuracies are highly correlated with the amount the displays appeared animate when variations of the display (coherent vs. scrambled and upright vs. inverted) were presented to the observers [23]. Consequently, the direction discrimination task has often been used to




Figure 2.2: Illustration of a point-light display. A set of 11 markers are used to represent the major joints of the human body (the head, shoulder, hip, two elbows, two wrists, two knees, and two ankles). Three frames from a sequence of such markers animated in accord with the motion of a walking person. Human observers perceive a temporal sequence of such frames as resulting from a walking human. Figure redrawn from [55].

study the perception of biological motion. Further studies on the direction discrimination task have indicated that the local motion, of feet in particular, play a vital role in the accuracy of direction identification [57, 115]. Experiments comparing the displays of naturally accelerating foot motions with those containing constant speeds revealed that the acceleration contained in the foot motion plays a significant role in direction discrimination accuracy [24, 49]. It has been shown that other animals (e.g., newly hatched chicks) show similar sensitivity to vertical acceleration, as they respond to upright point-light displays (of a hen) by aligning their bodies in the apparent direction of motion but not to inverted displays [117]. These findings suggest that biological motion perception is based on the vertical acceleration patterns the foot exhibits as an animal moves through the environment.

Psychophysical evidence suggests that humans can make inferences of animacy from the trajectory that an object traces [10]. Further evidence shows that humans can discriminate between symmetric and asymmetric trajectories [82]. Combining these pieces of evidence with the above reviewed work on point-light displays suggests that motion information similar to those found in biomechanics that indicate locomoting biological species (e.g., asymmetric oscillatory traces) may also be exploited in biological vision systems. Moreover, they suggest the potential efficacy of decomposing motion into intrinsic and extrinsic components in making such inferences.

## 2.4 Computational Vision

Some computational vision work was inspired by the ability of humans to perceive biological structure and motion in point-light displays viewed across time. Early work along these lines reconstructed the 3D structure and motion of animals using anatomical constraints by observing that animal limbs are (i) rigid, (ii) have a fixed length, and (iii) typically move in a single plane for extended periods of time [50]. While these anatomical constraints are generally true of legged terrestrial and aerial animals, they are not true for undulating animals. Furthermore, estimating the species-invariant animal pose in a non-intrusive way with a limited set of training data is a challenging task that further limits the exploitation of the proposed method from working on point-light displays.

Other work made use of 3D periodicity constraints [132]. The goal is to reconstruct a 3D structure from motion capture data of humans, rather than to detect general biological species in locomotion. Nevertheless, a walking human is described as a Fourier representation of (i) average posture, (ii) characteristic postures of the fundamental frequency, (iii) second harmonic of a discrete Fourier expansion, (iv) fundamental frequency, which does not explicitly model vertical acceleration that is present. While the notion of asymmetry is mentioned, it is in the context of left and right limb asymmetry and not within a stride as prevalent in animal locomotion.

Taking further inspiration from biology regarding how the mammalian visual system appears to process information in two parallel streams for form and motion [46], a corresponding algorithm was developed to infer biological shape and motion models from sparse point displays [45]. Specifically, the form pathway that analyzes the body shapes was modelled using Gabor-like filters [58] to obtain orientation details, max-like pooling [97] to provide position and scale-invariance, then Gaussian radial basis functions [89] to support selectivity towards complex shapes. The motion pathway was modelled using optical flow [44] patterns to mimic the direction sensitive and motion sensitive patterns in our brain. While their model provided a computational demonstration that the motion (dorsal) pathway is predominantly active (and that the form (ventral) pathway tends not to be activated) in the recognition of the point-light displays, it did not address how the motion model can be used to detect general biolocomotion nor the specific motion patterns that were learned for the categorization.

Yet another method considered decomposing motion exhibited by an animation (e.g., a person walking or strutting, a kangaroo or a rabbit hopping) into global and local components [103], akin to extrinsic and intrinsic motions used to describe perceptual organization in the biological visual system [55]. The global component is responsible for measuring the motion of the object’s centre of mass and the local component measures the rate of dispersion of the object about its centre of mass. Similar to other computational work to date, this algorithm did not address how a general biological species can be detected nor how a non-biolocomoting object can be rejected.

Periodic motion has been used in previous work to detect, track, and classify objects in videos (e.g., humans and dogs) [1, 16, 30, 83, 90, 95, 102, 116]. Across this research direction, various approaches have been proposed for periodic motion detection, including time-frequency analysis [16, 30, 100, 116], period trace [102], hypothesis testing on periodograms [95], and convolutional neural networks (CNNs) [71]. A limitation of these approaches is that they rely on non-trivial preprocessing of their input videos, including extraction of points corresponding to the major joints of the human body [116], conversion to figure-centric volumes [30, 90, 95], background-subtraction [16], require a precise slicing of a video along the XT-axis [83], or require static cameras [71] to obtain data indicative of periodic motion. Furthermore, the computational vision literature often has modelled a person’s walk via an inverted pendulum [1, 107], while a spring-mass system is a more accurate representation, as it accounts for the vertical ground-reaction force [43]. Significantly, none of these



approaches have been applied to the challenge of general biolocomotion detection of species on land, in air, and in water. Indeed, analysis of oscillation alone does not suffice for detection of biolocomotion, as not all oscillating objects are locomoting biological entities (e.g., person using a jump rope or a pendulum).

There has been growing interest in applications of computer vision to species classification and detection for wildlife. Correspondingly, several datasets that concentrate on imagery of wildlife in natural habitats (e.g., CUB200 [123], NABirds700 [52], iNat2017 [51], Snapshot Serengeti [109], Missouri Camera-Trap [131], and CCT-20 [6]) have been made available to the research community. CUB200, NABirds700, iNat2017, and CCT-20 are image datasets; and the Snapshot Serengeti dataset consists of image sequences collected from a camera trap, which are heat- or motion-activated cameras that capture a single image or a short sequences of images (1-5 frames with a frame rate of approximately 1 frame per second) at each trigger [84]. The interest of the current dissertation, however, lies in common videos rather than specialized ones, which are typically of longer duration and have higher frame rate (typically between 25-30). Furthermore, while datasets extracted from camera traps (e.g., Snapshot Serengeti, Missouri Camera-Trap, and CCT-20) pose various challenges similar to those in-the-wild, such as dynamic background, illumination changes, cluttered and dynamic scenes, they lack camera motion. Thus, some detection algorithms developed to perform reasonably on camera trap data are limited to videos with limited camera motion [6, 72]. Moreover, CUB200 and NABirds700 focus on fine-grained species classification, thus only consist of bird images, while

the goal of the present dissertation is to detect a wide range of species in locomotion. Thus, current algorithms developed to perform reasonably on wildlife datasets are either constrained to images [25], images that have been manually cropped to delineate the animals of interest [25, 128], species-specific [80, 125], or background-specific [130, 131] requiring sufficient training data to model those background in the test set [25].

Overall, while considerable computational work has addressed biological motion analysis or detection and classification of biological species in images, none has considered the detection of general biolocomotion in videos.

## 2.5 Summary

Despite the diverse locomotory designs that exist in different animals, there are significant mechanical and energetic similarities in the body and/or its parts for various types of locomotion. These underlying biomechanical constraints have been incorporated into minimalist models of legged animal locomotion in terms of a bouncing mass-spring model [13, 77, 93] that have been successfully applied to a wide range of terrestrial animals. Significantly, the implied overall positional advance of the body along with asymmetric oscillatory motion of its parts is present not just in terrestrial legged locomotion, but also extends to non-legged terrestrial [74], aerial [92], and aquatic [69, 111, 124] creatures.

Psychophysical studies suggest that the vertical acceleration pattern induced by gravitational and biomechanical constraints in terrestrial creatures play a significant role in the perception of biological motion [24, 49]. Vertical acceleration exhibited by these creatures as they push off against gravity causes their trajectories to trace asymmetric oscillatory patterns, while non-biological objects that locomote with oscillation display symmetric cycloidal patterns during their advance (e.g., rolling objects). Moreover, other studies from psychophysics have revealed that the human visual system decomposes the kinematics of an object into common translatory and residual motion (i.e., extrinsic and intrinsic motion) to understand the mechanics of a scene [55, 56]. The overall direction of an object (e.g., translation of a walker) and the local cues of the body (e.g., trajectory of the feet) tend to be different in locomoting biological entities as compared to non-biological objects.

While previous computational work has addressed animal motion analysis, none has addressed the detection of general biolocomotion in videos. Building on work from the biomechanics and perceptual psychophysics of biological motion, the remainder of this dissertation presents the first algorithm capable of spatiotemporally detecting biolocomotion in videos. Notably, the biomechanical and psychophysical principles motivate the definition of species-invariant biolocomotion signatures so that the approach benefits from not needing to model the wide range of within class variations of biological and non-biological objects in motion.

# Chapter 3

## Technical Approach

### 3.1 Introduction

The goal of a biolocomotion detection algorithm is to take an unprocessed video and output the spatiotemporal coordinates of biolocomotion; see Figure 3.1. To achieve this goal, distinctive properties of biolocomotion as seen in videos must be defined. A key trait that can be observed from biomechanics of animal locomotion and perception of biological motion is a directional trajectory modulated by asymmetric oscillation along with differences in its overall motion and its local cues. This observation is used as the basis for biolocomotion detection. In particular, a collection of trajectories across an image sequence that show an overall advance in spatial position of their tracked elements (i.e., *locomotion*) that exhibit *asymmetric oscillation* and overall motion (*extrinsic motion*) difference from its local cues (*intrinsic motion*) is detected; e.g., Figure 3.2.




Figure 3.1: General framework of the biolocomotion detection algorithm. The input and output of the proposed biolocomotion algorithm are marked in blue and the intermediate components that are involved in obtaining the spatiotemporal coordinates of locomoting biological objects in a video are marked in green.

This chapter unfolds in five sections. This first section has served to define the problem under consideration and outline the components necessary for biolocomotion detection. Section 3.2 describes the extraction of primitive features that can encapsulate critical information for modelling biolocomotion in terms of image point trajectories traced during biolocomotion. Section 3.3 presents algorithmic measures that map the extracted features to various components of the developed biolocomotion detector. Section 3.4 presents a sliding window realization of the approach for continuous video processing. Finally, Section 3.5 provides an overall summary of the approach.




Figure 3.2: Illustration of characteristic motion patterns exhibited by (a) biological objects in locomotion, (b) non-biological objects in locomotion, and (c) a non-biological object in oscillation. As objects move across time (indicated by varying grey levels, with lighter-to-darker moving forward in time), the means that biological objects use to locomote (e.g., legs for humans, wings for birds, body/fins for fish) exhibit an *asymmetric* sinusoidal pattern (with a point in yellow generating the red trace). Non-biological objects in locomotion, on the other hand, tend to exhibit either trajectories that lack oscillation as they move (e.g., car and plane), trajectories with *symmetric* oscillation (e.g., rolling ball), or the extrinsic motion (i.e., overall direction of motion) (blue) and its intrinsic motion (i.e., local cues) (red) tend to coincide (e.g., bouncing ball), while an oscillating non-biological object (e.g., metronome) does not have an accompanying overall positional advance. Note that the alignment of the extrinsic and intrinsic curves yield an alternating red and blue traces for locomoting non-biological objects, except those where the intrinsic component traces a cycloid.

## 3.2 Feature Extraction

In this dissertation, a collection of tracked point trajectories serve as inputs to the biolocomotion detector. Thus, in this section, a definition of point trajectories and subsequent post-processing that make the trajectories more amenable to further biolocomotion processing are provided.

### 3.2.1 Point Trajectories

Given an input video, point trajectories provide the path in which the tracked point travelled across time. In particular, point trajectories support quantitative measurement of key components that distinguish biolocomotion: spatial advance in overall position, asymmetric oscillatory traces, and a difference between extrinsic and intrinsic motions. Consequently, a collection of tracked point trajectories are used as inputs to the biolocomotion detector. There are variety of approaches to obtaining point trajectories available in the field of computer vision (e.g., KLT Trajectories [79], SIFT Trajectories [108], and Dense Trajectories [120, 122]). In this dissertation, the recovered trajectories are built on improved Dense Trajectories (iDTs) [121] that previously supported state-of-the-art performance amongst handcrafted algorithms for action recognition. This choice is made since biolocomotion detection itself can be conceptualized as a type of an action.

The iDTs are built on Dense Trajectories (DTs), which are obtained by densely sampling feature points on a grid space by  $S$  pixels over several spatial scales to en-

sure that the feature points are sampled from all spatial positions and scales. While denser sampling (e.g.,  $S = 2$  to sample every other pixel) offers better performance, it significantly increases computational complexity. Thus, DTs extracted at a sampling stride of  $S = 20$  are used in present work. Since points in homogeneous areas are difficult to track reliably, they are removed using the good-features-to-track criterion [104], which removes points with very small eigenvalues of the auto-correlation matrix. Feature points are tracked at each spatial scale separately. Each feature point  $\mathbf{x}(t) = (x(t), y(t))$  at frame  $t$  is tracked to the next frame  $t + 1$  by median filtering on a dense optical flow field  $\mathbf{u}(t) = (u(t), v(t))$ , where  $u(t)$  and  $v(t)$  are horizontal and vertical components of the optical flow, respectively. Specifically, given a point  $\mathbf{x}(t)$ , its tracked position in the next image frame is smoothed by applying a median filter on  $\mathbf{u}(t)$ :

$$\mathbf{x}(t + 1) = \mathbf{x}(t) + (M * \mathbf{u}(t))|_{\mathbf{x}(t)}, \quad (3.1)$$

where  $M$  is a  $3 \times 3$  median filtering kernel. Points of subsequent frames are concatenated to form trajectories  $\mathcal{T} = [\mathbf{x}(t) \ \mathbf{x}(t + 1) \ \dots \ \mathbf{x}(t + L)]$ . To overcome the drifting effect (i.e., points drifting from their true locations during the tracking process), the length of the trajectories are limited to  $L$  frames. Empirically, dense trajectories that span  $L = 15 - 20$  frames have been found effective in the action recognition literature [122]. Thus, trajectories of  $L = 15$  are used in the current work.

Dense trajectories can be improved by removing the global background motion created by camera motion. Here, camera motion is estimated by assuming that two consecutive frames are related by a homography, where the homography is es-



timated by finding correspondences between two frames. The correspondences can be found by: (i) extracting SURF features [5] and matching them based on the nearest neighbour rule, and (ii) sampling motion vectors from optical flow using the good-features-to-track criterion. The candidates from the two approaches are used to estimate the homography using RANSAC [41] to rectify the image to remove camera motion. Compared to the original flow, the rectified version suppresses the background camera motion and enhances the foreground moving objects. Trajectories generated by camera motion are removed by thresholding the displacement vectors of the trajectories in the warped flow field. If the displacement is too small, the trajectory is considered to be too similar to camera motion, and thus removed.

In this dissertation, a tracked point trajectory,

$$\mathcal{T}^k(t) = \left[ \mathbf{x}_1^k(t) \quad \mathbf{x}_2^k(t+1) \quad \cdots \quad \mathbf{x}_{L^k}^k(t+L^k-1) \right] \quad (3.2)$$

denotes trajectory  $k$  that begins at frame  $t$  with a temporal length of  $L^k$  and its  $l^{\text{th}}$  point for  $1 \leq l \leq L^k$  is specified by

$$\mathbf{x}_l^k(t) = (x_l^k(t), y_l^k(t)). \quad (3.3)$$

Note  $\mathcal{T}^k(t)$  and  $\mathcal{T}^k$  as well as  $\mathbf{x}_l^k(t)$ ,  $\mathbf{x}^k(t)$ , and  $\mathbf{x}_l^k$  will be used interchangeably for simplicity throughout this dissertation, where  $\mathbf{x}^k(t)$  will be used to emphasize a point present at frame  $t$  and  $\mathbf{x}_l^k$  to emphasize the  $l^{\text{th}}$  point of trajectory  $\mathcal{T}^k$ . Furthermore,  $x$  and  $y$  components will be referred to as horizontal and vertical components, re-

spectively, of the trajectory.

A **displacement vector** of trajectory  $\mathcal{T}^k$  at frame  $t$  from  $\Delta t$  previous frames is defined as

$$\Delta \mathbf{x}_{(t, \Delta t)}^k = \mathbf{x}^k(t) - \mathbf{x}^k(t - \Delta t). \quad (3.4)$$

For simplicity,  $\Delta \mathbf{x}^k(t)$  will be used to denote  $\Delta \mathbf{x}_{(t, 1)}^k$ .

The **arc length** of trajectory  $\mathcal{T}^k$  is defined as

$$\|\mathcal{T}^k\| = \sum_{l=1}^{L^k-1} \sqrt{(x_{l+1}^k - x_l^k)^2 + (y_{l+1}^k - y_l^k)^2}. \quad (3.5)$$

Select measurements, such as amplitude and asymmetry, are sensitive to the spatial direction of the trajectory as it unfolds across time. To ensure these calculations are robust to such situations, it is necessary to **detrend** trajectory  $\mathcal{T}^k$  to  $\tilde{\mathcal{T}}^k$ . Trajectory  $\mathcal{T}^k$  can be detrended by (i) determining the line of best fit using least squares, (ii) finding the angle,  $\theta$ , between the line of best fit and the positive x-axis of a 2D Cartesian plane, (iii) rotating trajectory  $\mathcal{T}^k$  by  $-\theta$ , then (iv) applying vertical translation such that its horizontal mean is 0 (i.e.,  $\text{mean}_{\forall l}(\tilde{y}_l^k) = 0$ ). Note that proper measurement of a trajectory's oscillation amplitude depends on the correct ordering of rotation and demeaning; see Figure 3.3.




Figure 3.3: Detrending a trajectory. To obtain an accurate measurement of the trajectory amplitude, the trajectory must be rotated then demeaned. Given a trajectory representative of a sinusoid rotated  $\theta = 45^\circ$  counterclockwise about the origin with an amplitude of  $A = 1$  (red), the amplitude of its demeaned trajectory (green) is  $\frac{1}{\sqrt{2}}$ , while the amplitude of a rotated then demeaned trajectory (blue) is 1.

### 3.2.2 Trajectory Post-Processing

Collections of point trajectories serve as input to the proposed biolocomotion detector. While iDTs have previously supported state-of-the-art performance in action recognition, it is necessary to further post-process the iDT results prior to passing them to the biolocomotion detector, as follows. First, they must be pruned to remove unuseful trajectories for biolocomotion. Second, they need to be clustered so that biolocomotion detection operates on sets of trajectories that are likely to correspond to a single or similarly moving entity in the world. Third, iDTs do not have adequate robustness to camera motion for present purposes. Correspondingly, they need to be

further stabilized. Fourth, it is useful to elongate them to allow for more temporal support in biolocomotion detection. In the remainder of this section, the entailed processing steps are outlined with details provided in Appendix A.

### **Pruning Trajectories**

Static trajectories and random trajectories are unlikely to provide meaningful information in identifying biolocomotion. Hence, trajectories that do not contain motion information or trajectories with sudden large displacements, which are likely to be erroneous, are removed before they are further processed. While original iDT calculations include measures to remove static and random trajectories [121], they are deemed either insufficient or too aggressive for current purposes. In response, variants on conditions considered in iDT calculations have been defined and employed.

### **Clustering Trajectories**


Given a set of trajectories, the trajectories are clustered into disjoint sets such that each cluster corresponds to a single or similarly moving object in the world. A variant on spectral trajectory clustering is employed as the original formulation [17] produced poor results for the considered trajectories. Correspondingly, alternative measures of positional and shape affinities are defined for trajectory pairs in spectral clustering. The output of this processing stage are the centres as well as the horizontal and vertical extents of ellipses that cover each cluster.

## Stabilizing Trajectories

Many videos in-the-wild have camera motion. While iDTs are designed for some robustness to camera motion, their processing is inadequate for current purposes for two reasons. (1) In some cases, objects do not exhibit their actual locomotion in the captured video because it was recorded to stabilize the object of interest in the field of view. In such cases, the tracked object would remain in the same position (exhibiting near zero displacement) within the image, while it is in locomotion in the real world (exhibiting non-zero displacement). By undoing the stabilization the camera operator has imposed, the tracked object in the image would be more representative of its locomotion, also exhibiting non-zero displacement; see Figure 3.4. The iDTs do not model such situations. (2) In other cases, iDTs simply contain too much residual motion arising from the camera to adequately capture signatures representative of biolocomotion. In response, the trajectories are stabilized to reveal the motion of the objects in the field of view with additional robustness to camera motion.



(a) Object-centric image sequence from operator imposed camera motion to track object of interest.



(b) Camera motion stabilization can reestablish object motion.

Figure 3.4: Camera motion compensation via stabilization. (a) Select frames demonstrating an object-centric stabilized video. The object of interest (dog) appears near the centre of the frame across the image sequence even though it is locomoting in the real-world. (b) Corresponding frames with camera motion stabilization. Camera motion that the camera operator has imposed to centre the object of interest within the field of view (i.e., pan) is undone, such that features extracted from the camera motion stabilized video is more representative of biolocomotion in the real-world.

## Elongating Trajectories

To overcome the drifting effect when tracking points, it is recommended that iDTs only span 15 – 20 frames for videos captured at 30 fps [121]. In contrast, biolocomotion detection benefits from longer trajectories, especially to provide sufficient time to diagnose asymmetric oscillatory behaviour. In response, iDTs are concatenated across frames based on spatial proximity and appearance similarity to obtain elongated trajectories.

### 3.3 Biolocomotion Detector

A collection of tracked point trajectories are used as the input to the proposed algorithm to analyze motion in terms of (i) locomotion, (ii) oscillation, (iii) asymmetry, and (iv) extrinsic motion dissimilarity<sup>1</sup>; see Figure 3.5. In this section, approaches to quantifying these components will be defined, followed by a way to combine these measures into a single biolocomotion detector. Throughout the section, each component will be defined while alluding to the simple real-life example in Figure 3.6 that variously exhibits locomotion without oscillation (toy car), oscillation without locomotion (pendulum), symmetric oscillation with locomotion (rolling ball), asymmetric oscillation with coinciding extrinsic and intrinsic motions (bouncing ball), and asymmetric oscillation with dissimilar extrinsic and intrinsic motions (person).

---

<sup>1</sup>For the sake of compactness, the dissimilarity between extrinsic and intrinsic motion will be referred to as **extrinsic motion dissimilarity** in the following.




Figure 3.5: Composition of components of the developed biocomotion detector. The detector consists of components that measure locomotion,  $\mathcal{M}_\Lambda$ , oscillation,  $\mathcal{M}_\Omega$ , asymmetry,  $\mathcal{M}_\Sigma$ , and extrinsic motion dissimilarity,  $\mathcal{M}_E$ . *Locomotion* is measured based on the aggregate displacement of the centre of mass, *oscillation* and *asymmetry* are calculated based on amplitude and area-under-the-curve of the trajectories, respectively, and *extrinsic motion dissimilarity* is measured based on the dissimilarity between extrinsic and intrinsic displacement vectors.



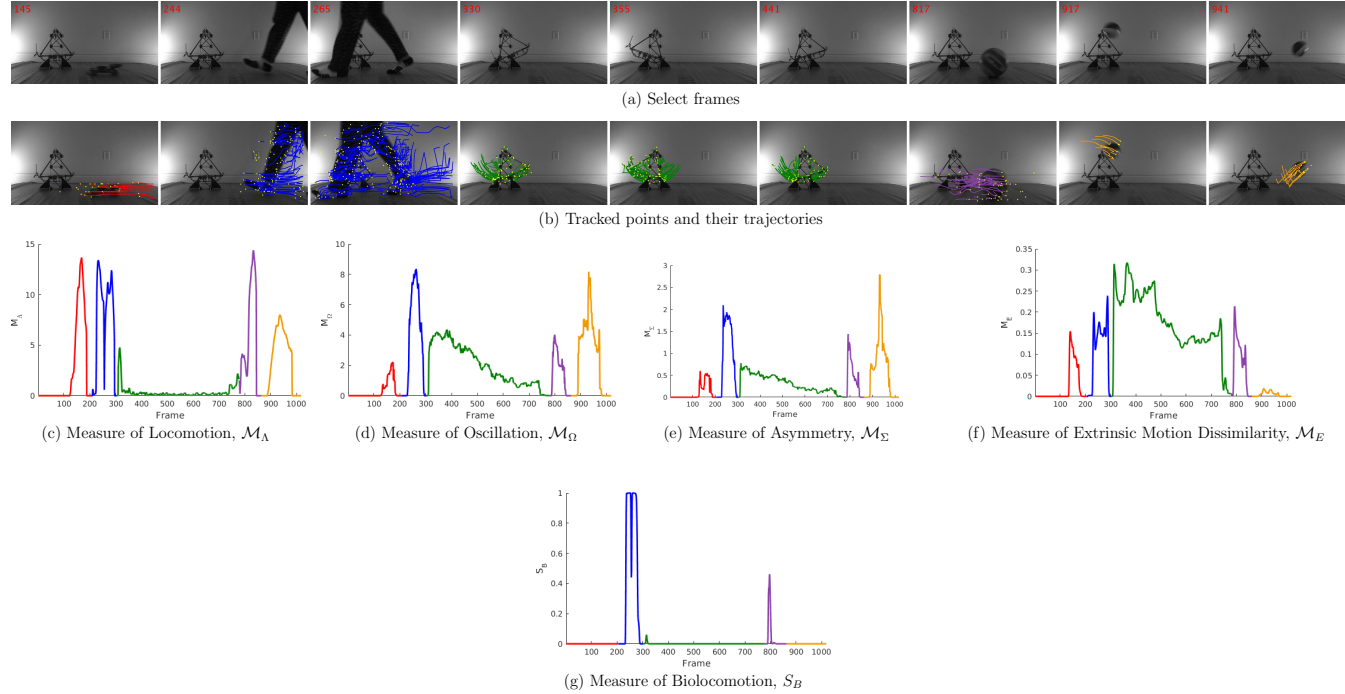


Figure 3.6: A simple real-world illustration of biolocomotion analysis. (a) Select frames demonstrating non-biological objects in locomotion (toy car at frame 145, rolling ball at frame 817, and bouncing ball at frames 917 and 941), a biological object in locomotion (person at frames 244 and 265), and a non-biological object in oscillation (pendulum at frames 330, 355 and 441). (b) Tracked points (yellow) and their trajectories overlaid with red, blue, green, purple, and orange indicating the toy car, person, pendulum, rolling ball, and bouncing ball, respectively. (c-g) Recovered measures of locomotion ( $\mathcal{M}_\Lambda$ ), oscillation ( $\mathcal{M}_\Omega$ ), asymmetry ( $\mathcal{M}_\Sigma$ ), extrinsic motion dissimilarity ( $\mathcal{M}_E$ ), and biolocomotion ( $S_B$ ) with respect to time (in frames).

### 3.3.1 Measure of Locomotion, $\mathcal{M}_\Lambda$

The amount an object displaces from one location to another in a given temporal window is quantified using a set of trajectories by calculating the **measure of locomotion**,  $\mathcal{M}_\Lambda$ ; see blue curves in Figure 3.2a,b or lack thereof in Figure 3.2c. The measure of locomotion can be evaluated by considering the norm (e.g.,  $L_1$ ,  $L_2$ , etc.) of the average displacements of the object over some temporal window. In the following, a measure of locomotion is described.




Figure 3.7: Centroid of points at frame  $t$ . (a) Input video at frame  $t$ . (b) Points present at frame  $t$  (yellow dots) can be extracted from a set of trajectories that are present at frame  $t$  (red lines). (c) The centroid (red cross) can be obtained by computing the trimmed mean of the points that are present at frame  $t$  (yellow dots in (b)).

The displacement of an object is measured by considering the displacements of the centroid of an object over some temporal window; see Figure 3.7. The **centroid of an object** at frame  $t$ ,  $\bar{\mathbf{x}}^c(t)$ , is calculated by considering the trimmed mean of points that are present at frame  $t$  for  $1 \leq t \leq T$  according to

$$\bar{\mathbf{x}}^c(t) = (\bar{x}^c(t), \bar{y}^c(t)) = \left( \frac{1}{|X(t)|} \sum_{x^k(t) \in X(t)} x^k(t), \frac{1}{|Y(t)|} \sum_{y^k(t) \in Y(t)} y^k(t) \right), \quad (3.6)$$

where  $X(t)$  and  $Y(t)$  are sets of points measured along the horizontal and vertical axes, respectively, present at frame  $t$  within 1.5 of the standard deviation from its respective mean. That is, suppose  $\mu_x(t)$  and  $\sigma_x(t)$  denotes the mean and the standard deviation, respectively, of trajectory points present at frame  $t$  measured along the horizontal axis. Then a set of points,  $X(t)$ , along the horizontal axis at frame  $t$  is defined as

$$X(t) = \{x^k(t) \mid \mu_x(t) - 1.5\sigma_x(t) \leq x^k(t) \leq \mu_x(t) + 1.5\sigma_x(t)\}. \quad (3.7)$$

To ensure that reliable data contributes to the calculation of  $\bar{x}^c(t)$ , only data within 1.5 of the standard deviation from the mean are considered. Discarding data that are outside 1.5 of the standard deviation from the mean ensures the final measurement,  $\bar{x}^c(t)$ , is based on approximately 86.67% of the total data while the lowest 6.68% and highest 6.68%, that may correspond to noise, are discarded.

Similarly, a set of points,  $Y(t)$ , along the vertical axis at frame  $t$  is defined as

$$Y(t) = \{y^k(t) \mid \mu_y(t) - 1.5\sigma_y(t) \leq y^k(t) \leq \mu_y(t) + 1.5\sigma_y(t)\} \quad (3.8)$$

for mean and standard deviation,  $\mu_y(t)$  and  $\sigma_y(t)$ , respectively, of trajectory points present at frame  $t$  along the vertical axis.

Consequently, the **centroid displacement vector** at frame  $t$  from a previous

frame at  $t - \Delta t$  is defined as

$$\Delta \bar{\mathbf{x}}_{(t, \Delta t)}^c = \bar{\mathbf{x}}^c(t) - \bar{\mathbf{x}}^c(t - \Delta t). \quad (3.9)$$

Note that  $\Delta \bar{\mathbf{x}}_{(t, \Delta t)}^c$  is similar to  $\Delta \bar{\mathbf{x}}_{(t, \Delta t)}^k$  for trajectory  $\mathcal{T}^k$  (3.4), except  $\Delta \bar{\mathbf{x}}_{(t, \Delta t)}^c$  is specific for centroids.

Finally, the measure of locomotion is obtained by combining the overall horizontal and vertical displacements of the centroids over a temporal window. Mathematically, it is defined as

$$\mathcal{M}_\Lambda(t) \equiv \sqrt{\left( \frac{1}{\min\{\omega, t\}} \sum_{n=\max\{t-\omega, 1\}}^t \Delta \bar{x}^c(n) \right)^2 + \left( \frac{1}{\min\{\omega, t\}} \sum_{n=\max\{t-\omega, 1\}}^t \Delta \bar{y}^c(n) \right)^2}, \quad (3.10)$$

for some constant  $\omega > 0$ , which represents the length of the temporal window. Empirically,  $\omega$  is set to the frame rate of the video. Since the goal is to quantify displacement rather than distance, framewise displacements are aggregated (without considering their absolute values) before overall magnitude is calculated.

It can be observed in Figure 3.6c that as the toy car, person, and the balls make a spatial advance, their locomotion measurements (red, blue, purple, and orange, respectively) are large, while the pendulum (green) has a low locomotion measure since its overall displacement is zero.

### 3.3.2 Measure of Oscillation, $\mathcal{M}_\Omega$

The amount of oscillation exhibited by an object can be quantitatively measured by considering the behaviour of the entire trajectory (as opposed to the centroid of the tracked points). In particular, the oscillatory behaviour induced by the motion of a localized body part (e.g., a limb) can be captured by considering its amplitude, where large amplitudes generally provide a stronger evidence of oscillation; cf. locomoting biological species and locomoting non-biological objects in Figure 3.2. Thus, to obtain a **measure of oscillation**,  $\mathcal{M}_\Omega$ , an aggregate of trajectory amplitudes are considered.

To calculate the amplitude,  $a^k$ , of trajectory  $\mathcal{T}^k$ , the trajectory must be detrended to  $\tilde{\mathcal{T}}^k$  such that the calculated amplitude is independent of the spatial direction the trajectory unfolds over time; see Figure 3.3. Once the trajectory of interest has been detrended, its amplitude can be computed in various ways: (i) frequencies, (ii) integrals, or (iii) differentials [60]. Based on preliminary experimentation, a frequency-based approach is employed; see Appendix B for further discussion of integral- and differential-based approaches. The amplitude of a detrended trajectory,  $\tilde{\mathcal{T}}^k$ , can be calculated using *frequencies* by obtaining the maximum magnitude of the Fourier Transformation of  $\tilde{\mathcal{T}}^k$ . That is,

$$a_{\mathcal{F}}^k = \max \left\{ \mathcal{F}(\tilde{\mathcal{T}}^k) \right\}, \quad (3.11)$$

where  $\mathcal{F}(\cdot)$  denotes the Fourier Transform. One maximum value is sufficient since

considered trajectories only span a few (i.e., 15) frames to compensate for the drifting effect.

With the amplitude estimated for each detrended trajectory, the overall measure of oscillation at frame  $t$ ,  $\mathcal{M}_\Omega(t)$ , is calculated by considering the weighted mean of the amplitudes. That is, for a detrended trajectory  $\tilde{\mathcal{T}}^k(t)$ , the measurement of oscillation is defined as

$$\mathcal{M}_\Omega(t) \equiv \frac{w_\nu(t)}{|A(t)|} \sum_{a^k(t) \in A(t)} a^k(t), \quad (3.12)$$

where  $a^k$  is the amplitude of trajectory  $\tilde{\mathcal{T}}^k$ ,  $A(t)$  is a set of amplitudes considered at frame  $t$ ,  $|A(t)|$  is the number of amplitudes considered at frame  $t$ , and  $w_\nu(t)$  is a weight assigned to account for the percentage of total trajectories present in a given frame, as follows.

To ensure that reliable data contributes to the calculation of  $\mathcal{M}_\Omega(t)$ , only amplitudes within 1.5 of the standard deviation from the mean are considered in  $A(t)$ , analogous to the computation of centroid of an object,  $\mathbf{x}^c(t)$ , as in (3.6). That is,

$$A(t) = \{a^k(t) \mid \mu_a(t) - 1.5\sigma_a(t) \leq a^k(t) \leq \mu_a(t) + 1.5\sigma_a(t)\}, \quad (3.13)$$

where  $\mu_a(t)$  and  $\sigma_a(t)$  are the mean and standard deviation, respectively, of the amplitudes present at frame  $t$ .

It is desirable to avoid having frames with exceptionally small number of trajectories unduly bias the overall measure of oscillation,  $\mathcal{M}_\Omega(t)$ . Thus,  $\mathcal{M}_\Omega$  is weighted by  $w_\nu(t)$ , which is defined based on the amount of data present at frame  $t$ ,  $|A(t)|$ , relative to the typical amount of data that is available in a frame across the entire video. A sigmoid function can be used to model a fair distribution based on the amount of data according to

$$w_\nu(t) = \frac{1}{1 + \exp[-(|A(t)| - |\mu_\nu - \sigma_\nu|)]}, \quad (3.14)$$

where  $\mu_\nu$  and  $\sigma_\nu$  are the mean and standard deviation, respectively, of the amount of data in a given set. Thus, a unit weight is assigned when there are more trajectories present at frame  $t$  relative to the average across the video, since it suggests that the frame contains a reliable amount of data. On the contrary, frames with exceptionally small number of trajectories are assigned a low weight through the sigmoid's saturation to prevent a few exceptional data points from dominating the final outcome of the overall measure,  $\mathcal{M}_\Omega$ .

It can be observed in Figure 3.6d that the oscillation measure,  $\mathcal{M}_\Omega$ , captures the amount of oscillation that is exhibited by each object in motion. Specifically, the toy car with near linear trajectories (red) has very low  $\mathcal{M}_\Omega$  values, while the person (blue), pendulum (green), and the rolling and bouncing balls (purple and orange, respectively) have relatively large  $\mathcal{M}_\Omega$  values. Furthermore, as the pendulum slows down to a gradual stop, its oscillation values also approach zero.

### 3.3.3 Measure of Asymmetry, $\mathcal{M}_\Sigma$

The resistive forces that species must battle as they move through an environment can be captured through an asymmetric motion trace these species exhibit during their path of motion (cf. locomoting biological species and rolling ball in Figure 3.2); this amount is quantitatively measured by considering the behaviour of the entire trajectory (similar to the measure of oscillation). In particular, the resistive force that must be battled can be calculated by measuring the magnitude of asymmetry of the individual trajectories, the **measure of asymmetry**,  $\mathcal{M}_\Sigma$ , where presence of asymmetry is more indicative of biolocomotion.

A variety of approaches can be employed to measure asymmetry: (i) direct calculation of vertical acceleration, (ii) skewness [15] of the trajectory, and (iii) comparison of the left and right areas under the trajectories from its highest peak. Preliminary investigations showed difficulties in reliably estimating asymmetry via direct calculation of vertical acceleration, by considering the second-derivative of the vertical components of the trajectories, as computational instability was encountered and amplified for each considered derivative. Similar challenges appeared in estimating asymmetry via skewness, as the third-order moment statistic was very sensitive to the presence of noise in the trajectories. Thus, the magnitude of asymmetry,  $s^k$ , of a detrended trajectory,  $\tilde{\mathcal{T}}^k$ , is quantified by finding its local maximum valued point and comparing the integrals from that value to its left and right minimal valued points.



Since the recovered trajectories traces the path of the object from frame  $t$  to  $t + L^k$ , the recovered trajectory is not guaranteed to contain a *full cycle* of the tracked motion from the lowest point to its next lowest point (e.g., stance-swing-stance phases of a walk). Rather, it is likely to track a cycle of the motion at random points (e.g., latter half of swing-stance-former half of swing phases of a walk). To increase the likelihood that the considered trajectory will always contain the desired information when computing area-under-the-curve (e.g., stance-swing-stance), the input trajectories are replicated. Subsequently, any part of the replicated trajectory is representative of the stance-swing-stance cycle; therefore, it can be used to extract asymmetry information. Furthermore, the detrended trajectory is vertically shifted such that its vertical minimum is zero to ensure the calculated area is non-negative. Then the integrals from the (local) maximum valued point is compared to its left and right minimal valued points. More precisely, let

$$l_L = \arg \min \tilde{y}^k \text{ and } l_R = \arg \min \tilde{y}^k \quad (3.15)$$

be the left and right minimal points of  $\tilde{\mathcal{T}}^k$ , respectively, for

$$l_{\max} = \arg \max \tilde{y}^k \quad (3.16)$$

such that  $l_L < l_{\max} < l_R$ ; see Figure 3.8. Then, the asymmetry magnitude of  $\tilde{\mathcal{T}}^k$  is computed by comparing the discrete approximation of the integrals from the left

and right of the highest peak as

$$s^k = \left| \sum_{l=l_L}^{l_{\max}} [\tilde{y}_l^k + \min \tilde{y}^k] - \sum_{l=l_{\max}}^{l_R} [\tilde{y}_l^k + \min \tilde{y}^k] \right|, \quad (3.17)$$

where  $\min \tilde{y}^k$  is the vertical minimum of  $\tilde{\mathcal{T}}^k$  and serves to shift vertically the  $\tilde{y}^k$ , so that the calculated values are non-negative. A large asymmetric magnitude,  $s^k$ , is indicative of biolocomotion.




Figure 3.8: Calculation of magnitude of asymmetry. The magnitude of asymmetry of a detrended trajectory  $\tilde{\mathcal{T}}^k$  can be computed by comparing the discrete approximation of integrals from the left of the highest peak (orange) and right of the highest peak (green) of a vertically shifted detrended trajectory such that its vertical minimum is 0.

To ensure that the measure of asymmetry is invariant to the stride length of the object,  $s^k$  is normalized by the arc length of  $\tilde{\mathcal{T}}^k$ ,  $\|\tilde{\mathcal{T}}^k\|$  as in (3.5), as

$$\tilde{s}^k = \frac{s^k}{\|\tilde{\mathcal{T}}^k\|}. \quad (3.18)$$

While normalizing the amplitude by the arc length could also make the measure of

oscillation, (3.12), invariant to the stride length of the object's motion, such normalization retracts informative data. Loss of information when normalizing amplitude by arc length can be understood through an analogy to various parts of a circle. That is, suppose the radius of a circle is to the amplitude of a detrended trajectory as circumference is to arc length, and area is to area-under-the-curve. Furthermore, let the radius of a circle be  $r$  such that its circumference is  $2\pi r$  and area is  $\pi r^2$ . Then normalizing the radius (or amplitude) by the circumference (or arc length) yields

$$\frac{r}{2\pi r} = \frac{1}{2\pi}, \quad (3.19)$$

where informative data,  $r$ , is lost. Normalizing the area (or area-under-the-curve) by the circumference (or arc length), however, yields

$$\frac{\pi r^2}{2\pi r} = \frac{r}{2}, \quad (3.20)$$

where informative data  $r$  is multiplied with a constant.

Similar to the measure of oscillation, the measure of asymmetry at frame  $t$ ,  $\mathcal{M}_\Sigma(t)$ , is obtained by aggregating the asymmetry magnitudes. Specifically, it is calculated by considering the weighted mean of the asymmetry magnitudes according to

$$\mathcal{M}_\Sigma(t) \equiv \frac{w_\nu(t)}{|S(t)|} \sum_{\tilde{s}^k \in S(t)} \tilde{s}^k, \quad (3.21)$$

where  $\tilde{s}^k$  is the normalized magnitude of asymmetry of  $\tilde{\mathcal{T}}^k$ ,  $S(t)$  is a trimmed set

of asymmetric values that are present at frame  $t$  analogous to (3.13),  $|S(t)|$  is its cardinality, and  $w_\nu$  is a weight that accounts for the percent of total asymmetry values considered in a given frame exactly analogous to the weights applied to the aggregation of oscillation measures as in (3.14).

It can be observed in Figure 3.6e that the person (blue) and bouncing ball (orange) exhibit large asymmetry measures,  $\mathcal{M}_\Sigma$ , as they have to fight against gravity during its lift phase, while the car (red), pendulum (green), and rolling ball (purple) do not. As a result,  $\mathcal{M}_\Sigma$  is lower for the toy car, pendulum, and rolling ball.

### 3.3.4 Measure of Extrinsic Motion Dissimilarity, $\mathcal{M}_E$

The biolocomotion detector can benefit from emphasizing the dissimilarity between the overall path of the object (extrinsic motion) and the trace exhibited by the individual body parts (intrinsic motion)<sup>2</sup>; cf. locomoting biological objects and bouncing ball in Figure 3.2.

The amount of deviation between extrinsic and intrinsic motion exhibited by an object, the **measure of extrinsic motion dissimilarity**,  $\mathcal{M}_E$ , is measured by comparing the dissimilarity between the overall path of the object (i.e., centroid) and the individual trajectory shape. In particular, the angle between the displacement vectors of a (non-detrended) trajectory  $\mathcal{T}^k$  and the centroid are compared. Suppose

---

<sup>2</sup>In the psychophysical literature, *intrinsic motion* is obtained by removing *common (extrinsic) motion* from the observed motion, via subtraction [24]. However, such subtraction limits the ability to compare common and relative motion. Thus, the individual traces exhibited by the body parts are used without the subtraction of common (extrinsic) motion for such comparison in this dissertation.

$\Delta \mathbf{x}_{(t_{ext}, \Delta t_{ext})}^k$  denotes a displacement vector of trajectory  $\mathcal{T}^k$  between frames  $t_{ext} - \Delta t_{ext}$  and  $t_{ext}$  as in (3.4), while  $\Delta \mathbf{x}_{(t_{ext}, \Delta t_{ext})}^c$  denotes a displacement vector of the centroid between frames  $t_{ext} - \Delta t_{ext}$  and  $t_{ext}$  as in (3.9), where  $t_{ext}$  is a frame at which a vertical extremum is reached by  $\mathcal{T}^k$  and  $\Delta t_{ext}$  is the temporal separation between vertical extrema; see Figure 3.9. Then the extrinsic motion dissimilarity at frame  $t_{ext}$  is calculated using the normalized inner product between extrinsic and intrinsic displacement vectors from frame  $t_{ext} - \Delta t_{ext}$  to  $t_{ext}$  according to

$$\begin{aligned} e_{(t_{ext}, \Delta t_{ext})}^k &= 1 - \frac{1}{2} \left( \frac{\langle \Delta \mathbf{x}_{(t_{ext}, \Delta t_{ext})}^k, \Delta \mathbf{x}_{(t_{ext}, \Delta t_{ext})}^c \rangle}{\| \Delta \mathbf{x}_{(t_{ext}, \Delta t_{ext})}^k \| \| \Delta \mathbf{x}_{(t_{ext}, \Delta t_{ext})}^c \|} + 1 \right) \\ &= \frac{1}{2} \left( 1 - \frac{\langle \Delta \mathbf{x}_{(t_{ext}, \Delta t_{ext})}^k, \Delta \mathbf{x}_{(t_{ext}, \Delta t_{ext})}^c \rangle}{\| \Delta \mathbf{x}_{(t_{ext}, \Delta t_{ext})}^k \| \| \Delta \mathbf{x}_{(t_{ext}, \Delta t_{ext})}^c \|} \right). \end{aligned} \quad (3.22)$$

The normalized inner product in (3.22) is increased by 1 then multiplied by a factor of  $\frac{1}{2}$  to ensure the considered inner product ranges between 0 and 1 instead of  $[-1, 1]$ . The range is altered via shift followed by a multiplicative constant instead of considering the absolute of the normalized inner product for the following reason. To reduce angular momentum of the body during locomotion, animals often accompany their (extrinsic) overall positional advance with (intrinsic) motion of body parts in the opposite direction (e.g., backward swing of arms during walking). Correspondingly, for present purposes, it is beneficial that antiparallel extrinsic vs. intrinsic motion be distinguished from parallel. Simply taking the absolute value of the normalized inner product fails to make this distinction, while formula (3.22) does. Further, the normalized inner product is subtracted from 1 to ensure a large corresponding value is indicative of biolocomotion, similar to other computed mea-

asures for biolocomotion (e.g., large amplitude values are indicative of biolocomotion). Finally, the measure of extrinsic motion dissimilarity for trajectory  $\mathcal{T}^k$  is obtained by considering the weighted combination of all normalized inner products of extrinsic and intrinsic displacement vectors according to

$$e^k = \frac{\sum_{t_{ext}} \Delta t_{ext} e^k_{(t_{ext}, \Delta t_{ext})}}{\sum_{t_{ext}} \Delta t_{ext}}. \quad (3.23)$$

The temporal separation between two extreme points,  $\Delta t_{ext}$ , is used as a weight, such that the difference between the considered extrinsic and intrinsic displacements are maximized. For example, comparing the extrinsic and intrinsic displacement vectors with  $\Delta t = 1$  is likely to be more alike than vectors with  $\Delta t = 15$  since the centroid is calculated from the motion of the object. This manipulation is analogous to discrete approximation for estimating the derivative of a mathematical function. A more accurate estimation of the slope is obtained when the discretization ( $\Delta t$ ) approaches zero. Similarly, extrinsic and intrinsic displacement vectors will be more alike as  $\Delta t \rightarrow 0$ , while the goal is to determine how different they are.

Similar to the measurement of oscillation and asymmetry, the measure of extrinsic motion dissimilarity,  $\mathcal{M}_E$ , is obtained by aggregating the dissimilarities of extrinsic and intrinsic values,  $e^k$ , via a weighted mean. That is, the measurement of extrinsic motion dissimilarity at frame  $t$ ,  $\mathcal{M}_E(t)$ , is defined as

$$\mathcal{M}_E(t) \equiv \frac{w_\nu}{|E(t)|} \sum_{e^k \in E(t)} e^k, \quad (3.24)$$




Figure 3.9: Calculation of extrinsic and intrinsic motion dissimilarity. Given trajectory  $\mathcal{T}^k$  (light red curve) from frame  $t_1$  to  $t_4$ , where vertical extreme points,  $\mathbf{x}(t_1), \dots, \mathbf{x}(t_4)$  (red dots), are achieved at frames  $t_1, \dots, t_4$ , the coordinates of the centroid at its respective frames are given by  $\mathbf{x}^c(t_1), \dots, \mathbf{x}^c(t_4)$  (blue dots). Correspondingly, the displacement vector of trajectory  $\mathcal{T}^k$  between extreme points  $\mathbf{x}^k(t_2)$  and  $\mathbf{x}^k(t_3)$  is given by  $\Delta \mathbf{x}_{(t_3, t_3-t_2)}^k$  (red solid arrow) and the displacement vector of the centroid at its corresponding frames is given by  $\Delta \mathbf{x}_{(t_3, t_3-t_2)}^c$  (blue solid arrow). The extrinsic motion dissimilarity at frame  $t_2$ ,  $e_{(t_3, t_3-t_2)}^k$ , is obtained by computing the normalized inner product of  $\Delta \mathbf{x}_{(t_3, t_3-t_2)}^k$  and  $\Delta \mathbf{x}_{(t_3, t_3-t_2)}^c$ . Finally, the measure of extrinsic motion for trajectory  $\mathcal{T}^k$ ,  $e^k$ , is obtained by computing the weighted sum of the normalized inner products between extrinsic and intrinsic displacement vectors at all extreme points  $t_1, \dots, t_4$  (i.e.,  $e_{(t_2, t_2-t_1)}^k$ ,  $e_{(t_3, t_3-t_2)}^k$ , and  $e_{(t_4, t_4-t_3)}^k$ ).

where  $e^k$  is the measurement of dissimilarity between extrinsic and intrinsic motion of trajectory  $\mathcal{T}^k$ ,  $E(t)$  is a trimmed set of dissimilarity calculations of trajectories present at frame  $t$  analogous to (3.13),  $|E(t)|$  is its cardinality, and  $w_\nu$  is a weight that accounts for the percent of total trajectories present in a given frame analogous to (3.14).

It can be observed in Figure 3.6f that the measurement of extrinsic motion dissimilarity is relatively large when the person, pendulum, and rolling ball are in motion

(blue, green, and purple, respectively), while the car and bouncing ball (red and orange, respectively) have low measures since their extrinsic and intrinsic motions coincide.

### 3.3.5 Biocomotion Detector, $S_B$

Once the raw measurements for each of the biocomotion components have been computed, they must be normalized into a common range such that each component has equal contribution to the overall measure of biocomotion. Then they must be combined into a single measure.

Each raw measurement,  $\mathcal{M}_\chi$  for  $\chi \in \{\Lambda, \Omega, \Sigma, E\}$ , is normalized to range between  $[0, 1]$ . This conversion allows each measure to be treated as a confidence score. The raw measurement is converted into a confidence score using a sigmoidal function according to

$$S_\chi(t) = S(\mathcal{M}_\chi(t); \tau_\chi, k_\chi) \equiv \frac{1}{1 + \exp[-k_\chi(\mathcal{M}_\chi(t) - \tau_\chi)]} \quad (3.25)$$

for some values  $\tau_\chi$  and  $k_\chi$ . The sigmoid function restricts the range to the desired interval  $[0, 1]$ , with a smooth asymptotic behaviour as the extreme values are approached.

Treating the scores of locomotion, oscillation, asymmetry, and extrinsic motion akin to independent probabilities, the overall biocomotion confidence score is ob-



tained by taking the product of each component's confidence scores

$$S_B(t) = \prod_{\chi \in X(t)} S(\mathcal{M}_\chi(t); \tau_\chi, k_\chi), \quad (3.26)$$

where  $X(t)$  is a set of measurements to be considered in the final biolocomotion calculation.

Since asymmetric motion trace is most applicable when objects move orthogonal to a resistive force (e.g., gravity), the measurement of asymmetry,  $\mathcal{M}_\Sigma$ , is constrained to situations when objects move in the orthogonal direction. For example, as the foot pushes off against gravity for terrestrial animals, it results in a rapid rising trajectory, which is followed by a swing phase that moves entirely under the influence of gravity to yield a more elongated trace; e.g., see [24, 55]. This phenomenon is less apparent when motion is parallel to gravity and similar pattern holds for other animals; e.g., see [7, 9, 42]. Assuming the resistive forces are aligned with the image vertical, motion orthogonal to that direction,  $\mathbf{r} = (1, 0)$ , should show an asymmetric trace. To capture the motion direction, the normalized inner product of the centroid displacement vector,  $\Delta \bar{\mathbf{x}}_{(t, \Delta t)}^c$  as in (3.9), and a reference vector,  $\mathbf{r}$ , is used to determine if the general direction of the object motion is perpendicular to gravity. More specifically, the average absolute value of the normalized inner products within a temporal window is calculated as

$$w_\perp(t) = \frac{1}{k} \sum_{\Delta t=1}^k \left\{ \left| \frac{\langle \Delta \mathbf{x}_{(t, \Delta t)}^c, \mathbf{r} \rangle}{\|\Delta \mathbf{x}_{(t, \Delta t)}^c\| \|\mathbf{r}\|} \right| \right\}, \quad (3.27)$$

where the absolute value of the normalized inner product maintains invariance to the direction of motion (left vs. right or toward vs. against gravity). Since we want to evaluate the *general* path of the object, displacement vectors comparing centroid at frame  $t$  to numerous centroids in previous frames are compared (i.e.,  $\Delta t = [1, k]$ , where  $k$  is empirically set to the frame rate of the video); see Figure 3.10. Thus,

$$X(t) = \begin{cases} \{\Lambda, \Omega, \Gamma\}, & \text{if } w_{\perp}(t) < \tau_{\perp} \\ \{\Lambda, \Omega, \Sigma, \Gamma\}, & \text{if } w_{\perp}(t) \geq \tau_{\perp} \end{cases}, \quad (3.28)$$

where  $\tau_{\perp}$  is empirically set to 0.9.




Figure 3.10: Computing displacement vectors for various temporal windows. To determine the *direction* of the object, displacement vectors (blue arrows) across  $k + 1 = 4$  time steps,  $t_1, \dots, t_4$ , are considered along the path of the object (light blue curve). Each displacement vector,  $\Delta \mathbf{x}_{(t_4, t_4-t_1)}^c, \dots, \Delta \mathbf{x}_{(t_4, t_4-t_3)}^c$ , is compared with a reference vector,  $\mathbf{r} = (1, 0)$  (black arrow), via a normalized inner product. The absolute value of these differences are averaged to determine the overall direction of the object at frame  $t_4$ ,  $w_{\perp}(t_4)$ .

By construction,  $S_B(t)$  is in the range of  $[0, 1]$  with large values more indicative of biolocomotion. As the product of the confidence scores for each component are taken, objects that lack locomotion, oscillation, asymmetry, or difference in extrinsic and intrinsic motion (e.g., pendulum, car, rolling ball, and bouncing ball) possess

low  $S_B$  values, while objects that exhibit locomotion, asymmetric oscillation, and difference in extrinsic and intrinsic motion (e.g., walking person) retain large  $S_B$  values; see Figure 3.6g. Hence,  $S_B$  (3.26) is highly indicative of biolocomotion.

### 3.4 Adaptive Sliding Temporal Window

Since the input to the biolocomotion detector is a cluster of trajectories, it is not necessary to process an entire video at once. Instead, an adaptively defined sliding temporal window is used to allow for incremental video processing. More specifically, the video is temporally segmented into disjoint sets of frames and each segment is processed successively. The segment lengths are defined adaptively to break the video into temporal extents, where the overall frame-to-frame image motion maintains the same coarse direction. Such subdivision provides a natural way to break the video into segments dominated by a single direction of camera motion (or in the absence of camera motion) with a single overall direction of object motion, e.g., the overall locomotion (3.10). The overall direction of motion for frame  $t$  is determined by using displacement vectors that are present at that frame. Specifically, for displacement vector  $\Delta \mathbf{x}^k(t)$  of trajectory  $\mathcal{T}^k$  at frame  $t$  as in (3.4), the direction of the displacement for  $\mathcal{T}^k$  is defined as

$$\theta^k(t) = \arctan \left( \frac{\Delta y^k(t)}{\Delta x^k(t)} \right), 0^\circ \leq \theta^k(t) \leq 360^\circ. \quad (3.29)$$

Then, the overall displacement direction is calculated by taking the average of  $\theta^k(t)$  according to

$$\Theta(t) = \frac{1}{N(t)} \sum_k \theta^k(t), \quad (3.30)$$

where  $N(t)$  is the number of trajectories. The **adaptive Sliding Temporal Windows** (aSTW) are then defined to partition the video at times where  $\Theta(t)$  and  $\Theta(t+1)$  differ significantly (i.e., if  $|\Theta(t) - \Theta(t+1)| > \tau_\theta$ , then segment). Empirically,  $\tau_\theta$  is set to  $\frac{3}{2}f$  for frame rate  $f$ .

### 3.5 Algorithmic Summary

Algorithm 3.1 provides a summary of the overall approach to biolocomotion detection. Note that the approach both localizes and labels regions of biolocomotion of an input video.

```

Input : Video,  $V$ 

Output: (i) Biocomotion confidence score,  $S_B$  (3.26), for each recovered
           trajectory cluster,
           (ii) the spatiotemporal location of the centre of each cluster, and
           (iii) the horizontal and vertical extents of the cluster at each frame

1 Extract and prune iDTs (§3.2.2 and A.1);
2 Estimate average direction of motion for a crude division of temporal
  windows,  $aSTW$  (§3.4);
3 for  $aSTW \in V$  do
4   Cluster (A.2), stabilize (A.3), and elongate trajectories (A.4);
5   for  $cluster$  do
6     Measure locomotion,  $\mathcal{M}_\Lambda$  (§3.3.1);
7     Measure oscillation,  $\mathcal{M}_\Omega$  (§3.3.2);
8     Measure asymmetry,  $\mathcal{M}_\Sigma$  (§3.3.3);
9     Measure extrinsic motion dissimilarity,  $\mathcal{M}_E$  (§3.3.4);
10    Measure biocomotion,  $S_B$  (§3.3.5) ;
11  end
12 end

```

**Algorithm 3.1:** Summary of the proposed biocomotion detector algorithm

# Chapter 4

## Empirical Evaluation

### 4.1 Introduction

Quantitative analysis of an algorithm can assist in the development and understanding of its strengths and weaknesses. A benchmark dataset can provide a way for comparative analysis with other algorithms or identification of important/unimportant components within an algorithm. To provide a benchmark for biolocomotion detection in videos, an extant dataset from a related field is exploited and a new dataset is also constructed to further test the robustness of the algorithm of interest.

This chapter unfolds in six sections. This section has served to motivate the need for empirical evaluation. Section 4.2 describes the datasets considered for the task of biolocomotion. Section 4.3 describes the performance metrics used to quantitatively compare different components within an algorithm and across algorithms.

Section 4.4 describes how the set of parameters used for the presented algorithm is determined. Section 4.5 presents an alternative approach to biolocomotion detection that is based on generic handcrafted features combined with learning. Since it appears that no previous algorithms for biolocomotion detection have been developed, this alternative approach provides a basis of comparison for the main approach presented in Chapter 3. Section 4.6 provides qualitative and quantitative results along with discussions comparing individual components of the proposed and alternative approaches as well as more general comparisons between the two approaches. Finally, Section 4.7 provides an overall summary of the evaluations.

## 4.2 Datasets

While there are many benchmark datasets in the field of computer vision that can be used for quantitative evaluation of various detection algorithms in the context of humans (e.g., CMU Crowded Videos [62], UCF Sports [98, 105], UCF101 [106], ActivityNet [48], J-HMDB [53], etc.), there are none available for biolocomotion detection. Consequently, a novel dataset, Biological Object in Locomotion Detection (BOLD) dataset, is developed and biolocomotion annotations for an extant camouflaged animals dataset (CAD) [8] are provided. This section contains a description of the considered datasets as well as biolocomotion annotation.

## 4.2.1 Biological Objects in Locomotion Detection (BOLD)

### Dataset

The Biological Objects in Locomotion Detection (**BOLD**) dataset builds on an extant dataset, A2D [126]. A2D is taken as a starting point as it provides diversity of objects in motion (i.e., more than humans) and also includes labelled regions where they are present. However, it lacks generality to validate biolocomotion detection. Specifically, A2D does not consider undulation nor swim as one of its actions and its species diversity remains too limited for present purposes. Thus, a richer set that contains more variety of objects and locomotion types to provide stronger justification of a biolocomotion detection algorithm is needed. Consequently, the BOLD dataset is constructed to provide greater diversity in terms of object type and modes of locomotion. BOLD builds on a subset of videos from A2D and is substantially augmented with various videos from YouTube. It expands the diversity of objects by including: *reptiles, cetaceans, seals, fish, stingray, eel, sea snakes, snakes, insects, spiders, scorpions, lobsters, trains, motorbikes, submarines, airplanes, helicopters, rockets, metronomes, and pendulums* and modes of locomotion by adding: *swim, undulate, and oscillate*; a comparison of A2D and BOLD is visualized in Figure 4.1 and a detailed breakdown is provided in Table 4.1.

BOLD contains 1,348 videos split into 1,078 training and 270 test videos. Videos range between 22 – 504 frames and on average span 143.98 frames. Bounding box annotations are provided for at least a frame per video, up to 18 frames in longer videos; see Figure 4.2.



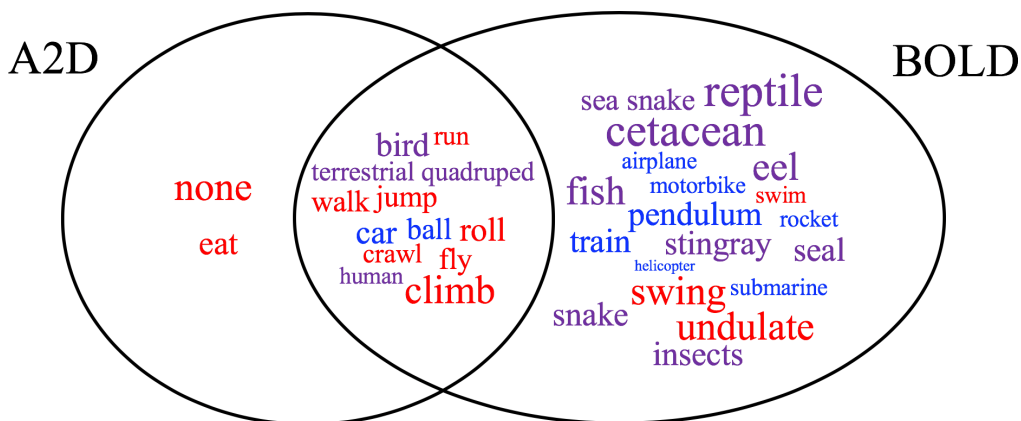




Figure 4.1: Venn diagram of motions and species that are contained in A2D and BOLD. The wordle in each component of the Venn diagram illustrates a set of actions (red), biological objects (purple), and non-biological objects (blue) that are contained in  $A2D \setminus BOLD$ ,  $A2D \cap BOLD$ , and  $BOLD \setminus A2D$ .

		Motion Type										
Object Type		climb	crawl	fly	jump	roll	run	swim	undulate	walk	locomotion	swing
Biological Species	human (adult, baby)	33	27	-	36	33	33	20	-	31	-	-
	terrestrial quadruped mammal (cat, dog)	40	25	-	37	18	67	-	-	35	-	-
	bird (incl. penguin)	35	-	100	33	10	-	-	-	26 + 8	-	-
	reptile (alligator, chameleon, crocodile, lizard, turtle)	-	25	-	-	-	-	15	-	-	-	-
	cetacean (dolphin, shark, whale)	-	-	-	-	-	-	23	-	-	-	-
	seal	-	-	-	-	-	-	18	-	-	-	-
	fish	-	-	-	-	-	-	20	-	-	-	-
	stingray	-	-	-	-	-	-	20	-	-	-	-
	eel	-	-	-	-	-	-	-	20	-	-	-
	sea snake	-	-	-	-	-	-	-	20	-	-	-
	snake	-	-	-	-	-	-	-	63	-	-	-
	insects, spiders, scorpions, lobster	-	25	-	-	-	-	-	-	-	-	-
	Non-biological Objects	ball	-	-	25	50	50	-	-	-	-	-
car		-	-	-	50	50	-	-	-	-	25	-
train		-	-	-	-	-	-	-	-	-	33	-
motorbike		-	-	-	-	-	-	-	-	-	33	-
submarine		-	-	-	-	-	-	25	-	-	-	-
airplane		-	-	27	-	-	-	-	-	-	9	-
helicopter		-	-	25	-	-	-	-	-	-	-	-
rocket		-	-	25	-	-	-	-	-	-	-	-
oscillating stuff (metronome, pendulum, boat)		-	-	-	-	-	-	-	-	-	-	25
	Total (Biological)	108+0	52 + 50	100 + 0	106 + 0	61 + 0	100 + 0	0 + 116	0 + 103	92 + 8	-	-
	Total (Non-biological)	-	-	25 + 77	100 + 0	100 + 0	-	0 + 25	-	-	25 + 75	25
												619 + 277
												250 + 202


Table 4.1: Summary of the number of videos in each motion and object class for the BOLD dataset. The numbers indicate the number of videos in each object-motion class, hyphen (-) indicates lack of videos in the respective categories. Numbers in orange indicate videos extracted from A2D, and green indicates new videos that were obtained from YouTube.



(a) Humans locomoting in various forms (e.g., climb, swim, roll, and walk).



(b) Variations in biological species (e.g., snake, turtle, terrestrial quadruped, and bird).



(c) Various non-biological objects in locomotion (e.g., airplane, motorcycle, ball, and submarine).

Figure 4.2: Illustration of videos (top) in BOLD and its biolocomotion annotations indicated with overlaid red boxes (bottom). Select frames of videos from BOLD demonstrating variation in modes of locomotion exhibited by (a) humans, (b) other biological species, and (c) non-biological objects.

## 4.2.2 Camouflage Animals Dataset (CAD)

The Camouflage Animals Dataset (CAD) [8] was created for motion segmentation with the challenge of detecting camouflaging animals in motion. It contains nine videos of animals in motion: *chameleon*, *frog*, *glowworm beetle*, 4 *scorpions*, *snail*, and *stick insect*. All videos in CAD are test videos. As CAD lacks training data, transfer learning is applied, where training data from BOLD is used to test videos in CAD. Videos on average span 92.89 frames and range between 30 – 218 frames. Dense-pixel level annotations are provided for at least 12 frames; see Figure 4.3.




Figure 4.3: Illustration of select videos (top) in CAD [8] and its fine-grained biolocomotion annotations overlaid in red on its respective raw image frames (bottom).


## 4.2.3 Biolocomotion Annotations

To determine how well a proposed algorithm performs in the task of biolocomotion detection, each annotated region is supplied with a biolocomotion label. Thus, each object is spatially identified, via bounding boxes (i.e., BOLD) or fine-grained segmentation (i.e., CAD), for select frames and categorized into its appropriate class


(i.e., biolocomotion or non-biolocomotion). The annotated regions serve as the set of positives (or negatives) for ensuing experiments. While the main focus is placed on biolocomotion, finer grained annotations are provided for more general use. In particular, all labelled objects are identified as being in one of six categories:

- (i) biological objects in locomotion (i.e., biolocomotion),
- (ii) biological objects that are oscillating but not locomoting,
- (iii) biological objects that are neither locomoting nor oscillating,
- (iv) non-biological objects that are only locomoting,
- (v) non-biological objects that are only oscillating, and
- (vi) non-biological objects that are neither locomoting nor oscillating.


Visualization of these classes can be seen in Figure 4.4.



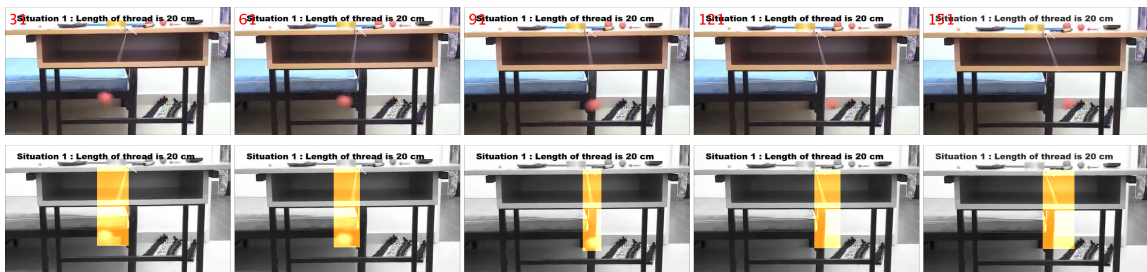
(a) biocomotion (red) and biological objects not in locomotion nor oscillation (blue).



(b) biological object in oscillation.



(c) non-biological object in locomotion (purple) and non-biological object neither in locomotion nor oscillation (yellow).



(d) non-biological object in oscillation (orange).

Figure 4.4: Biocomotion annotation illustrations. Select videos demonstrating the provided annotations with overlaid boxes: biocomotion (red), biological objects in oscillation (green), biological objects not in locomotion nor oscillation (blue), non-biological objects in locomotion (purple), non-biological objects in oscillation (orange) and non-biological objects in neither locomotion nor oscillation (yellow).

Here, locomotion corresponds to non-zero displacement exhibited by the object in the real-world. For example, a cat (biological object) walking (action related to locomotion) on a treadmill results in zero displacement in the real-world, thus is classified as non-biolocomotion.

In addition to the six object-action labels, each video also has an associated label for camera motion according to five categories:

- (i) translation,
- (ii) rotation,
- (iii) zoom,
- (iv) more than one of the above, and
- (v) no camera motion.

These labels could be used in future work to study algorithm performance as a function of camera motion.

#### **4.2.4 Summary**

Table 4.2 provides a summary of the key features of the considered datasets for the task of biolocomotion detection. Also, for the sake of size consistency, all videos are resized to a fixed height of 320 pixels and the widths are adjusted accordingly to maintain the original aspect ratio.

<b>Dataset</b>	<b>BOLD</b>	<b>CAD [8]</b>
<b>Task</b>	biolocomotion detection	causal motion segmentation
<b>Source</b>	YouTube	
<b>Objects</b>	human, terrestrial quadruped, bird, reptile, cetacean, seal, fish, stingray, eel, sea snake, insects, spiders, scorpion, lobster, ball, car, train, motorbike, submarine, airplane, helicopter, rocket, oscillating stuff	chameleon, frog, glowworm beetle, scorpion, snail, stick insect
<b>Number of videos</b>	1,348 videos <ul style="list-style-type: none"> <li>• 1,078 training</li> <li>• 270 test</li> </ul>	9 videos <ul style="list-style-type: none"> <li>• 0 training</li> <li>• 9 test</li> </ul>
<b>Duration</b>	<ul style="list-style-type: none"> <li>• 22 – 504 frames</li> <li>• avg. 143.98 frames</li> </ul>	<ul style="list-style-type: none"> <li>• 30 – 218 frames</li> <li>• avg. 92.89 frames</li> </ul>
<b>Number of videos with biolocomotion</b>	882	9
<b>Challenges</b>	Variations in viewpoint, camera motion present, background and foreground clutter present	
<b>Groundtruth labelling</b>	<ul style="list-style-type: none"> <li>• bounding boxes</li> <li>• 1 – 18 frames</li> <li>• avg. 4.03 frames</li> </ul>	<ul style="list-style-type: none"> <li>• fine-grained</li> <li>• 12 – 43 frames</li> <li>• avg. 21.22 frames</li> </ul>

Table 4.2: Summary of the datasets.

### 4.3 Evaluation Metrics

To quantitatively evaluate the algorithms of interest, two standard detection metrics are applied to the test sets. To evaluate the value of each component of an algorithm, detection results are plotted as precision-recall (PR) curves. To compare between opposing algorithms, average precision (AP) scores are reported. In the following, detection results (e.g., cluster outputs with biolocomotion score  $S_B$  from Algorithm 3.1) are referred to as *biolocomotion proposals*, which are regions likely to contain biolocomotion, associated with confidence scores.

Precision-recall (**PR**) curves are obtained by calculating precision and recall values at various confidence thresholds. *Precision* is the percentage of correctly assigned pixels, which is defined as

$$precision = \frac{TP}{TP + FP}, \quad (4.1)$$

where TP denotes the number of correctly predicted pixels with respect to the groundtruth (i.e., true positive) and FP denotes the number of incorrectly predicted pixels with respect to the groundtruth (i.e., false positive). *Recall* is the percentage of detected pixels with respect to the labelled groundtruth pixels, which is defined as

$$recall = \frac{TP}{TP + FN}, \quad (4.2)$$

where FN denotes the number of incorrectly missed pixels (i.e., false negative). A set of pixels within the proposal are considered a positive, which outputs a mask,  $M_P$ ,



if the biolocomotion score  $S_B$  is greater than or equal to some threshold  $\tau_B \in [0, 1]$ . PR curves are generated by computing the precision and recall values at various detection thresholds,  $\tau_B$ , for each annotated frame with precision and recall plotted along the ordinate and abscissa, respectively.

Average precision (**AP**) is a standard evaluation methodology used in action detection (e.g., [48, 54]) that evaluates how well an action proposal (i.e., region likely to contain an action) is ranked for the specified action class. Similarly, AP can be used for biolocomotion detection to determine how well each biolocomotion proposal is ranked. AP is defined as

$$AP = \frac{\sum_{k=1}^n (P(k) \times rel(k))}{\sum_{k=1}^n rel(k)}, \quad (4.3)$$

where  $n$  is the total number of proposals,  $P(k)$  is the precision at cutoff  $k$  of the list of proposals, and  $rel(k)$  is an indicator function which equals 1 if the  $k^{th}$  ranked proposal is a true positive and 0 otherwise. The denominator in (4.3) represents the total number of true positives in the list. To determine whether the proposal should be considered a true or false positive, the Intersection over Union (**IoU**) between the predicted mask and the groundtruth is considered, which is defined as

$$IoU(M_P, M_{GT}) = \frac{M_P \cap M_{GT}}{M_P \cup M_{GT}}, \quad (4.4)$$

where  $M_P$  denotes the predicted mask and  $M_{GT}$  denotes the groundtruth mask. A predicted mask is considered correct if  $IoU(M_P, M_{GT})$  is greater than or equal to

some constant,  $\iota \in [0, 1]$ . AP scores are calculated as  $\iota$  varies. Note that a typical action detection problem considers mean AP (mAP) scores as there are multiple classes to be considered, but the problem of interest considers AP as there is only one class (i.e., biolocomotion).

The AP metric evaluates the *quality* of biolocomotion detection results by ranking the proposals using biolocomotion scores. However, its consideration of positives only disregards missed regions, limiting a thorough investigation of an algorithm. The PR curve, on the other hand, considers both positives and negatives of a detection algorithm. The recall value in a PR curve, however, is often given less importance in detection work [14, 26, 73, 86], since a perfect recall value (i.e.,  $recall = 1$ ) is attained if the predicted mask is larger than the groundtruth (i.e.,  $M_{GT} \subseteq M_P$ ), which can be easily achieved by simply selecting the whole image. This often comes at a cost of a reduced precision value, while a good detection algorithm should be able to (spatiotemporally) locate the biolocomoting object as accurately as possible (i.e., with high precision) [73, 86]. Thus, PR curves become insufficient for comparison, especially if the recall values of the results reside in very little overlapping ranges. Consequently, AP values and PR curves are used in tandem for the evaluation of biolocomotion algorithms, where PR curves serve as a reliable measure for comparing various components of an algorithm and AP values serve as a good metric for comparing the quality of biolocomotion proposals output by different algorithms.

## 4.4 Algorithm Parameter Values

The parameters,  $\tau_\chi$  and  $k_\chi$ , necessary for normalizing the raw measurements,  $\mathcal{M}_\chi$ , into confidence scores,  $S_\chi$  in (3.25), are determined via 1D grid search. The 1D grid search is performed by obtaining PR curves for each component on the BOLD training data for numerous thresholds. The *locomotion*,  $\Lambda$ , component is compared with annotations labelled as (biological or non-biological) locomotion; the *oscillation*,  $\Omega$ , component is compared with annotations labelled as biolocomotion or (biological or non-biological) oscillation; the *asymmetry*,  $\Sigma$ , and *extrinsic motion dissimilarity*,  $E$ , components are compared with annotations labelled as biolocomotion. The values that optimize the area under the PR curve are chosen. Since CAD lacks training data, the same set of parameters as in BOLD are used in CAD. Table 4.3 provides a summary of the empirically set parameters for evaluation.

Measure ( $\chi$ )	$\tau_\chi$	$k_\chi$
Locomotion ( $\Lambda$ )	$51/10 = 5.1$	$(10/51) \ln 99$
Oscillation ( $\Omega$ )	$27/10 = 2.7$	$(10/27) \ln 99$
Asymmetry ( $\Sigma$ )	$3/2 = 1.5$	$(2/3) \ln 99$
Extrinsic Motion Dissimilarity ( $E$ )	$21/200 = 0.105$	$(200/21) \ln 99$

Table 4.3: Parameters used to obtain confidence score,  $S_\chi$ , for measure  $\mathcal{M}_\chi$  for  $\chi \in \{\Lambda, \Omega, \Sigma, E\}$ .

To ensure the predicted masks and groundtruth masks are compared fairly, a tight bounding box fit around the ellipse is used as predicted masks for BOLD since its annotations are bounding boxes.

## 4.5 Alternative Algorithm

Since biolocomotion detection in videos is a new field of research in computer vision, there are no previous algorithms for comparison. Nevertheless, it remains of interest to have some basis for comparison in evaluating the approach to biolocomotion detection presented in Chapter 3. In response, an alternative algorithm has been assembled from components of extant algorithms from the related-field of video-based action analysis. In particular, an extant algorithm representative of state-of-the-art in handcrafted approaches to action proposals [118] and action recognition [121] is considered and adapted to biolocomotion detection. Along with their reliance on generic handcrafted features, these approaches also make use of learning to derive mid-level feature encodings and final classification; see Figure 4.5. Beyond their documented strong performance in action analysis [118, 121], these approaches have the added comparative benefit in that their input are also dense trajectories, similar to the proposed algorithm.

An algorithm based on a ConvNet with deep-learning is not considered as it is not feasible for biolocomotion detection given the relatively small amounts of training data available (cf. action recognition training on UCF101 [106] and kinetics [21], where the datasets are 1-2 orders of magnitude greater). Therefore, the proposed approach is compared to a handcrafted approach with learned mid-level feature encoding and classification.

The action proposal [118] and recognition [121] algorithms are combined as fol-




Figure 4.5: General framework of the alternative algorithm [118]. For each video, a set of proposals [118] are generated and a generic set of descriptors that typically appear with iDTs (e.g., trajectory shape [121], HOG [31], HOF [68], and MBH [32]) are extracted. The descriptors are encoded (e.g., via iFV [87]) into a learned representation specific to biolocomotion detection. The encoded descriptors that result from a set of training videos are used to train a classifier (e.g., SVM [28]), which is used to obtain a biolocomotion score (e.g., via Platt scaling [88]) for encoded descriptors from a test video. The input (raw video) and output (biolocomotion score) are marked in blue and the intermediate steps are marked in green.

lows. A biolocomotion score is associated to each generated *action proposal based on trajectories* (**APT**) [118] by (i) recovering a set of descriptors that typically appear with (improved) dense trajectories (i.e., trajectory shape [120], HOG [31], HOF [68], and MBH [32]), (ii) encoding them using improved Fisher Vectors (iFVs) [87], and (iii) obtaining a biolocomotion score from a classifier (i.e., analogous to the iDT approach to action detection [121]).

Trajectory shape, HOG, HOF, and MBH are standard descriptors used in video analysis tasks, especially action recognition (e.g., [38, 120]). *Trajectory shape* features are obtained by concatenating the normalized displacement vectors of tracked trajectories [120]. *Histogram of Oriented Gradients* (HOG) [31] are descriptors that store spatially oriented gradients of image intensity along a discrete set of directions to capture appearance information. *Histogram of Optical Flow* (HOF) [68] captures local motion of the pattern by quantizing the orientation of the optical flow vec-

tors. *Motion Boundary Histogram* (MBH) [32] captures spatial change of motion and is obtained by computing spatial derivatives along a discrete set of directions for both the horizontal and vertical components of the optical flow field. The employed trajectory shape, HOG, HOF, and MBH parameter values (e.g., discretization sampling) are the same as used in their previous applications to action recognition (e.g., [38, 120]).

The iFVs are computed by (i) reducing the dimension of the features by a factor of two via Principal Component Analysis (PCA), (ii) estimating a Gaussian Mixture Model (GMM) with  $K = 128$  components, representing each proposal by a  $2DK$  FV, where  $D$  is the dimension of the descriptor after PCA has been applied, and (iii) converting to iFVs by applying power and L2 normalization to the FVs, cf. [121]. The feature types are combined by concatenating their iFVs. Finally, given such a test feature vector, a linear SVM [28] (previously built on the training data with loss trade-off parameter set to  $C = 1$ ) assigns a biolocomotion probability score based on Platt scaling [88].

## 4.6 Results

This section provides results and discussions on the proposed and alternative biolocomotion algorithm as evaluated on the BOLD dataset and CAD. First, quantitative results, using AP scores, comparing the proposed to the alternative approach are provided. Second, results that evaluate the performance of each component of the

proposed and alternative approaches using PR curves are provided. Third, qualitative results from the proposed biolocomotion detection algorithm are provided.

#### 4.6.1 Proposed Algorithm vs. Alternative Algorithm

Table 4.4 shows AP scores as Intersection over Union (IoU),  $\iota$ , varies. The proposed algorithm maintains its superior performance over the alternative by a relatively large margin in the majority of the cases (i.e., difference of at least 0.1209 for BOLD at  $\iota \in [0.1, 0.5]$  and CAD at  $\iota \in \{0.1, 0.2\}$ ), a small superiority in two cases (i.e., difference of 0.0028 and 0.0401 for CAD at  $\iota = 0.4$  and  $\iota = 0.5$ , respectively) and performs slightly worse in a single case (i.e., difference of 0.0079 for CAD at  $\iota = 0.3$ ). More typically, the performance margin is in favour of the proposed approach (i.e., in nine of the ten cases) demonstrating the superiority of the proposed approach over the alternative algorithm.

Dataset	Algorithm	$\iota$				
		0.1	0.2	0.3	0.4	0.5
BOLD	Proposed	<b>0.6391</b>	<b>0.5362</b>	<b>0.4404</b>	<b>0.3362</b>	<b>0.2239</b>
	Alternative [118]	0.4566	0.3242	0.2294	0.1535	0.1030
CAD [8]	Proposed	<b>0.8014</b>	<b>0.5166</b>	0.2458	<b>0.1551</b>	<b>0.1275</b>
	Alternative [118]	0.4183	0.3262	<b>0.2537</b>	0.1523	0.0874

Table 4.4: Average Precision (AP) scores as a function of Intersection over Union (IoU),  $\iota$ , for the BOLD dataset and CAD.

## 4.6.2 Component Analysis

### Proposed Algorithm


The effect of each component in the biocomotion algorithm (i.e.,  $\Lambda$ ,  $\Omega$ ,  $\Sigma$ , and  $E$ ) can be assessed through an ablation study; see Figure 4.6. Here, the ablation study is designed to examine the effect of adding each component of the presented algorithm in sequence ( $\Lambda$ ,  $\Lambda + \Omega$ ,  $\Lambda + \Omega + \Sigma$ , and  $\Lambda + \Omega + \Sigma + E$ ). A complete study that contains all  $2^4 - 1$  combinations can be found in Appendix C. The ablation study is conducted only on the BOLD dataset, as it contains a better balanced set of biological and non-biological object videos as well as objects in locomotion and non-locomotion (as opposed to CAD).

It can be observed in Figure 4.6a that the recall values of the proposed method tend to be skewed to the left (i.e., ranging between  $[0, 0.4]$ ). This fact is largely due to the condition that labels a cluster that occupies the majority of the field of view as a cluster corresponding to camera motion; see Section 3.2.2 and Appendix A.3. As a result, the number of missed pixels would increase, especially if the object occupies the majority of the frame, which in turn reduces the recall rate.

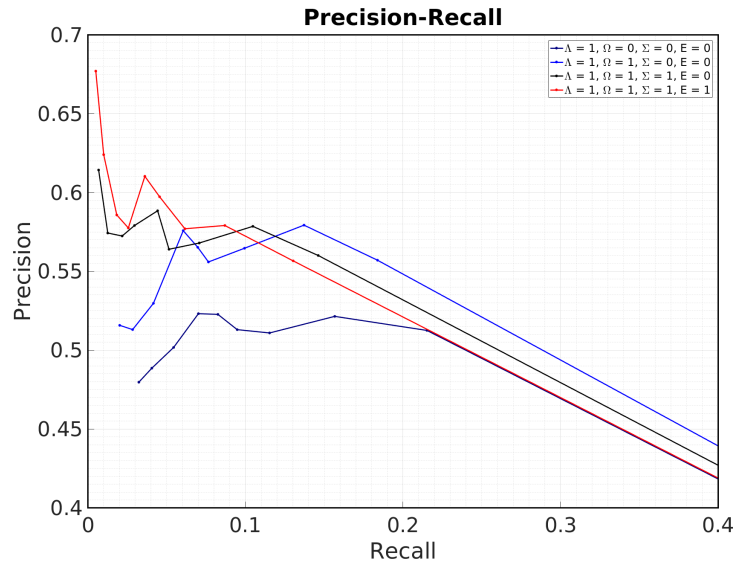
It can be seen in Figure 4.6b that adding each component increases the precision value with a slight decrease in recall for a specified  $\tau_B$ . Specifically, adding oscillation to locomotion (cf. navy and blue) improves the average of the precision differences by  $+0.0383$ ; adding asymmetry to locomotion and oscillation (cf. blue and black) by  $+0.0243$ ; and adding extrinsic motion dissimilarity to locomotion, oscillation, and



asymmetry (cf. black and red) by +0.0186. The benefit of each component can be observed through the vertical shift upwards and is confirmed via positive average precision differences. The relatively low rate of improvement by adding the extrinsic motion dissimilarity component can be attributed to the low thresholds that were used (i.e.,  $\tau_E = 0.105$  and  $k_E = 43.76$ ). This choice ensures that a relatively small number of FN is introduced. Nevertheless, this study confirms the benefits of each component in the proposed biomechanics and psychophysics driven algorithm.



(a) PR curve in full scale.



(b) PR curve for recall  $\in [0, 0.4]$  and precision  $\in [0.4, 0.7]$ .

Figure 4.6: Precision-Recall (PR) curves of the proposed algorithm. ‘1’ and ‘0’ indicate the enabled and disabled components, respectively. (a) a full PR curve is provided with both precision and recall in the ranges  $[0, 1]$  and (b) the same PR curve, as (a), restricted to recall values within  $[0, 0.4]$  and precision within  $[0.4, 0.7]$  as the majority of the points are focused in that select region of the PR curve.

## Alternative Algorithm

Figure 4.7 provides an evaluation of the alternative algorithm to gauge the role of appearance, motion, and its combination by considering iFV-encodings of HOG (green), trajectory shape + HOF + MBH (blue), and trajectory shape + HOG + HOF + MBH (red), respectively. It can be observed that considering appearance information alone (i.e., HOG) yields the most inferior result in detecting biolocomotion across both datasets. This result is potentially due to two possibilities: (1) There are so many in-class variations in appearance (e.g., humans, birds, fish, cars, and planes) that the considered training data does not suffice to capture all variations. (2) The detected biological object is not guaranteed to be locomoting. Moreover, augmenting motion information with appearance information adds no apparent benefit. These results confirm that eschewing appearance information and obtaining results strictly based on motion information is the desirable path to take when detecting biolocomotion in videos.





Figure 4.7: Precision-Recall (PR) curves for alternative algorithm.

### 4.6.3 Qualitative Results

Figure 4.8 provides qualitative detection results on a set of seven example frames from the BOLD dataset for the proposed algorithm. Various examples of locomoting biological objects are shown in Figure 4.8a. While the species type varies (e.g., dog, bird, fish, and human), the way in which these species make a positional advance in the world (i.e., by running, flying, swimming, or climbing) all include oscillatory behaviour with extrinsic motion dissimilarity and an added asymmetry in its oscillation if moving orthogonal to the resistive force (i.e., running dog, flying bird, and swimming fish). As a result, they project oscillatory trajectories as they locomote inducing signs of biolocomotion. On the contrary, a non-biological object (e.g., ball) locomoting is shown in Figure 4.8b. While this object is making a spatial advance (i.e., by rolling) and even includes oscillation (in its cycloidal trajectory trace), since it does not exhibit *asymmetric* oscillatory behaviour, it does not yield large  $S_B$  values to indicate biolocomotion. Conversely, an example of a biological object (bird) not locomoting is also shown in Figure 4.8c. Even though the bird exhibits some oscillatory behaviour (i.e., jumping up and down), it does not locomote and hence is correctly not detected as biolocomotion. While the measure of locomotion,  $\mathcal{M}_\Lambda$ , is able to capture locomotion of (biological or non-biological) objects in the majority of the cases, it fails to capture objects primarily moving along the optical axis (i.e., towards (or away) from the camera); as depicted in Figure 4.8d. As a result, a low signal for locomotion is obtained, resulting in a false negative biolocomotion detection. Future research can address this limitation by considering measurements that are indicative of change in the depth of an object within a temporal window.



(a) Biological objects in locomotion



(b) Non-biological object in locomotion



(c) Biological object not in locomotion



(d) Biological object in locomotion not detected as locomoting

Figure 4.8: Example results from the proposed biocomotion detection algorithm. (a) Illustration of true detection results on different types of biological objects (dog, bird, fish, and human) locomoting in various ways (run, fly, swim, and climb). (b) Illustration of correctly undetected non-biological object (ball) in locomotion. (c) Illustration of correctly undetected biological object (bird) not in locomotion. (d) Illustration of a falsely undetected biological object (cat) in locomotion. Detected regions are indicated by red bounding boxes that cover the trajectory clusters that led to the detection.

## 4.7 Summary

This chapter has provided qualitative and quantitative results to validate the proposed algorithm's ability to detect biolocomotion and its superiority over an alternative approach that combines handcrafted features with learning. The results confirmed that biolocomotion can be modelled by measuring spatial advance of an object (locomotion) in conjunction with asymmetric oscillatory patterns and extrinsic motion dissimilarity. Furthermore, the results indicate that eschewing appearance information and relying solely on motion information in detecting biolocomotion yields superior performance. These studies were enabled by the introduction of the first dataset built for biolocomotion detection, BOLD.

# Chapter 5

## Conclusion

### 5.1 Summary

This dissertation has shed light on a new research topic in the field of computer vision: biolocomotion detection in videos. While previously not considered in computer vision, biolocomotion detection is an important topic not only because it presents an interesting and a well-defined challenge, but also for its wide spread applicability as a precursor for subsequent analysis (e.g., tracking and action recognition). Moreover, research here has the potential to provide a computational model for the apparent ability of biological systems to detect biolocomotion from visual data.

Inspired by biomechanical properties of animals in locomotion and the perception of biological motion in psychophysics, a unified computational algorithm to detect biolocomotion in videos has been developed. In particular, the developed approach



exploits the distinctive signature of asymmetric oscillatory pattern as biological entities make a positional advance, which also tends to yield dissimilar extrinsic and intrinsic motions. Since the developed algorithm is based on distinctive *motion* patterns of locomoting biological objects, it enjoys the benefit of not needing to model the wide range of variations that exist between and within biological and non-biological objects.

A novel dataset, BOLD, has been assembled and biolocomotion annotations on BOLD as well as on an extant camouflage animals dataset have been supplied to provide evaluation benchmarks for the developed and future biolocomotion detection algorithms. To demonstrate the efficacy of the developed approach, an alternative approach based on generic handcrafted features with learning also has been developed and compared. Quantitative results indicate that the proposed algorithm considerably outperforms the alternative approach. These results support the hypothesis that biolocomotion can be detected reliably based on the biomechanically and psychophysically motivated signature of positional advance with asymmetric oscillation and extrinsic motion dissimilarity.

## 5.2 Future Work

In light of the work presented in this dissertation, several directions for future work can be considered.

While the proposed approach has demonstrated its ability to detect a wide range of locomoting biological species in videos, it has also revealed its limitations in detecting objects that move parallel to the optical axis (i.e., towards (or away) from the camera, as in Figure 4.8d). An algorithm that can estimate the change in depth of such locomoting species could provide a solution to such scenarios. A potential approach could be based on measuring motion along the optical axis based on the relative rate of the visually apparent expansion (or contraction) of an object across time, as evidenced in biological systems (e.g., [96]) and employed in collision avoidance systems for autonomous vehicles (e.g., [63]). Combining the result with the currently proposed measurement of locomotion could better encompass the full three-dimensional displacement of biological species in motion.

One of the motivations for the computational study of biolocomotion detection is that it can serve as a precursor for many related tasks. Once a living creature has been detected, it can be recognized (e.g., human, cat, bird), its action can be classified (e.g., running, flying, jumping), or even more subtle distinctions can be made (e.g., gender, age, emotional state, personality trait) [113]. Correspondingly, an interesting direction for future research is to couple the biolocomotion detector with such subsequent processes.

Another source of motivation for the computational study of biolocomotion detection is that it might yield a computational theory of how biological systems solve this problem. Toward that end, one shortcoming of the current algorithm is that it is

not subject to the inversion effect, whereby inverted biolocomotion displays are not easily recognized by animals [115, 117]. An interesting direction for future research is to consider how this constraint could be incorporated into the detector and what advantages it might provide. Further, the specific mechanisms that are incorporated in the current detector could provide the basis for psychophysical experiments to further probe how biological systems operate.

Finally, various combinations of the components introduced in this dissertation might be combined in different ways to detect more general types of motion, not restricted to locomotion. For example, studies have shown that animacy can be perceived from the motion of simple rigid geometric shapes, such as triangles, circles, and squares [47]; see Figure 5.1. Therefore, future research can consider how the biolocomotion detection components defined in this dissertation can be leveraged to detect more general animacy.




Figure 5.1: Illustration of perception of animacy. Select frames of rigid geometric shapes in motion displaying perception of animacy. The way geometric figures move provides sufficient information to give observers the perception of animacy. For example, an observer can interpret the display as the red triangle chasing after the blue triangle and the green circle moving around an obstacle (black opened box). Figure redrawn from [101].

# Appendix A

## Trajectory Processing

Motivated by its state-of-the-art performance amongst handcrafted algorithms for action recognition, improved Dense Trajectories (iDTs) [121] are used as input to the proposed biolocomotion detector. For the best performance, trajectories had to be further processed prior to passing them to the biolocomotion detector, as follows. First, trajectories must be pruned to remove unuseful trajectories. Second, trajectories need to be clustered such that biolocomotion detection operates on sets of trajectories that are likely to correspond to a single or similarly moving entity in the world. Third, while iDTs are claimed to be robust to camera motion, trajectories need to be stabilized to (i) remove object-centric stabilization induced by the camera operator or (ii) remove residual motion arising from camera motion that contain corrupted data for measurement. Finally, select trajectories in a video can be elongated, without increasing its chance for drifting, to allow for more temporal support in biolocomotion detection. In this appendix, the details of each trajectory

processing step is described.

## A.1 Pruning Trajectories

To ensure that a meaningful set of trajectories are obtained, static trajectories that do not contain motion information and random trajectories are removed in the post-processing stage. While the need and conditions to remove static and random trajectories are addressed in the original iDT work [121], their solution was insufficient for present purposes. In particular, the removal of static trajectories via standard deviation and the removal of random trajectories via arc length was insufficient. In addition, the removal of trajectories if the trajectory exceeded some pre-defined threshold was too aggressive in that they removed trajectories that corresponded to objects moving quickly (e.g., cars). Thus, more variations and less aggressive measures to prune static and random trajectories that would yield meaningful information in identifying biolocomotion are considered in this section.

### A.1.1 Removing Static Trajectories

Static trajectories can be removed using (i) standard deviation, (ii) arc length, and (iii) removing trajectories with  $m$  relatively static points. Given a trajectory  $\mathcal{T}^k$ , its mean and standard deviation are defined as

$$(\mu_x^k, \mu_y^k) = \left( \frac{1}{L^k} \sum_{l=1}^{L^k} x_l^k, \frac{1}{L^k} \sum_{l=1}^{L^k} y_l^k \right) \quad (\text{A.1})$$

and

$$(\sigma_x^k, \sigma_y^k) = \left( \sqrt{\frac{1}{L^k} \sum_{l=1}^{L^k} |x_l^k - \mu_x^k|^2}, \sqrt{\frac{1}{L^k} \sum_{l=1}^{L^k} |y_l^k - \mu_y^k|^2} \right), \quad (\text{A.2})$$

respectively. A trajectory is considered static if the standard deviation in the horizontal and vertical directions are below some threshold,  $\tau_x$  and  $\tau_y$  (i.e., if  $(\sigma_x^k < \tau_x) \wedge (\sigma_y^k < \tau_y)$ , then  $\mathcal{T}^k$  is static). The thresholds can be set to some pre-defined values like  $\tau_x = \tau_y = \sqrt{3}$ , which were empirically found to yield acceptable results.

Trajectory  $\mathcal{T}^k$  is considered static if its arc length,  $\|\mathcal{T}^k\|$  as in (3.5), is shorter than some threshold (i.e., if  $\|\mathcal{T}^k\| < \tau$ , then  $\mathcal{T}^k$  is static). Again,  $\tau$  can be set with some pre-defined value, like  $\tau = 10$ , which has shown to yield good results.

Lastly, a trajectory is considered static if it contains  $m$  relatively static points. A point is considered static if its position has not changed in the next frame (i.e.,  $(|x_{l+1}^k - x_l^k| \approx 0) \wedge (|y_{l+1}^k - y_l^k| \approx 0)$ ). In practice, a small threshold,  $\tau_s$ , is assigned to determine if a point is *relatively* static (i.e.,  $(|x_{l+1}^k - x_l^k| < \tau_s) \wedge (|y_{l+1}^k - y_l^k| < \tau_s)$ ). The number of relatively static points that a trajectory contains can be determined by summing the number of points that satisfy the relatively static condition. That is, if

$$\sum_{l=1}^{L^k-1} [(|x_{l+1}^k - x_l^k| < \tau) \wedge (|y_{l+1}^k - y_l^k| < \tau)] > m, \quad (\text{A.3})$$

then  $\mathcal{T}^k$  is considered static. Empirically,  $\tau = 1$  and  $m = \frac{2}{3}L^k$  has shown to prune static trajectories well.

### A.1.2 Removing Random Trajectories

Trajectories with sudden large displacements are likely to be erroneous. Thus, it would be helpful to remove them before further processing the data. Random trajectories can be removed using (i) standard deviation, (ii) arc length, and (iii) the ‘68-95-99.7’ rule [119].

Using the standard deviation as defined in equation (A.2), a trajectory is likely to be erroneous if  $(\sigma_x > \tau_x) \vee (\sigma_y > \tau_y)$ , since it implies that  $\mathcal{T}^k$  has a sudden large displacement in either the horizontal or vertical direction. Empirically,  $\tau_x = \tau_y = 75$  has shown to prune random trajectories, while retaining non-random ones.

Random trajectories are pruned using arc length as done in the original evaluation [121]. More specifically, if the displacement vector between two consecutive frames is larger than 70% of the overall displacement, then it indicates a sudden large displacement, which is likely to be a random trajectory. Formally, if  $\sqrt{(x_{l+1}^k - x_l^k)^2 + (y_{l+1}^k - y_l^k)^2} > 0.7\|\mathcal{T}^k\| \forall k$ , then  $\mathcal{T}^k$  is likely to be a random trajectory.

Finally, the standard deviation of the trajectories is used more adaptively via the ‘68-95-99.7’ rule. That is, if  $(\mu_x - \tau_\sigma\sigma_x < x_l^k < \mu_x + \tau_\sigma\sigma_x) \wedge (\mu_y - \tau_\sigma\sigma_y < y_l^k < \mu_y + \tau_\sigma\sigma_y)$  for  $1 \leq l \leq L^k$ , it means that all the points of the trajectory lie within 68%, 95%, or 99.7% of the data if  $\tau_\sigma = 1, 2$ , or  $3$ , respectively. Otherwise, a point is considered “random” within the trajectory. Empirically,  $\tau_\sigma = 3$  retained useful



trajectories while removing random ones.

## A.2 Clustering Trajectories

To ensure that the developed algorithm is applicable in the real world, it should be robust to various factors that could occur in videos from the wild, such as objects occurring simultaneously in a single frame (see Figure A.1), presence of camera motion, etc. One of the ways in which these factors can be dealt with is to cluster them and subsequently handle them separately. Thus, in this section, a method to cluster trajectories is introduced followed by a way to identify and remove camera motion in videos.



Figure A.1: Comparison of single object per frame vs. multiple objects per frame.

Given a set of trajectories  $\mathcal{T} = \{\mathcal{T}^1, \dots, \mathcal{T}^N\}$ , the goal is to cluster them into disjoint sets that correspond to similarly moving objects in the world. A variant on spectral trajectory clustering is to be employed since the original formulation [17] produced poor results with iDTs for the biolocomotion detection task. This shortcoming is likely due to the fact that the original intent of the algorithm was to cluster

trajectories that were mostly linear, while the current goal is to cluster oscillatory trajectories. Correspondingly, novel measures of positional and shape affinity,  $d_p$  and  $d_s$ , respectively, for trajectory pairs are defined to perform spectral clustering. The output of this processing stage are the centres as well as the horizontal and vertical extents of ellipses that cover each cluster.

Let  $\mathcal{T}^i$  and  $\mathcal{T}^j$  represent two trajectories that span the same temporal length  $L$  (i.e.,  $L = L^i = L^j$ ) in  $\mathcal{T}$ . Then the *proximity* of two trajectories can be measured by considering a weighted Euclidean distance between the average spatiotemporal coordinates of the trajectories,

$$d_p(\mathcal{T}^i, \mathcal{T}^j) = \alpha \left[ \frac{(\bar{x}^i - \bar{x}^j)^2}{w} + \frac{(\bar{y}^i - \bar{y}^j)^2}{h} \right] + (1 - \alpha) \left[ \frac{(\bar{t}^i - \bar{t}^j)^2}{f} \right], \quad (\text{A.4})$$

where  $\bar{x}^k$ ,  $\bar{y}^k$ , and  $\bar{t}^k$  denote spatial and temporal averages, respectively, of trajectory  $\mathcal{T}^k$  for  $k \in \{i, j\}$ . That is,  $\bar{\mathbf{x}}^k = (\bar{x}^k, \bar{y}^k, \bar{t}^k) = (\frac{1}{L} \sum_{l=1}^L x_l^k, \frac{1}{L} \sum_{l=1}^L y_l^k, \frac{1}{L} \sum_{l=1}^L t_l^k)$  with  $\mathbf{x}_l^k = (x_l^k, y_l^k, t_l^k)$  the spacetime coordinates of the  $l^{\text{th}}$  point in trajectory  $k$  at frame  $t_l$ . To account for variations in resolution and frame rate of videos, the spatial distances are divided by the width,  $w$ , or height,  $h$ , and the temporal distance is divided by the frame rate,  $f$ . Greater emphasis on spatial distance over temporal distance (or vice versa) can be applied by the use of a constant term,  $\alpha \in \mathbb{R}$ , where  $0 \leq \alpha \leq 1$ ;  $\alpha = 0.5$  is used in practice. The formulation in (A.4) ensures for a pair of spatiotemporally close trajectories  $d_p(\mathcal{T}^i, \mathcal{T}^j) \rightarrow 0$ , while those that are far tend to  $\infty$ .

To compare the *shape* of a pair of trajectories,  $\mathcal{T}^i$  and  $\mathcal{T}^j$ , subsections of trajec-

tories (or trajectory displacements) can be compared by considering either (i) the squared  $L_2$ -norm or (ii) the angle of the normalized inner product. While both are valid methods of comparison, considering the angle ensures that the overall *shape* of the trajectory is compared while being minimally affected by the length of the trajectory displacements. Abstracting away the magnitude of the displacements between frames would ensure robustness to a potentially varying amount of displacement exhibited by different limbs of a biological object. That is, for displacement vector  $\Delta \mathbf{x}_l^k$  of trajectory  $\mathcal{T}^k$  at the  $l^{\text{th}}$  point for  $1 \leq l \leq L^k$  and  $k \in \{i, j\}$ , its angle of the normalized inner product of two vectors is defined as

$$\Theta(\Delta \mathbf{x}_l^i, \Delta \mathbf{x}_l^j) = \arccos \left( \frac{\langle \Delta \mathbf{x}_l^i, \Delta \mathbf{x}_l^j \rangle}{\|\Delta \mathbf{x}_l^i\| \|\Delta \mathbf{x}_l^j\|} \right). \quad (\text{A.5})$$

The angle of the normalized inner product to consider is bound between  $[0, 180]$ , where similar vectors tend to 0 while very different vectors tend to 180, instead of a bound between  $[-1, 1]$ , where two alike vectors tend to 1 and two dissimilar vectors tend to  $-1$ .

While the individual angular based difference between vectors could be combined across trajectories by simply taking their mean or maximum, taking a weighted sum according to the percentage of the overall trajectory length of the displacements is advantageous as longer portions of the trajectory will count more heavily. Corre-

spondingly, a weight can be defined as

$$w(\Delta \mathbf{x}_l^i, \Delta \mathbf{x}_l^j) = \frac{\|\Delta \mathbf{x}_l^i\|}{\sum_l \|\Delta \mathbf{x}_l^i\|} \frac{\|\Delta \mathbf{x}_l^j\|}{\sum_l \|\Delta \mathbf{x}_l^j\|}. \quad (\text{A.6})$$

Finally, while the overall shape difference between a pair of trajectories can be compared by taking the weighted sum of the individual similarities between vectors as

$$d_s(\mathcal{T}^i, \mathcal{T}^j) = \sum_{n=1}^{L-1} w(\Delta \mathbf{x}_l^i, \Delta \mathbf{x}_l^j) \Theta(\Delta \mathbf{x}_l^i, \Delta \mathbf{x}_l^j), \quad (\text{A.7})$$

the phase shifts that could occur between opposing limbs (e.g., left and right arms of a human walker) must be accounted for. To obtain a phase invariant shape measure, the minimum of all phase shifts of  $\mathcal{T}^j$  with respect to  $\mathcal{T}^i$  are considered

$$d_{s'}(\mathcal{T}^i, \mathcal{T}^j) = \min_{l=0}^{L-2} \left\{ \sum_{l=1}^{L-1} w(\Delta \mathbf{x}_l^i, \Delta \mathbf{x}_{l-l}^j) \Theta(\Delta \mathbf{x}_l^i, \Delta \mathbf{x}_{l-l}^j) \right\}, \quad (\text{A.8})$$

which is used as the shape similarity between two trajectories.

The two traits that distinguish trajectories from another,  $d_p$  and  $d_{s'}$ , can be transformed into affinities for spectral clustering via

$$W(\mathcal{T}^i, \mathcal{T}^j) = \exp(-\kappa_s \cdot d_p(\mathcal{T}^i, \mathcal{T}^j) - d_{s'}(\mathcal{T}^i, \mathcal{T}^j)), \quad (\text{A.9})$$

where  $\kappa_s \in \mathbb{R}$  is some constant, which is empirically set to  $\kappa_s = 0.01$ .

Once the affinity matrix has been defined, spectral clustering can be employed.

To compensate for potential noise in the defined affinity matrix, eigendecomposition is performed on a normalized Laplacian

$$V^\top \Lambda V = D^{-\frac{1}{2}}(D - W)D^{-\frac{1}{2}}, \quad (\text{A.10})$$

where  $D$  is the degree matrix of  $W$ , and  $V$  and  $\Lambda$  correspond to matrices containing the eigenvector and eigenvalues, respectively [17]. Then  $k$ -means clustering is performed on the eigenvectors that correspond to  $k$  smallest nonzero eigenvalues.

To allow spectral clustering to be as automated as possible, that is to limit the process of manually selecting the number of clusters for the  $k$ -means clustering step, a cluster validity index is employed. A cluster validity index is a measure that compares the compactness of a cluster (or the intra-class relationship between features *within* a cluster) and the separatedness of clusters (or the inter-class relationship of features *between* clusters). There are various cluster validity indices (e.g., Dunn [37] and Davies-Bouldin [33]) that can be used to identify the ideal number of clusters [34]. Here, *silhouette* value [99] is employed.

For simplicity, suppose  $\mathcal{T}^i$  is a trajectory that belongs to cluster  $\mathcal{C}^i$  and none other (i.e.,  $\mathcal{T}^i \in \mathcal{C}^i$  and  $\mathcal{T}^i \notin \mathcal{C}^{k_i}$ , where  $\mathcal{C}^{k_i}$  denotes any cluster that is not  $\mathcal{C}^i$  and  $\cap_i \mathcal{C}^{k_i} = \emptyset$ ); see Figure A.2. Then the *within* similarity of  $\mathcal{T}^i$  to all the other trajectories in  $\mathcal{C}^i$  can be calculated by taking the average similarity of  $\mathcal{T}^i$  to all other

trajectories in  $\mathcal{C}^i$  according to

$$a(\mathcal{T}^i) = \frac{1}{|\mathcal{C}^i|} \sum_{\mathcal{T}^j} W(\mathcal{T}^i, \mathcal{T}^j) \quad \forall \mathcal{T}^j \in \mathcal{C}^i, \quad (\text{A.11})$$

where  $W(\mathcal{T}^i, \mathcal{T}^j)$  is the similarity score between  $\mathcal{T}^i$  and  $\mathcal{T}^j$  as in (A.9). The *between* similarity of  $\mathcal{T}^i$  to any other cluster  $\mathcal{C}^{k_i}$  is calculated in a similar manner,

$$d(\mathcal{T}^i, \mathcal{C}^{k_i}) = \frac{1}{|\mathcal{C}^{k_i}|} \sum_{\mathcal{T}^j} W(\mathcal{T}^i, \mathcal{T}^j) \quad \forall \mathcal{T}^j \in \mathcal{C}^{k_i}. \quad (\text{A.12})$$

Once the *between* similarity of  $\mathcal{T}^i$  to all other clusters,  $\mathcal{C}^{k_i} \forall i$ , have been computed, the second best cluster for  $\mathcal{T}^i$  is chosen by taking the maximum of the between clusters according to

$$b(\mathcal{T}^i) = \max_{\mathcal{C}^{k_i} \forall i} \{d(\mathcal{T}^i, \mathcal{C}^{k_i})\}. \quad (\text{A.13})$$

The *silhouette* of  $\mathcal{T}^i$  is obtained by combining  $a(\mathcal{T}^i)$  and  $b(\mathcal{T}^i)$  according to

$$s(\mathcal{T}^i) = \begin{cases} 1 - \frac{a(\mathcal{T}^i)}{b(\mathcal{T}^i)}, & \text{if } a(\mathcal{T}^i) < b(\mathcal{T}^i) \\ 0, & \text{if } a(\mathcal{T}^i) = b(\mathcal{T}^i), \\ \frac{b(\mathcal{T}^i)}{a(\mathcal{T}^i)} - 1, & \text{if } a(\mathcal{T}^i) > b(\mathcal{T}^i) \end{cases} \quad (\text{A.14})$$




Figure A.2: Silhouettes [99]. The quality of the clusters is quantified using silhouette values. The *within* similarity of  $\mathcal{T}^i \in \mathcal{C}^i$ , denoted  $a(\mathcal{T}^i)$ , is taken as the average similarity of  $\mathcal{T}^i$  to all other trajectories in  $\mathcal{C}^i$  (*orange lines*). The *between* similarity of  $\mathcal{T}^i$  to other clusters  $\mathcal{C}^{k_i}$  (*blue lines*), denoted  $d(\mathcal{T}^i, \mathcal{C}^{k_i})$ , is calculated by taking the average similarity of  $\mathcal{T}^i$  to all objects in cluster  $\mathcal{C}^{k_i}$ . The second most suitable cluster for  $\mathcal{T}^i$  is chosen by selecting the cluster with the largest between similarity (e.g.,  $b(\mathcal{T}^i) = \max\{d(\mathcal{T}^i, \mathcal{C}^{k_1}), d(\mathcal{T}^i, \mathcal{C}^{k_2})\}$ ). The *within*,  $a(\mathcal{T}^i)$ , and *between*,  $b(\mathcal{T}^i)$ , values are summarized into a single value using a silhouette value  $s(\mathcal{T}^i)$ .

which can be summarized into a single formula as

$$s(\mathcal{T}^i) = \frac{a(\mathcal{T}^i) - b(\mathcal{T}^i)}{\max\{a(\mathcal{T}^i), b(\mathcal{T}^i)\}}. \quad (\text{A.15})$$

Since it is unclear how  $a(\mathcal{T}^i)$  should be defined if there is only one trajectory in cluster  $\mathcal{C}^i$ , by convention,  $s(\mathcal{T}^i)$  is set to 0.

It is worth noting that  $s(\mathcal{T}^i)$  is defined in a way such that  $-1 \leq s(\mathcal{T}^i) \leq 1$ .  $s(\mathcal{T}^i) \rightarrow 1$  if  $a(\mathcal{T}^i)$  is large and  $b(\mathcal{T}^i)$  is small, which implies that the *within* similarity,  $a(\mathcal{T}^i)$ , is greater than the *between* similarity,  $b(\mathcal{T}^i)$ , implying that  $\mathcal{T}^i$  has been well-clustered.  $s(\mathcal{T}^i) \rightarrow 0$  if  $a(\mathcal{T}^i) \approx b(\mathcal{T}^i)$ , which implies it is unclear whether  $\mathcal{T}^i$  should have been clustered into  $\mathcal{C}^i$  or  $\mathcal{C}^k$ .  $s(\mathcal{T}^i) \rightarrow -1$  if  $a(\mathcal{T}^i)$  is small and  $b(\mathcal{T}^i)$  is large, which implies that the *within* similarity,  $a(\mathcal{T}^i)$ , is less than the *between* similarity,  $b(\mathcal{T}^i)$ , implying that it would have been better to classify  $\mathcal{T}^i$  into cluster  $\mathcal{C}^k$  rather than  $\mathcal{C}^i$ . Hence,  $\mathcal{T}^i$  has been misclassified.

The quality of the final clustering result is evaluated by computing the *overall average silhouette*,  $S^k$ , for  $k$  clusters as follows

$$S^k = \frac{1}{N} \sum_{i=1}^N s(\mathcal{T}^i) \quad \forall \mathcal{T}^i, \quad (\text{A.16})$$

where  $N$  is the number of trajectories in cluster  $\mathcal{C}^k$ . The cluster that yields the largest overall average silhouette,  $S^k$ , is indicative of the most optimal number of partitions. Silhouettes offer a quantitative analysis of the final outcome of a cluster-



ing algorithm without the knowledge of which clustering algorithm was actually used to acquire such result [99]. Since the most computationally expensive step of spectral clustering lies in the construction of the affinity matrix,  $W$ , numerous clusters can be evaluated, and the optimal number of clusters can be identified, rather efficiently, by selecting the cluster with the largest average silhouette width,  $S$ .

Finally, to recover the extent of each cluster, GMMs are used to model the shape of the clusters at frame  $t$ . Specifically, the points that are present at frame  $t$  are represented by GMMs with varying number of components (e.g., two to five components) to yield a candidate of shapes for the cluster. Among the candidate of GMMs, the GMM with the smallest Akaike Information Criterion (AIC) value [2] is selected to represent the cluster of points at frame  $t$ . To ensure temporal smoothness of the estimated GMM-based masks, each frame is weighted according to its neighbouring frames. That is, suppose  $B_t^k$  represents a binary mask of cluster  $k$  at frame  $t$ . Then its smoothed map is a weighted sum of its neighbouring frames

$$\tilde{B}_t^k = \sum_{t'=1}^T w(t, t') B_{t'}^k \quad (\text{A.17})$$

for  $1 \leq t' \leq T$ , where  $w(t, t') = \exp(-\lambda(t - t')^2)$  for constant  $\lambda$ , which controls the amount of contribution the neighbouring frames have on frame  $t$ . Empirically,  $\lambda$  is set to 0.1. The weighted sum of the binary masks results in a real-valued map.

Therefore, the final binary mask,  $\mathcal{B}_t^k$ , is defined as

$$\mathcal{B}_t^k(\mathbf{x}) = \begin{cases} 1, & \text{if } \tilde{B}_t^k(\mathbf{x}) > \tau_d \\ 0, & \text{otherwise} \end{cases}, \quad (\text{A.18})$$

at some spatial coordinate  $\mathbf{x}$ , where  $\tau_d$  is empirically set to 0.7.

### A.3 Stabilizing Trajectories

While iDTs are designed to have a certain level of robustness to camera motion, the wide range of pans and zooms present in videos in-the-wild makes the entailed processing inadequate. Other recent approaches to image-based camera motion estimation and cancellation [85] have also been inadequate. Thus, an approach to camera stabilization that corrects the overall global shift between frames (approximation to pan and tilt), linear expansion/contraction about a centre (approximation to zoom) as well as estimation of frames without camera motion is needed. In this section, a method to camera stabilization using trajectories is described.

#### Estimating Global Shift

Trajectories that arise as a result of overall global translation (e.g., horizontal, vertical, and/or its combination) between frames should have similar shapes and therefore be grouped into a single cluster according to processing described in §A.2. Under the assumption that the field of view is not dominated by the object of inter-

est (e.g., human, car), trajectories corresponding to camera motion would be well-dispersed throughout the field of view. Thus, the cluster whose trajectories occupy a widespread area of the frame can be taken as the cluster indicative of camera motion manifest as global shift. That is, cluster  $\mathcal{C}^k$  containing a set of trajectories correspond to global shift if

$$\left(\frac{w}{2} - v_x(t) < \tau_{cx}\right) \vee \left(\frac{h}{2} - v_y(t) < \tau_{cy}\right), \quad (\text{A.19})$$

where  $(v_x(t), v_y(t))$  correspond to horizontal and vertical extents of cluster  $\mathcal{C}^k$  at frame  $t$ , while  $w$  and  $h$  correspond to the width and height of the image, respectively. The thresholds,  $(\tau_{cx}, \tau_{cy})$ , to determine if the trajectories occupy the majority of the frame is defined as a ratio with respect to the resolution of the image (e.g.,  $(\tau_{cx}, \tau_{cy}) = (\frac{w}{15}, \frac{h}{15})$ ).

Once the cluster of trajectories corresponding to camera motion has been identified, global shift can be estimated by calculating the median displacement of the trajectories on a frame-by-frame basis. That is, for displacement vector  $\Delta \mathbf{x}_{(t, \Delta t)}^k$  of trajectory  $\mathcal{T}^k$  between frames  $t$  and  $t - \Delta t$ , global shift at frame  $t$  is defined as

$$(x_c(t), y_c(t)) = (\text{med}_{x^k(t)}\{\Delta x^k(t)\}, \text{med}_{y^k(t)}\{\Delta y^k(t)\}). \quad (\text{A.20})$$

Finally, with the global shift estimated, motion caused by camera motion can be removed by subtracting the estimated shift from each trajectory.

## Estimating Global Expansion/Contraction

The displacement vectors in frames with global expansion (or contraction) motion (i.e., arising from camera zoom in (or zoom out)) exhibit unique patterns; see Figure A.3. To handle frames with global expansion (or contraction), pattern of expansion (or contraction) must be detected, the scale of expansion (or contraction) imposed by the camera operator must be removed, then shifted to match the focus of expansion (or contraction). In the following, each of the steps necessary to remove global expansion (or contraction) will be detailed by describing the case where the focus of expansion (or contraction) is at the centre of image first, followed by an arbitrary focus of expansion (or contraction) location.

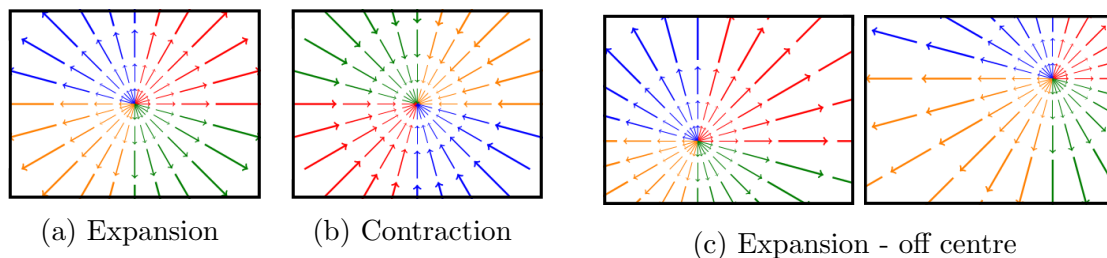


Figure A.3: Illustration of trajectories in presence of linear expansion or contraction. The displacement vectors of each trajectory are grouped based on the quadrant that it belongs to (red for displacement vectors in quadrant  $Q_I$ , blue for  $Q_{II}$ , orange for  $Q_{III}$  and green for  $Q_{IV}$ ). A frame that contains expansion (i.e., zoom in) exhibits the following characteristic:  $Q_{II}$  is to the left of  $Q_I$ ,  $Q_{III}$  is to the left of  $Q_{IV}$ ,  $Q_{II}$  is above  $Q_{III}$  and  $Q_I$  is above  $Q_{IV}$ . Conversely, a frame that contains contraction (i.e., zoom out) exhibits the following pattern:  $Q_{IV}$  is to the left of  $Q_{III}$ ,  $Q_I$  is to the left of  $Q_{II}$ ,  $Q_{IV}$  is above  $Q_I$  and  $Q_{III}$  is above  $Q_{II}$ .

To identify the pattern indicative of expansion (or contraction), all displacement

vectors present at frame  $t$  must be assigned to a quadrant based on its orientation then subsequently consider the spatial arrangements of the assigned vectors. A displacement vector,  $\Delta \mathbf{x}_{(t,\Delta t)}^k$ , for trajectory  $\mathcal{T}^k$  present at frame  $t$  with a temporal difference of  $\Delta t$  is assigned to a quadrant according to

$$\Delta \mathbf{x}^k(t) \in \begin{cases} Q_I, & \text{if } 0^\circ \leq \arctan\left(\frac{\Delta y_{(t,\Delta t)}^k}{\Delta x_{(t,\Delta t)}^k}\right) < 90^\circ \\ Q_{II}, & \text{if } 90^\circ \leq \arctan\left(\frac{\Delta y_{(t,\Delta t)}^k}{\Delta x_{(t,\Delta t)}^k}\right) < 180^\circ \\ Q_{III}, & \text{if } 180^\circ \leq \arctan\left(\frac{\Delta y_{(t,\Delta t)}^k}{\Delta x_{(t,\Delta t)}^k}\right) < 270^\circ \\ Q_{IV}, & \text{if } 270^\circ \leq \arctan\left(\frac{\Delta y_{(t,\Delta t)}^k}{\Delta x_{(t,\Delta t)}^k}\right) < 360^\circ \end{cases}. \quad (\text{A.21})$$

The centre of each quadrant is calculated according to

$$(x_q(t), y_q(t)) = \left( \frac{1}{|q|} \sum_{\forall \Delta x_{(t,\Delta t)}^k \in q} \Delta x^k(t), \frac{1}{|q|} \sum_{\forall \Delta y_{(t,\Delta t)}^k \in q} \Delta y^k(t) \right) \quad (\text{A.22})$$

for  $q \in \{Q_I, Q_{II}, Q_{III}, Q_{IV}\}$ . The pattern indicative of expansion (or contraction) is classified using the relative positions of the quadrant centres. A frame is considered to contain expansion if  $Q_{II}$  is to the left of  $Q_I$ ,  $Q_{III}$  is to the left of  $Q_{IV}$ ,  $Q_{II}$  is above  $Q_{III}$  and  $Q_I$  is above  $Q_{IV}$ ; see Figure A.3a. Conversely, a frame is considered to contain contraction if  $Q_{IV}$  is to the left of  $Q_{III}$ ,  $Q_I$  is to the left of  $Q_{II}$ ,  $Q_{IV}$  is above  $Q_I$  and  $Q_{III}$  is above  $Q_{II}$ ; see Figure A.3b. The state of expansion (or

contraction) at frame  $t$  can be determined according to

$$t_s = \begin{cases} 1, & \text{if } (x_{Q_{II}}(t) < x_{Q_I}(t)) \wedge (x_{Q_{III}}(t) < x_{Q_{IV}}(t)) \\ & \wedge (y_{Q_{III}}(t) < y_{Q_{II}}(t)) \wedge (y_{Q_{IV}}(t) < y_{Q_I}(t)) \\ -1, & \text{if } (x_{Q_{IV}}(t) < x_{Q_{III}}(t)) \wedge (x_{Q_I}(t) < x_{Q_{II}}(t)) \\ & \wedge (y_{Q_{IV}}(t) < y_{Q_I}(t)) \wedge (y_{Q_{III}}(t) < y_{Q_{II}}(t)) \\ 0, & \text{otherwise} \end{cases}, \quad (\text{A.23})$$

where

$$t_s = \begin{cases} 1 & \text{implies expansion} \\ -1 & \text{implies contraction} \\ 0 & \text{neither expansion nor contraction} \end{cases}. \quad (\text{A.24})$$

The scale of expansion (or contraction),  $\rho_t$ , is determined by considering the magnitude of the displacement vector,

$$\rho_t = \max_{\forall k|_t} \{ \|\Delta \mathbf{x}_{(t,\Delta t)}^k\| \}. \quad (\text{A.25})$$

The largest magnitude is considered since it provides the most computationally stable measure of scale change between frames. For example, the error in an estimated scale change of 5 when the actual scale change is 5.1 is relatively small compared to an estimated scale change of 1 when the actual scale change is 0.9, even though the error is 0.1 in both cases. Here, reliance on a single (maximal) value is reasonable, because

unreliable trajectories should already have been removed by the pruning in Section A.1. Consequently, the scale of expansion (or contraction) at frame  $t$  is determined by comparing the ratio of the widths between frames (cf. Figure A.4),

$$s_t = \frac{w_t}{w'_{t+1}} = \begin{cases} \frac{w_t}{w_t + 2\sqrt{\rho t}}, & \text{if } t_s = 1 \\ \frac{w_t}{w_t - 2\sqrt{\rho t}}, & \text{if } t_s = -1, \\ 1, & \text{if } t_s = 0 \end{cases} \quad (\text{A.26})$$

where  $w_t$  is the width at frame  $t$  and  $w'_{t+1}$  is the estimated width at frame  $t + 1$ . Note that the calculation for computing the scale of expansion (or contraction) is not limited to widths, but can be found in a similar manner using heights. Finally, the scale of expansion (or contraction),  $s_t$ , is used to remove the expansion (or contraction) imposed by the camera operator by scaling each of the tracked point,  $\mathbf{x}^k$ , of trajectory  $\mathcal{T}^k$  at frame  $t$  according to

$$\hat{\mathbf{x}}^k(t) = s_t \mathbf{x}^k(t). \quad (\text{A.27})$$

In the general case, where the focus of expansion (or contraction) is not at the centre of the image, as shown in Figure A.3c, multiplying each frame by  $s_t$  will incorrectly remove the expansion (or contraction) and further distort the imagery. Thus, the centre of the expansion (or contraction),  $\mathbf{x}_z(t) = (x_z(t), y_z(t))$ , must be taken into account, which is computed using the median of the quadrant centres,




Figure A.4: Scale factor. The scale factor between frames can be determined using the width and the maximum magnitude of the displacement vectors by comparing the width of the old frame (solid rectangle) to the new frame (dashed rectangle).

(A.22), according to

$$\mathbf{x}_z(t) = (\text{med}_{v_q}\{x_q(t)\}, \text{med}_{v_q}\{y_q(t)\}). \quad (\text{A.28})$$

Finally, once the scale change and the shift in zoom centre has been determined, expansion (or contraction) is reverted from the recovered trajectories by multiplying the recovered iDTs by the scale change, then shifting by the centre,

$$\hat{\mathbf{x}}^k(t) = s_t \mathbf{x}^k(t) - \mathbf{x}_z(t). \quad (\text{A.29})$$



## Frames without Camera Motion

Presence of camera motion can be identified by detecting trajectories that are likely to correspond to camera motion. Conversely, absence of camera motion can be identified by detecting trajectories that are likely to correspond to static motion. Frames that are predominantly occupied by static trajectories (see Section A.1.1) are likely to correspond to frames without camera motion. That is, suppose  $\{(x^k(t), y^k(t)) | \mathcal{T}^k\}$  denotes a set of points present at frame  $t$  that belong to static trajectory  $\mathcal{T}^k$ . Frame  $t$  does not contain camera motion if

$$\left(\frac{w}{2} - \dot{u}_x(t) < \dot{\tau}_x\right) \wedge \left(\frac{h}{2} - \dot{u}_y < \dot{\tau}_y\right), \quad (\text{A.30})$$

where  $(\dot{\tau}_x, \dot{\tau}_y)$  correspond to horizontal and vertical vertices of the static trajectories that occupy the frame at frame  $t$ . The vertices of the static trajectories is estimated by taking the standard deviation,  $(\dot{u}_x(t), \dot{u}_y(t)) = (2\dot{\sigma}_x(t), 2\dot{\sigma}_y(t))$ , of the points that are present at frame  $t$ . The thresholds,  $(\dot{\tau}_x, \dot{\tau}_y)$ , are defined in terms of the resolution of the video (e.g.,  $(\dot{\tau}_x, \dot{\tau}_y) = (\frac{w}{15}, \frac{h}{15})$ ). The AND condition is used in (A.30) instead of OR as in (A.19), because (A.30) will null frames that were originally identified to contain camera motion. Thus, a stronger condition,  $\wedge$ , is applied to ensure that frames with camera motion are not incorrectly identified as a static frame. These frames are further confirmed as static if the absolute of the average displacement of the static trajectories are less than some threshold. That is, if

$$\left(\left|\frac{1}{N(t)} \sum_{k=1}^{N(t)} \Delta \dot{x}_{(t,1)}^k\right| < \dot{\tau}_x\right) \wedge \left(\left|\frac{1}{N(t)} \sum_{k=1}^{N(t)} \Delta \dot{y}_{(t,1)}^k\right| < \dot{\tau}_y\right), \quad (\text{A.31})$$

where  $N(t)$  is the number of static trajectories that are present at frame  $t$ , and  $(\hat{\tau}_x, \hat{\tau}_y) = (0.1, 0.1)$  is some small pre-defined value, then frame  $t$  is a static frame.

### **Integrated Trajectory Stabilization**

To stabilize a general set of input trajectories that may contain global shift (pan/tilt), linear expansion/contraction (zoom in/out), its combination or neither, a way to integrate global shift, expansion/contraction estimations and cancellations is considered, as follows. First, all frames are checked for absence of camera motion. Frames that meet the diagnostics of no camera motion are not processed further for trajectory stabilization. Second, all remaining frames are processed to stabilize the trajectories against global shift (pan/tilt). Third, the frames are processed to stabilize the trajectories against expansion/contraction (zoom in/out).

## **A.4 Elongating Trajectories**

To overcome the drifting effect when tracking points, trajectories estimated as iDTs are recommended to span  $L = 15$  or  $20$  frames for videos with a frame rate of 30 fps [121]. This limitation yields short-term trajectories, which are often insufficient to obtain a reliable measurement of trajectory characteristics, such as oscillation amplitude. Thus, a more reliable set of trajectories for the biolocomotion detector are obtained by concatenating iDTs across frames to gain elongated trajectories. In particular, a set of trajectories that span a fixed number of frames is concatenated by considering (i) the spatiotemporal proximities and (ii) the appearance likeness

of the trajectories. That is, for trajectory  $\mathcal{T}^k(t)$  that begins at frame  $t$  and spans  $L^k$  frames, the proximity criterion is considered by seeking trajectories that *begin* at frame  $s \in [t + L^k, t + L^k + \delta t]$  for some constant  $\delta t$ , whose spatial distance between the estimated point of  $\mathcal{T}^k$  at frame  $s$  and the first point of  $\mathcal{T}^j$  is less than some pre-defined constant,  $\tau_{sp}$ , according to

$$d_{sp}(\hat{\mathbf{x}}_{L^k+dt}^k, \mathbf{x}_1^j) < \tau_{sp}, \quad (\text{A.32})$$

where  $\hat{\mathbf{x}}_{L^k+dt}^k$  is an estimated point of  $\mathcal{T}^k$  at frame  $s = t + L^k + dt$  for  $0 \leq dt \leq \delta t$ . The spatial distance between two points,  $\mathbf{x}_i^j$  and  $\mathbf{x}_l^k$ , is defined using the  $L_2$ -norm according to

$$d_{sp}(\mathbf{x}_i^j, \mathbf{x}_l^k) = d_{sp}((x_i^j, y_i^j), (x_l^k, y_l^k)) = \sqrt{(x_l^k - x_i^j)^2 + (y_l^k - y_i^j)^2}. \quad (\text{A.33})$$

The appearance similarity between two trajectories can be considered using HOGs [31]. That is, if the distance between the corresponding HOGs,  $H^k$  and  $H^j$  for trajectories  $\mathcal{T}^k$  and  $\mathcal{T}^j$ , respectively, is less than some pre-defined constant,  $\tau_H$ :

$$d(H^k, H^j) < \tau_H, \quad (\text{A.34})$$

then trajectories  $\mathcal{T}^k$  and  $\mathcal{T}^j$  are likely to correspond to the same point of an object. The distance between HOGs,  $d(H^k, H^j)$ , is defined using the  $L_2$ -norm according to

$$d(H^k, H^j) = \sqrt{\sum_i (H^j(i) - H^k(i))^2}, \quad (\text{A.35})$$

where  $H^k = [H^k(i)] \in \mathbb{R}^n$ . Trajectory  $\mathcal{T}^j$  that satisfies both (A.32) and (A.34) is concatenated with trajectory  $\mathcal{T}^k$ .

Trajectory concatenation is not limited to the connection of two constant length trajectories, but can be applied iteratively to concatenate arbitrarily many trajectories provided (A.32) and (A.34) are satisfied. Once two trajectories are concatenated, HOG is updated to the trajectory that occurred (temporally) later. That is, if  $\mathcal{T}^{kj}$  denotes the concatenation of trajectories  $\mathcal{T}^k$  and  $\mathcal{T}^j$ , where  $\mathcal{T}^k$  occurred earlier than  $\mathcal{T}^j$ , then  $H^{kj} = H^j$ .

Currently, constant value of  $\delta t = \frac{f}{6}$  for frame rate  $f$  is used as the temporal threshold,  $\tau_{sp} = 2.5 * S$  for sampling width  $S$  as the threshold for the spatial distance,  $\tau_H = 0.75$  as the threshold for appearance likeness, and  $n = 96 = 8 \times 2 \times 2 \times 3$  for the dimension of the HOG descriptor.

# Appendix B

## Alternative Oscillation Measures

In addition to the frequency-based calculation presented in the main dissertation (Section 3.3.2) to obtain the oscillation measure,  $\mathcal{M}_\Omega$ , the amplitude of a detrended trajectory can also be calculated using integrals and differentials.

The amplitude of a trajectory can be computed via *integrals* by considering the average of the maximum value and the absolute of the minimum value of the detrended trajectory if the overall integral is close to zero; see Figure B.1a. That is, for a detrended trajectory  $\tilde{\mathcal{T}}^k$ , its amplitude is defined as the average of the maximum and minimum values of  $\tilde{\mathcal{T}}^k$ , according to

$$a_\Sigma^k = \exp\left(-\frac{(\sum_l \tilde{y}_l^k)^2}{2}\right) \left(\frac{\max_l \tilde{y}_l^k + |\min_l \tilde{y}_l^k|}{2}\right). \quad (\text{B.1})$$

The exponential function in (B.1) ensures that the amplitude is more recognized as the overall sum is close to zero, hence the use of an exponential weighting function.

Alternatively, the amplitude of a trajectory can be evaluated using *differentials* by considering the average of vertical values where its slope is horizontal; see Figure B.1b. That is, the amplitude of a detrended trajectory  $\tilde{\mathcal{T}}^k$  is defined as a weighted sum of vertical values whose first-order derivative is near zero according to

$$a_{\Delta}^k = w_c \frac{\sum_l w_s \tilde{y}_l^k}{\sum_l w_s}, \quad (\text{B.2})$$

where  $w_s$  is a weight that accounts for the slope and  $w_c$  is a weight that accounts for the periodicity of  $\tilde{\mathcal{T}}^k$ . These weights are defined using Gaussian-like functions, as follows.

To ensure a larger weight is assigned to values with near horizontal slopes (i.e.,  $(\tilde{\mathcal{T}}^k)' \approx 0$ ), a Gaussian-like function centred at 0 with a height of 1 is used to define  $w_s$  according to

$$w_s = \exp\left(-\frac{[(\tilde{y}_l^k)']^2}{2(0.1)^2}\right). \quad (\text{B.3})$$

$w_s$  assigns a value close to 1 if the slope of  $\tilde{\mathcal{T}}^k$  is near zero. Empirically, a width of 0.1 was found effective for accepting how horizontal  $\tilde{\mathcal{T}}^k$  is.

To ensure periodicity of  $\tilde{\mathcal{T}}^k$ ,  $w_c$  is a function defined to favour even number of concavities. Thus,  $w_c$  is defined as the sum of second-derivative of  $\tilde{\mathcal{T}}^k$  using a

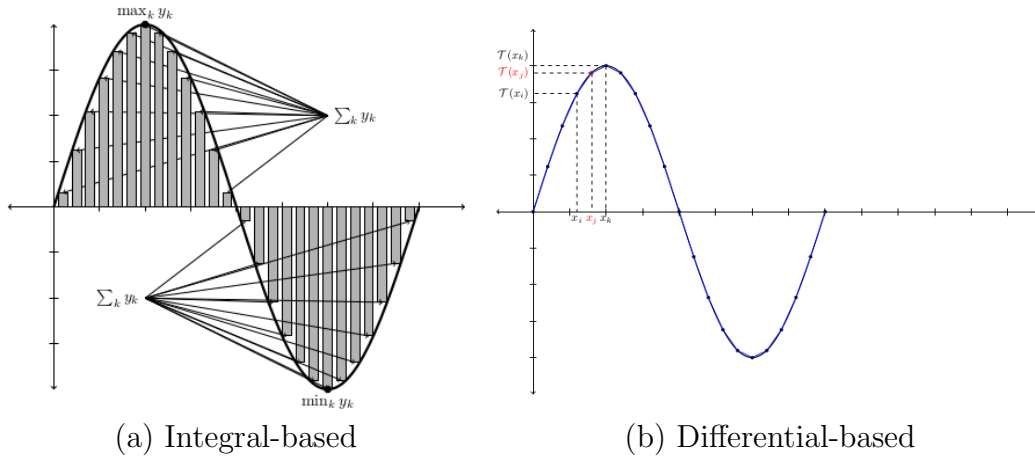


Figure B.1: Alternative measures of amplitude. See accompanying text in this appendix for details.

Gaussian-like function centred at 0 with a height of 1 according to

$$w_c = \exp\left(-\frac{[\sum_l w_s(\tilde{y}_l^k)']^2}{2}\right). \quad (\text{B.4})$$

$w_c$  assigns a value close to 1 when there are even number of concavities (i.e., the sum of the second derivative of  $\tilde{\mathcal{T}}^k$  is approximately 0). Empirically, a width of 1 was found effective for determining the concavity evenness of  $\tilde{\mathcal{T}}^k$ .


Based on preliminary experiments, the frequency-based approach was selected over the integral- and differential-based approaches for the work considered in this dissertation.

# Appendix C

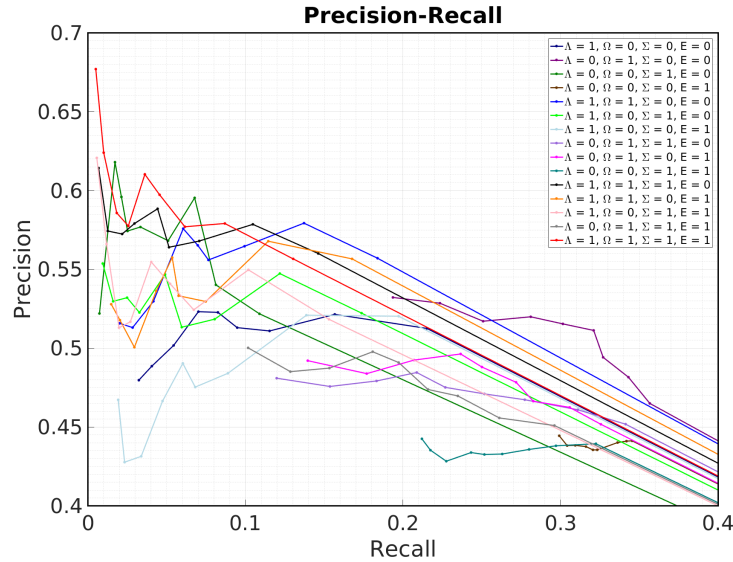
## Additional Empirical Results

In addition to the quantitative results presented in the main dissertation for the proposed biolocomotion detection algorithm (Section 4.6.2), this appendix presents more detailed results on the performance of individual algorithm components and their combinations on the BOLD dataset. Figure C.1 presents PR curves for all algorithm component combinations. Table C.1 shows precision values as detection threshold varies for all algorithm component combinations.





(a) PR curve in full scale.



(b) PR curve for recall  $\in [0, 0.4]$  and precision  $\in [0.4, 0.7]$ .

Figure C.1: Precision-Recall (PR) curves of the proposed algorithm. ‘1’ indicates enabled and ‘0’ indicates disabled components of locomotion ( $\Lambda$ ), oscillation ( $\Omega$ ), asymmetry ( $\Sigma$ ), and extrinsic motion dissimilarity ( $E$ ). (a) a full PR curve is provided with both precision and recall in the ranges  $[0, 1]$  and (b) the same PR curve, as (a), restricted to recall values within  $[0, 0.4]$  and precision within  $[0.4, 0.7]$  as the majority of the points are focused in a select region of the PR curve.

Components				$\tau_B$									
$\Lambda$	$\Omega$	$\Sigma$	$E$	0	0.1	0.2	0.3	0.4	0.5	0.6	0.7	0.8	0.9
1	0	0	0	0.1123	0.5125	0.5214	0.5109	0.5130	0.5227	0.5231	0.5017	0.4886	0.4797
0	1	0	0	0.1121	0.4650	0.4816	0.4942	0.5113	0.5154	0.5199	0.5171	0.5285	0.5322
0	0	1	0	0.1123	0.5218	0.5402	0.5953	0.5682	0.5769	0.5743	0.5959	0.6180	0.5220
0	0	0	1	0.1123	0.4412	0.4410	0.4401	0.4355	0.4355	0.4376	0.4384	0.4382	0.4444
1	1	0	0	0.1123	0.5571	0.5793	0.5647	0.5559	0.5652	0.5758	0.5297	0.5131	0.5157
1	0	1	0	0.1123	0.5222	0.5473	0.5184	0.5133	0.5467	0.5226	0.5320	0.5297	0.5537
1	0	0	1	0.1123	0.5203	0.5209	0.4840	0.4753	0.4904	0.4665	0.4314	0.4277	0.4672
0	1	1	0	0.1123	0.4518	0.4609	0.4673	0.4710	0.4751	0.4846	0.4792	0.4757	0.4810
0	1	0	1	0.1123	0.4517	0.4624	0.4664	0.4784	0.4880	0.4963	0.4924	0.4839	0.4921
0	0	1	1	0.1123	0.4394	0.4381	0.4358	0.4329	0.4326	0.4338	0.4283	0.4353	0.4425
1	1	1	0	0.1123	0.5601	0.5785	0.5680	0.5640	0.5884	0.5791	0.5724	0.5743	0.6143
1	1	0	1	0.1123	0.5566	0.5678	0.5296	0.5332	0.5575	0.5362	0.5006	0.5177	0.5279
1	0	1	1	0.1123	0.5183	0.5497	0.5243	0.5408	0.5547	0.5167	0.5130	0.5660	0.6207
0	1	1	1	0.1123	0.4509	0.4557	0.4697	0.4738	0.4909	0.4977	0.4874	0.4851	0.5002
1	1	1	1	0.1123	0.5566	0.5791	0.5770	0.5973	0.6103	0.5775	0.5858	0.6240	0.6770

Table C.1: Precision values at specified threshold  $\tau_B$  for components: locomotion ( $\Lambda$ ), oscillation ( $\Omega$ ), asymmetry ( $\Sigma$ ), and extrinsic motion dissimilarity ( $E$ ).

# References

- [1] J.K. Aggarwal and Q.Cai. Human Motion Analysis: A Review. *Computer Vision and Image Understanding*, 73(3):428–440, 1999.
- [2] H. Akaike. Information Theory and an Extension of the Maximum Likelihood Principle. In *2nd International Symposium on Information Theory*, pages 199–213, 1973.
- [3] R.M. Alexander. *Principles of Animal Locomotion*. Princeton University Press, 2003.
- [4] C.D. Barclay, J.E. Cutting, and L.T. Kozlowski. Temporal and spatial factors in gait perception that influence gender recognition. *Perception & Psychophysics*, 23(2):145–152, 1978.
- [5] H. Bay, T. Tuytelaars, and L. Van Gool. Surf: Speeded up robust features. In *European Conference on Computer Vision (ECCV)*, 2016.
- [6] S. Beery, G. Van Horn, and P. Perona. Recognition in Terra Incognita. In *European Conference on Computer Vision (ECCV)*, 2018.

- [7] A.M. Berg and A.A. Biewener. Wing and Body Kinematics of Takeoff and Landing Flight in the Pigeon (*Columba livia*). *Journal of Experimental Biology*, 213:1651–1658, 2010.
- [8] P. Bideau and E. Learned-Miller. It’s Moving! A Probabilistic Model for Causal Motion Segmentation in Moving Camera Videos. In *European Conference on Computer Vision (ECCV)*, 2016.
- [9] A.A. Biewener. *Animal Locomotion*. Oxford University Press, 2003.
- [10] G.P. Bingham. Dynamics and the Problem of Visual Event Recognition. In R.F. Port and T. van Gelder, editors, *Mind as Motion: Explorations in the Dynamics of Cognition*. The MIT Press, 1995.
- [11] R. Blake. Cats Perceive Biological Motion. *Psychological Science*, 4(1):54–57, 1993.
- [12] R. Blickhan and R.J. Full. Locomotion Energetics of the Ghost Crab: II. Mechanics of the Centre of Mass During Walking and Running. *Journal of Experimental Biology*, 130:155–174, 1987.
- [13] R. Blickhan and R.J. Full. Similarity in Multilegged Locomotion: Bouncing like a Monopode. *Journal of Comparative Physiology A*, 173:509–517, 1993.
- [14] A. Borji, M.M. Cheng, H. Jiang, and J. Li. Salient Object Detection: A Benchmark. *IEEE Transactions on Image Processing*, 24(12), 2015.
- [15] E.J. Borowski and J.M. Borwein. *The HarperCollins Dictionary of Mathematics*. HarperCollins, 1991.

- [16] A. Briassouli and N. Ahuja. Extraction and Analysis of Multiple Periodic Motions in Video Sequences. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 29(7):1244–1261, 2007.
- [17] T. Brox and J. Malik. Object Segmentation by Long Term Analysis of Point Trajectories. In *European Conference on Computer Vision (ECCV)*, pages 282–295, 2010.
- [18] M.A. Brubaker, D.J. Fleet, and A. Hertzmann. Physics-Based Person Tracking Using Simplified Lower-Body Dynamics. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007.
- [19] M. Calisti, G. Picardi, and C. Laschi. Fundamentals of Soft Robot Locomotion. *Journal of the Royal Society Interface*, 14:1–16, 2017.
- [20] K. Cannons. A Review of Visual Tracking. Technical Report CSE-2008-07, York University, 2008.
- [21] J. Carreira and A. Zisserman. Quo Vadis, Action Recognition? A New Model and the Kinetics Dataset. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4724–4733, 2017.
- [22] D.H.F. Chang, H. Ban, Y. Ikegaya, I. Fujita, and N.F. Troje. Cortical and Subcortical Responses to Biological Motion. *Neuroimage*, 174:87–96, 2018.
- [23] D.H.F. Chang and N. Troje. Perception of Animacy and Direction from Local Biological Motion Signals. *Journal of Vision*, 8(5):1–10, 2008.

- [24] D.H.F. Chang and N.F. Troje. Acceleration carries the local inversion effect in biological motion perception. *Journal of Vision*, 9(19):1–17, 2009.
- [25] G. Chen, T.X. Han, Z. He, R. Kays, and T. Forrester. Deep Convolutional Neural Network Based Species Recognition for Wild Animal Monitoring. In *IEEE International Conference on Image Processing (ICIP)*, 2014.
- [26] M.M. Cheng, N.J. Mitra, X. Huang, P.H.S. Torr, and S.M. Hu. Global Contrast Based Salient Region Detection. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 37(3), 2015.
- [27] S. Collins, A. Ruina, R. Tedrake, and M. Wisse. Efficient Bipedal Robots Based on Passive-Dynamic Walkers. *Science*, 307(5712):1082–1085, 2005.
- [28] C. Cortes and V. Vapnik. Support-Vector Networks. *Machine Learning*, 20(3):273–297, 1995.
- [29] R.D. Crowninshield and R.A. Brand. A Physiologically based Criterion of Muscle Force Prediction in Locomotion. *Journal of Biomechanics*, 14(11):793–801, 1981.
- [30] R. Cutler and L.S. Davis. Robust real-time periodic motion detection, analysis, and applications. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 22:781–796, 2000.
- [31] N. Dalal and B. Triggs. Histogram of Oriented Gradients for Human Detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, volume 1, pages 886–893, 2005.

- [32] N. Dalal, B. Triggs, and C. Schmid. Human Detection Using Oriented Histograms of Flow and Appearance. In *European Conference on Computer Vision (ECCV)*, pages 428–441, 2006.
- [33] D.L. Davies and D.W. Bouldin. A Cluster Separation Measure. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 1(2), 1979.
- [34] R.C. de Amorim and C. Hennig. Recovering the Number of Clusters in Data sets with Noise Features using Feature Rescaling Factors. *Information Sciences*, 324:126–145, 2015.
- [35] W. Dittrich, S. Lea, J. Barrett, and P Gurr. Categorization of Natural Movements By Pigeons: Visual Concept Discrimination and Biological Motion. *Journal of the Experimental Analysis of Behavior*, 70(3):281–299, 1998.
- [36] W.H. Dittrich, T. Troscianko, S.E. Lea, and D. Morgan. Perception of emotion from dynamic point-light displays represented in dance. *Perception*, 25(6):727–738, 1996.
- [37] J.C. Dunn. A Fuzzy Relative of the ISODATA Process and Its Use in Detecting Compact Well-Separated Clusters. *Journal of Cybernetics*, 3(3):32–57, 1974.
- [38] C. Feichtenhofer, A. Pinz, and A. Zisserman. Convolutional Two-Stream Network Fusion for Video Action Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [39] M. Fiaz, A. Mahmood, and S.K. Jung. Tracking Noisy Targets: A Review of Recent Object Tracking Approaches. In *eprint arXiv:1802.03098*, 2018.

- [40] M. Fiaz, A. Mahmood, and S.K. Jung. The 2019 DAVIS Challenge on VOS: Unsupervised Multi-Object Segmentation. In *eprint arXiv:1905.00737*, 2019.
- [41] M.A. Fischler and R.C. Bolles. Random Sample Consensus: A Paradigm for Model Fitting with Applications to Image Analysis and Automated Cartography. *Communications of the ACM*, 24(6), 1981.
- [42] F.E. Fish, C.M. Schreiber, K.W. Moored, G. Liu, H. Dong, and H. Bart-Smith. Hydrodynamic Performance of Aquatic Flapping: Efficiency of Underwater Flight in the Manta. *Aerospace*, 3(20), 2016.
- [43] H. Geyer, A. Seyfarth, and R. Blickhan. Compliant Leg Behaviour Explains Basic Dynamics of Walking and Running. *Proceedings of the Royal Society B: Biological Sciences*, 273(1603):2861–2867, 2006.
- [44] M.A. Giese and T. Poggio. Biologically Plausible Neural Model for the Recognition of Biological Motion and Actions. CBCL Paper 219, Massachusetts Institute of Technology, 2002.
- [45] M.A. Giese and T. Poggio. Neural Mechanisms for the Recognition of Biological Movements. *Nature Reviews Neuroscience*, 4:179–192, 2003.
- [46] M.A. Goodale and A.D. Milner. Separate visual pathways for perception and action. *Trends in Neurosciences*, 15(1):20–25, 1992.
- [47] F. Heider and M. Simmel. An Experimental Study of Apparent Behavior. *The American Journal of Psychology*, 57(2):243–259, 1944.



- [48] F.C. Heilbron, V. Escorcia, B. Ghanem, and J.C. Niebles. ActivityNet: A Large-Scale Video Benchmark for Human Activity Understanding. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 961–970, 2015.
- [49] M. Hirai, D.H. Chang, D.R. Saunders, and N.F. Troje. Body Configuration Modulates the Usage of Local Cues to Direction in Biological-Motion Perception. *Psychological Science*, 22(12):1543–1549, 2011.
- [50] D.D. Hoffman and B.E. Flinchbaugh. The Interpretation of Biological Motion. *Biological Cybernetics*, 42:195–204, 1982.
- [51] G. Van Horn, O. Mac Aodha, Y. Song, Y. Cui, C. Sun, A. Shepard, H. Adam, P. Perona, and S. Belongie. The iNaturalist Species Classification and Detection Dataset. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 8769–8778, 2018.
- [52] G. Van Horn, S. Branson, R. Farrell, S. Haber, J. Barry, P. Ipeirotis, P. Perona, and S. Belongie. Building a bird recognition app and large scale dataset with citizen scientists: The fine print in fine-grained dataset collection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 595–604, 2015.
- [53] H. Jhuang, J. Gall, S. Zuffi, C. Schmid, and M.J. Black. Towards Understanding Action Recognition. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3192–3199, 2013.

- [54] Y.-G. Jiang, J. Liu, A. Zamir, G. Toderici, I. Laptev, M. Shah, and R. Sukthankar. THUMOS challenge: Action recognition with a large number of classes. <http://crcv.ucf.edu/THUMOS14/>, 2014.
- [55] G. Johansson. Visual Perception of Biological Motion and a Model for Its Analysis. *Perception & Psychophysics*, 14(2):201–211, 1973.
- [56] G. Johansson. Vector Analysis in Visual Perception of Rolling Motion. *Psychological Research*, 36(4):311–319, 1974.
- [57] M.H. Johnson. Biological Motion: A Perceptual Life Detector? *Current Biology*, 16(10):376–377, 2006.
- [58] J.P. Jones and L.A. Palmer. An Evaluation of the Two-Dimensional Gabor Filter Model of Simple Receptive Fields in Cat Striate Cortex. *Journal of Neurophysiology*, 58(6):1233–1258, 1987.
- [59] S.M. Kang and R.P. Wildes. Review of Action Recognition and Detection Methods. Technical report, York University, 2016.
- [60] S. Karra, M. Jelali, M.N. Karim, and A. Horch. Detection of oscillating control loops. In *Detection and Diagnosis of Stiction in Control Loops*. Springer, London, 2010.
- [61] J. Katz and D. Weihs. Large amplitude unsteady motion of a flexible slender propulsor. *Journal of Fluid Mechanics*, 90(4):713–723, 1979.
- [62] Y. Ke, R. Sukthankar, and M. Hebert. Event Detection in Crowded Videos. In *IEEE International Conference on Computer Vision (ICCV)*, pages 1–8, 2007.

- [63] M. Kilicarslan and J.Y. Zheng. Predict Vehicle Collision by TTC From Motion Using a Single Video Camera. *IEEE Transactions on Intelligent Transportation Systems*, 20(2):522–533, 2018.
- [64] Y. Kong and Y. Fu. Human Action Recognition and Prediction: A Survey. In *eprint arXiv:1806.11230*, 2018.
- [65] L.T. Kozlowski and J.E. Cutting. Recognizing the sex of a walker from a dynamic point-light display. *Perception & Psychophysics*, 21(6):575–580, 1977.
- [66] M. Kristan, J. Matas, A. Leonardis, T. Vojir, R. Pflugfelder, G. Fernández, G. Nebehay, F. Porikli, and L. Cehovin. A Novel Performance Evaluation Methodology for Single-Target Trackers. *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, 38(11):2137–2155, 2016.
- [67] M. LaBarbera. Why the Wheels Won’t Go. *The University of Chicago Press Journals for the American Society of Naturalists*, 121(3):395–408, 1983.
- [68] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld. Learning Realistic Human Actions from Movies. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [69] G.V. Lauder and J.L. Tangorra. Fish Locomotion: Biology and Robotics of Body and Fin-Based Movements. In R. Du and P. Valdivia y Alvarado Z. Li, K. Youcef-Toumi, editors, *Robot Fish*. Springer-Verlag Berlin Heidelberg, 2015.

- [70] M. Leo, G. Medioni, M. Trivedi, T. Kanade, and G.M. Farinella. Computer Vision for Assistive Technologies. *Computer Vision and Image Understanding*, 154:1–15, 2017.
- [71] O. Levy and L. Wolf. Live Repetition Counting. In *IEEE International Conference on Computer Vision (ICCV)*, pages 3020–3028, 2015.
- [72] K.H. Lin, P. Khorrami, J. Wang, M. Hasegawa-Johnson, and T.S. Huang. Foreground object detection in highly dynamic scenes using saliency. In *IEEE International Conference on Image Processing (ICIP)*, pages 1125–1129, 2014.
- [73] T. Liu, J. Sun, N.N. Zheng, X. Tang, and H.Y. Shum. Learning to Detect A Salient Object. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2007.
- [74] H. Marvi and D. Hu. Friction enhancement in concertina locomotion of snakes. *Journal of the Royal Society Interface*, 1:1–14, 2012.
- [75] G. Mather and S. West. Recognition of Animal Locomotion from Dynamic Point-Light Displays. *Perception*, 22:759–766, 1993.
- [76] T. McGeer. Passive Dynamic Walking. *The International Journal of Robotics Research*, 9(2):62–82, 1990.
- [77] T.A. McMahon. The Role of Compliance in Mammalian Running Gaits. *The Journal of Experimental Biology*, 115:263–282, 1985.
- [78] T.A. McMahon and G.C. Cheng. The Mechanics of Running: How Does Stiffness Couple with Speed? *Journal of Biomechanics*, 23:65–78, 1990.

- [79] R. Messing, C. Pal, and H. Kautz. Activity Recognition Using the Velocity Histories of Tracked Keypoints. In *IEEE International Conference on Computer Vision (ICCV)*, pages 104–111, 2009.
- [80] A. Miguel, S. Beery, E. Flores, L. Klemesrud, and R. Bayrakcismith. Finding areas of motion in camera trap images. In *IEEE International Conference on Image Processing (ICIP)*, pages 1334–1338, 2016.
- [81] T.B. Moeslund, G. Thomas, and A. Hilton. *Computer Vision in Sports*. Springer, 2014.
- [82] M.M. Muchisky and G.P. Bingham. Trajectory Forms as a Source of Information about Events. *Perception & Psychophysics*, 64(1):15–31, 2002.
- [83] S.A. Niyogi and E.H. Adelson. Analyzing and recognizing walking figures in xyt. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1994.
- [84] A.F. O’Connell, J.D. Nichols, and K. Ullas. *Camera Traps in Animal Ecology: Methods and Analyses*. Springer, 2011.
- [85] D. Park, C.L. Zitnick, D. Ramanan, and P. Dollar. Exploring Weak Stabilization for Motion Feature Extraction. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2882–2889, 2013.
- [86] F. Perazzi, P. Krahenbuhl, Y. Pritch, and A. Hornung. Saliency Filters: Contrast based Filtering for Salient Region Detection. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2012.

- [87] F. Perronnin, J. Sanchez, and T. Mensink. Improving the Fisher kernel for large-scale image classification. In *European Conference on Computer Vision (ECCV)*, 2010.
- [88] J. Platt. Probabilistic outputs for support vector machines and comparisons to regularized likelihood methods. In A. Smola, P. Bartlerr, and B. Schölkopf, editors, *Advances in Large Margin Classifiers*. MIT Press, 1999.
- [89] T. Poggio and S. Edelman. A Network that Learns to Recognize Three-Dimensional Objects. *Nature*, 343:263–266, 1990.
- [90] R. Polana and R.C. Nelson. Detection and recognition of periodic, nonrigid motion. *International Journal of Computer Vision (IJCV)*, 23(3):261–282, 1997.
- [91] F.E. Pollick, C. Fidopiastis, and V. Braden. Recognising the style of spatially exaggerated tennis serves. *Perception*, 30(3):323–338, 2001.
- [92] S.J. Portugal, T.Y. Hubel, J. Fritz, S. Heese, D. Trobe, B. Voelkl, S. Hailes, A.M. Wilson, and J.R. Usherwood. Upwash Exploitation and Downwash Avoidance by Flap Phasing in Ibis Formation Flight. *Nature*, 505:399–402, 2014.
- [93] M. Raibert, M. Chepponis, and B. Brown. Running of four legs as though they were one. *IEEE Transactions on Robotics and Automation*, 2:70–82, 1986.
- [94] M.H. Raibert. *Legged Robots that Balance*. The MIT Press, 1986.

- [95] Y. Ran, I. Weiss, Q. Zheng, and L.S. Davis. Pedestrian Detection via Periodic Motion Analysis. *International Journal of Computer Vision (IJCV)*, 71(2):143–160, 2007.
- [96] D. Regan and K.I. Beverley. Looming detectors in the human visual pathway. *Vision Research*, 18(4):415–421, 1978.
- [97] M. Riesenhuber and T. Poggio. Hierarchical Models of Object Recognition in Cortex. *Nature Neuroscience*, 2:1019–1025, 1999.
- [98] M. Rodriguez, J. Ahmed, and M. Shah. Action MACH: A Spatio-temporal Maximum Average Correlation Height Filter for Action Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2008.
- [99] P.J. Rousseeuw. Silhouettes: A Graphical Aid to the Interpretation and Validation of Cluster Analysis. *Journal of Computational and Applied Mathematics*, 20:53–65, 1987.
- [100] T.F.H. Runia, C.G.M. Snoek, and A.W.M. Smeulders. Real-World Repetition Estimation by Div, Grad and Curl. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9009–9017, 2018.
- [101] B.J. Scholl and P.D. Tremoulet. Perceptual Causality and Animacy. *Trends in Cognitive Sciences*, 4(8):299–309, 2000.
- [102] S.M. Seitz and C.R. Dyer. View-Invariant Analysis of Cyclic Motion. *International Journal of Computer Vision (IJCV)*, 25(3):231–251, 1997.

- [103] E. Shavit and A. Jepson. Motion understanding from qualitative visual dynamics. In *IEEE Workshop on Qualitative Vision*, pages 82–88, 1993.
- [104] J. Shi and C. Tomasi. Good Features to Track. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 1994.
- [105] K. Soomro and A.R. Zamir. *Computer Vision in Sports*, chapter 9 - Action Recognition in Realistic Sports Videos, pages 181–208. Springer, 2014.
- [106] K. Soomro, A.R. Zamir, and M. Shah. UCF101: A Dataset of 101 Human Actions Classes from Videos in the Wild. Technical Report CRCV-TR-12-01, University of Central Florida, 2012.
- [107] M. Srinivasan and A. Ruina. Computer Optimization of a Minimal Biped Model Discovers Walking and Running. *Nature*, 439:72–75, 2006.
- [108] J. Sun, X. Wu, S. Yan, L.F. Cheong, T.S. Chua, and J. Li. Hierarchical Spatio-Temporal Context Modeling for Action Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2004–2011, 2009.
- [109] A. Swanson, M. Kosmala, C. Lintott, R. Simpson, A. Smith, and C. Packer. Snapshot Serengeti, High-Frequency Annotated Camera Trap Images of 40 Mammalian Species in an African Savanna. *Scientific Data*, 2, 2015.
- [110] G. Thomas, R. Gade, T.B. Moeslund, P. Carr, and A. Hilton. Computer Vision for Sports: Current Applications and Research Topics. *Computer Vision and Image Understanding*, 159:3–18, 2017.



- [111] M.S. Triantafyllou, G.S. Triantafyllou, and D.K.P. Yue. Hydrodynamics of Fishlike Swimming. *Annual Review of Fluid Mechanics*, 32:33–53, 2000.
- [112] N.F. Troje. Decomposing Biological Motion: A Framework for Analysis and Synthesis of Human Gait Patterns. *Journal of Vision*, 2:371–387, 2002.
- [113] N.F. Troje. Biological Motion Perception. In *The Senses: A Comprehensive References*, pages 231–238. Elsevier, 2008.
- [114] N.F. Troje. What is Biological Motion?: Definition, Stimuli, and Paradigms. In *Social Perception: Detection and Interpretation of Animacy, Agency, and Intention*. M.D. Rutherford and V.A. Kuhlmeier, 2012.
- [115] N.F. Troje and C. Westhoff. The Inversion Effect in Biological Motion Perception: Evidence for a Life Detector. *Current Biology*, 16:821–824, 2006.
- [116] P.S. Tsai, M. Shah, K. Keiter, and T. Kasparis. Cyclic motion detection for motion based recognition. *Pattern Recognition*, 27(12):1591–1603, 1994.
- [117] G. Vallortigara and L. Regolin. Gravity bias in the interpretation of biological motion by inexperienced chicks. *Current Biology*, 16(8):R279–R280, 2006.
- [118] J.C. van Gemert, M. Jain, E. Gati, and C.G.M. Snoek. APT: Action localization Proposals from dense Trajectories. In *Proceedings of the British Machine Vision Conference (BMVC)*, pages 177.1–177.12, 2015.
- [119] D.D. Wackerly, W. Mendenhall III, and R.L. Scheaffer. *Mathematical Statistics with Applications*. Thomson Brooks/Cole, seventh edition, 2008.

- [120] H. Wang, A. Klaser, C. Schmid, and C.L. Liu. Action recognition by dense trajectories. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3169–3176, 2011.
- [121] H. Wang, A. Klaser, C. Schmid, and C.L. Liu. Action recognition with improved trajectories. In *IEEE International Conference on Computer Vision (ICCV)*, page 3551—3558, 2013.
- [122] H. Wang, A. Klaser, C. Schmid, and C.L. Liu. Dense Trajectories and Motion Boundary Descriptors for Action Recognition. *International Journal of Computer Vision (IJCV)*, 103:60–79, 2013.
- [123] P. Welinder, S. Branson, T. Mita, C. Wah, F. Schroff, S. Belongie, and P. Perona. Caltech-ucsd birds 200. Technical Report CNS-TR-2010-001, California Institute of Technology, 2010.
- [124] L. Wen and G. Lauder. Understanding Undulatory Locomotion in Fishes using an Inertia-Compensated Flapping Foil Robotic Device. *Bioinspiration & Biomimetics*, 8, 2013.
- [125] M.J. Wilber, W.J. Scheirer, P. Leitner, B. Heflin, J. Zott, D. Reinke, D.K. Delaney, and T.E. Boulton. Animal recognition in the mojave desert: Vision tools for field biologists. In *IEEE Workshop on Applications of Computer Vision (WACV)*, 2013.

- [126] C. Xu, S.H. Hsieh, C. Xiong, and J.J. Corso. Can Humans Fly? Actions Understanding with Multiple Classes of Actors. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2264–2273, 2015.
- [127] R. Yao, G. Lin, S. Xia, J. Zhao, and Y. Zhou. Video Object Segmentation and Tracking: A Survey. In *eprint arXiv:1904.09172*, 2019.
- [128] X. Yu, J. Wang, R. Kays, P.A. Jansen, T. Wang, and T. Huang. Automated identification of animal species in camera trap images. *EURASIP Journal on Image and Video Processing*, 52, 2013.
- [129] Y. Yu and B. Tong. A Flow Control Mechanism in Wing Flapping with Stroke Asymmetry During Insect Forward Flight. *Acta Mechanica Sinica*, 21(3):218–227, 2005.
- [130] Z. Zhang, T.X. Han, and Z. He. Coupled Ensemble Graph-Cuts and Object Verification for Animal Segmentation from Highly Cluttered Videos. In *IEEE International Conference on Image Processing (ICIP)*, pages 2830–2834, 2015.
- [131] Z. Zhang, Z. He, G. Cao, and W. Cao. Animal Detection From Highly Cluttered Natural Scenes Using Spatiotemporal Object Region Proposals and Patch Verification. *IEEE Transactions on Multimedia*, 18(10):2079–2092, 2016.
- [132] Z. Zhang and N.F. Troje. 3D Periodic Human Motion Reconstruction from 2D Motion Sequences. *Neural Computation*, 19, 2007.