

ADDRESSING ZONING ABBREVIATION INCONSISTENCIES
IN THE GREATER TORONTO AREA (GTA)
THROUGH A UNIVERSAL ONTOLOGY

ARIAN MESBAHIAN

A THESIS SUBMITTED TO
THE FACULTY OF GRADUATE STUDIES
IN PARTIAL FULFILLMENT OF THE REQUIREMENTS
FOR THE DEGREE OF
MASTER OF APPLIED SCIENCE

GRADUATE PROGRAM IN CIVIL ENGINEERING
YORK UNIVERSITY
TORONTO, ONTARIO

January 2026

© Arian Mesbahian, 2026

ABSTRACT

Urban development in the Greater Toronto Area (GTA) faces growing challenges due to inconsistencies in zoning abbreviations and terminology across municipal jurisdictions. Similar zoning labels such as “R2” in Toronto and “R2-S” in Markham carry different land use permissions and development standards, creating semantic misalignments that confuse planners, developers, and regulators. These discrepancies lead to inefficiencies, legal ambiguities, and prolonged approval processes, ultimately hindering coordinated urban growth. At the same time, Ontario’s housing shortage has become a critical issue. The provincial government’s target of building 1.5 million new homes by 2031, or approximately 150,000 annually, remains far from reach, with only 81,300 units projected for 2024. Key obstacles include limited servicing allocations for water, sewage, and stormwater infrastructure, as well as slow and fragmented municipal approval systems. When combined with inconsistent zoning terminology, these challenges contribute to significant delays in housing delivery and regional planning.

To address these systemic barriers, this study introduces Zonology, a machine readable, ontology based framework designed to semantically standardize zoning terminologies across municipalities. Using the City of Toronto and the City of Markham as a proof of concept case studies, Zonology harmonizes more than 60 zoning categories and aligns over 150 land use permissions through a shared semantic structure. The ontology is developed using the Linked Open Terms (LOT) methodology and follows FAIR principles Findable, Accessible, Interoperable, and Reusable to ensure transparency, compatibility, and reusability across urban data systems.

Zonology models zoning designations, permitted land uses, development standards, and spatial relationships by integrating legal by-laws, planning regulations, and geospatial data. It enables semantic querying and interoperability with Geographical Information Systems (GIS) platforms, smart city applications, and automated planning systems. By bridging semantic and regulatory gaps, the framework enhances data driven decision making, improves intermunicipal collaboration, and supports consistent, scalable zoning governance. The development of Zonology represents a significant advancement in urban informatics and planning modernization, offering a replicable foundation for harmonized, intelligent, and future-ready urban development across the GTA and beyond.

ACKNOWLEDGMENTS

I would like to extend my sincere appreciation to my supervisor, Dr. Mojgan Jadidi, whose guidance, encouragement, and expertise have been central to the completion of this thesis. Her thoughtful advice and unwavering support have greatly shaped the direction and quality of this research.

My gratitude also goes to my committee member, Dr. Amin Sarang, for his constructive feedback and insightful comments, which provided valuable perspectives and strengthened the overall analysis of my work.

I am also thankful to my examiners, Dr. Kevin Zhang and Dr. Evalyna Bogdan, for their generously devoting the time to review this thesis and for offering important feedbacks that have helped refine the outcomes of this study.

I would like to also thank Mitacs, AECO Innovation Lab, and LandLogic for their support through funding, internship opportunities, and access to the LandLogic platform, which greatly contributed to the development of this research. I am especially grateful to Dr. Arash Shahi for his mentorship, guidance, and continued encouragement throughout this work.

Lastly, I wish to acknowledge the ongoing support of my family, friends, and colleagues. Their patience, understanding, and encouragement have been a constant source of motivation throughout this academic journey.

TABLE OF CONTENTS

ABSTRACT	ii
ACKNOWLEDGMENTS.....	iv
TABLE OF CONTENTS	v
LIST OF TABLES.....	ix
LIST OF FIGURES	xi
ABBREVIATIONS	xiii
CHAPTER 1: INTRODUCTION	1
1.1. MOTIVATION OF RESEARCH	1
1.2. PROBLEM STATEMENT	4
1.3. RESEARCH OBJECTIVES.....	8
1.4. RESEARCH CONTRIBUTIONS	9
1.5. STRUCTURE OF THESIS	10
CHAPTER 2: Litreture review.....	11
2.1. EXISTING URBAN AND INFRASTRUCTURE ONTOLOGIES.....	11
2.2. INTEGRATION OF EXISITNG ONTOLOGIES INTO THE ZONING ONTOLOGY FRAMEWORK.....	17
CHAPTER 3: METHODOLOGY.....	21
3.1. ontology specification	22

3.1.1. Acquiring domain knowledge	26
3.1.2. Creation of ORSD.....	27
3.1.3. Formulate Competency Questions (CQs)	29
3.2. Ontology conceptualization.....	33
3.2.1. Identifying Core Zoning Terms and Class Hierarchy	34
3.2.2. Define Object and Data Properties	36
3.2.3. Aligning terminology with real word planning workflows.....	42
3.2.4. Relationships (Object and Data properties) and Modeling patterns	44
3.3. Ontology refinement and Alignment.....	48
3.3.1. Reuse Elements From Existing Ontologies And Refine Structure	50
3.3.2. Prepare for OWL-DL compliance and finalize properties.....	51
3.4. Ontology implementation.....	52
3.5. ontological model optimization and redundancy deduction.....	56
3.5.1. Normalization of class and property hieracies.....	57
3.5.2. Aligment with External Vocabularies	58
3.5.3. Impact on Query Selectivity and Clarity	59
CHAPTER 4: Result and discussion	61
4.1. Case Studies	62
4.1.1. Listing Toronto Zoning Abbreviations with Their Full Names	63
4.1.2. Retrieving the Full Name for the EMP-GE Abbreviation	64

4.1.3. Identifying Parent and Equivalent Classes for R1	66
4.1.4. Retrieving Building Structures Permitted in Residential Zones	68
4.1.5. Listing Markham Zoning Abbreviations with Their Full Names	69
4.1.6. Identifying Parent and Equivalent Classes for Industrial (I) Zones.....	70
4.1.7. Identifying Parent and Equivalent Classes for Residential Semi-Detached (RS) Zones	71
4.1.8. Identifying Parent and Equivalent Classes for Mixed-Use (MU) Zones	72
4.1.9. Retrieving Building Structures Permitted in Industrial Zones.....	74
4.1.10. Retrieving Building Structures Permitted in Employment Zones	75
4.1.11. Over-Strict Equivalence in Industrial Zones (Query Blocked).....	76
4.1.12. Property Range Misspecification in Employment Zones (Query Blocked).....	77
4.1.13. Missing Conjunctive Semantics for Mixed Use Zones (Query Missed Result)	78
4.2. Evaluation metric	79
4.3. semantic reasoning and logical coherence assesment	81
4.3.1. Reasoner Configuration and Method	81
4.3.2. Outcomes of Consistency Checking	83
4.3.3. Interpretation of findings	85
CHAPTER 5: Conclusion	90

5.1. Summary	90
5.2. limitation of work	93
5.3. Future Research	95
Refferences	97

LIST OF TABLES

Table 2.1 Comparative Review of Existing Urban and Infrastructure Ontologies and Their Reuse in the Proposed Zoning Ontology	16
Table 2.2 Gap Analysis of Existing Urban/Planning Ontologies	20
Table 3.1 The proposed framework to develop Zonology	22
Table 3.2 Comparison of zoning abbreviations for the permitted use “Detached Dwelling” in Aurora and Ajax, highlighting the variation in terminology for the same land use category.	23
Table 3.3 Variation in zoning terminology between Toronto and Markham for the same permitted use (“Detached Dwelling”), highlighting the inconsistency in abbreviation systems across municipalities.	25
Table 3.4 Cross reference of zoning terminology between Toronto and Markham, demonstrating how functionally similar land uses are represented through different zone labels and abbreviation systems.	26
Table 3.5 Competency Question Categories and Examples Used to Guide Ontology Development	29
Table 3.6 Redundancy reduction actions and observed effects	58
Table 4.1 Summary of Competency Questions Used in the Evaluation	62
Table 4.2 Retrieval metrics evaluating the performance of the Zonology query suite across precision, recall, F1 score, and execution reliability.	81
Table 4.3 Summary of DL reasoning outcomes.	85
Table 4.4 Manual retrieval vs ontology backed retrieval (qualitative)	87

Table 4.5 The propsoed Zonology compariosn with exsting urban and planning ontology models	87
---	----

LIST OF FIGURES

Figure 3.1 Zonology ontology model for zoning information that includes core classes, constraints, and relationships.	29
Figure 3.2 Class hierarchy of the zoning ontology in Protégé.	34
Figure 3.3 Two core anchors of the model: Zone (spatial unit) and ZoningDesignation (regulatory category).	35
Figure 3.4 Extended class hierarchy of the zoning ontology in Protégé.	36
Figure 3.5 Main zoning designation families and their sub-categories.	37
Figure 3.6 Authority class in Protégé, showing hierarchy and annotation comment defining its role in planning and development regulation.	38
Figure 3.7 Class hierarchy of LandUse categories in Protégé.	40
Figure 3.8 Municipality class in Protégé with member instances representing GTA cities and regions.	42
Figure 3.9 Zoning Abbreviations class hierchy in Protege	43
Figure 3.10 Example of a subject–predicate–object relationship for Toronto land use permissions	44
Figure 3.11 Object property hierarchy in Protégé, showing zoning-related properties such as hasZoningCharacteristic, hasZoningIndex, and definedBy. Quantities and Constraints	46
Figure 3.12 OntoGraf visualization of core classes in Protégé.	49
Figure 3.13 OntoGraf visualization highlighting the Zoning Abbreviation class and its relations in Protégé.	52
Figure 3.14 Data property hierarchy in Protégé, showing hasUnit and hasValue.	53

Figure 3.15 Ontology-based alignment of Toronto’s “CL” and Markham’s “COM” zoning abbreviations, showing semantic equivalence through shared land use classes and property links.	55
Figure 3.16 OntoGraf visualization of zoning ontology classes and relationships in Protégé.	56
Figure 4.3 . SPARQL query and results listing Toronto zoning abbreviations with their full names.	64
Figure 4.4 SPARQL query and result identifying the full meaning of the zoning abbreviation EMP-GE.	66
Figure 4.5 SPARQL query and result identifying the equivalent zoning class of R1.	67
Figure 4.1 Mapping zoning abbreviations between municipalities (example: CL in Toronto and COM in Markham).	80
Figure 4.2 class in Protégé, showing its equivalency to TR and subclass relation to R. ..	84
Figure 4.6 Comparison of traditional versus ontology-based methods for identifying zoning equivalencies (example: RM1 in Markham vs. Toronto).	89

ABBREVIATIONS

AI – Artificial Intelligence

API – Application Programming Interface

ABox – Assertional Box

AEC- Architecture, Engineering and Construction

BIM- Buidling Inforamtion Modeling

CI – Continuous Integration

CSV – Comma-Separated Values

CQ – Competency Question

DAG – Directed Acyclic Graph

DL – Description Logic

ETL – Extract, Transform, Load

FAIR – Findable, Accessible, Interoperable, Reusable

FOL – First-Order Logic

F1 – Harmonic mean of precision and recall (F-measure)

FSI- Floor Space Index

GFA – Gross Floor Area

GIS – Geographic Information Systems

GML – Geography Markup Language

HTML – Hypertext Markup Language

HTTP – Hypertext Transfer Protocol

IC – Institutional Community (in zoning, but also used technically)

IDE – Integrated Development Environment

IRI – Internationalized Resource Identifier

ISO – International Organization for Standardization

JSON – JavaScript Object Notation

JSON-LD – JSON for Linked Data

LCBS- Land-Based Classification Standards

LOD – Linked Open Data

LOT – Linked Open Terms

NLP – Natural Language Processing

NS – Namespace

OCL – Object Constraint Language

ODP – Ontology Design Pattern

OGC – Open Geospatial Consortium

OLAP – Online Analytical Processing

OLTP – Online Transaction Processing

OMA – Ontology Modular Architecture

OMA-S – Ontology Modular Architecture – Structural

OMA-F – Ontology Modular Architecture – Functional

OMT – Ontology Mapping Tool

OPA – Official Plan Amendment (planning, but also used technically as reference in thesis)

ORS - Ontology Requirements Specification Document

OWL – Web Ontology Language

OWL 2 DL – Web Ontology Language 2 (Description Logic profile)

OWL 2 QL – Web Ontology Language 2 (Query Language profile)

OWL 2 RL – Web Ontology Language 2 (Rule Language profile)

PDF – Portable Document Format

R1, R2, R3... – Residential classes (also used in ontology case studies)

RA – Residential Apartment

RDF – Resource Description Framework

RDS – RDF Schema

REST – Representational State Transfer (API style)

RPC – Remote Procedure Call

SAAS – Software as a Service

SCA – Special Control Area (planning context but used in ontology testing)

SDI – Spatial Data Infrastructure

SHACL – Shapes Constraint Language

SKOS – Simple Knowledge Organization System

SML – Semantic Modeling Language

SPARQL – SPARQL Protocol and RDF Query Language

SQL – Structured Query Language

SP – Special Purpose (planning term, also ontology modeling)

SPA – Site Plan Approval (planning, used in ontology case studies)

SPIN – SPARQL Inferencing Notation

SPQ – SPARQL Query (shorthand in evaluation figures)

SWRL – Semantic Web Rule Language

TBox – Terminological Box

TTL – Turtle RDF serialization format

UML – Unified Modeling Language

URI – Uniform Resource Identifier

URL – Uniform Resource Locator

UX – User Experience

W3C – World Wide Web Consortium

WKT – Well-Known Text (for geometry representation)

WMS – Web Map Service

WFS – Web Feature Service

XML – eXtensible Markup Language

XSD – XML Schema Definition

ZBA- Zoning By-law Amendment

CHAPTER 1: INTRODUCTION

1.1. MOTIVATION OF RESEARCH

Ontario and particularly the Greater Toronto Area (GTA) is experiencing a severe housing supply crisis. Rapid population growth, constrained land availability, and escalating construction costs have created a mismatch between demand and new housing construction. In 2024, Ontario produced approximately 81,300 new housing units far below the provincial target of 125,000 units per year needed to meet the goal of 1.5 million new homes by 2031 (Ministry of Municipal Affairs and Housing (MAH), 2024). This persistent shortfall contributes to rising home prices and rental rates, threatening housing affordability and economic stability (Canada Mortgage and Housing Corporation (CMHC), 2024; Statistics Canada, 2024).

From a civil engineering standpoint, the housing crisis is not merely a design or financial problem it is deeply linked to how and when construction can begin. Before a project reaches the construction stage, developers must navigate an extensive and multi layered permitting and approvals process. This process involves coordination between engineers, planners, legal consultants, and municipal authorities, and includes several major applications such as the Official Plan Amendment (OPA), Zoning By-law Amendment (ZBA), and Site Plan Approval (SPA) (MAH, 2020; City of Toronto, 2023).

As the Greater Toronto Area (GTA) continues to experience rapid population growth and intensification, the need for consistent interpretation and enforcement of zoning by-laws has

become increasingly critical. Studies on Canadian land use planning have shown that fragmented municipal governance structures such as those found in the GTA create substantial inconsistencies in regulatory frameworks, leading to operational inefficiencies and uneven development outcomes (Hodge & Gordon, 2013; Slack & Bird, 2013). Because each municipality maintains its own zoning by-law, terminology, and classification system, planners and practitioners lack a unified structure for comparing or interpreting zoning designations across jurisdictions. This absence of a standardized semantic framework has been identified in planning literature as a major barrier to regional coordination, especially in metropolitan regions where growth pressures require integrated decision-making (Taylor, 2019). Moreover, research in urban informatics demonstrates that technology enabled planning systems and Geographical Information Systems (GIS) platforms require consistent, machine readable zoning information to function effectively; without it, interoperability issues arise, leading to data mismatch, manual reconciliation, and reduced analytical accuracy (Kuhn et al., 2014; Silvennoinen et al., 2023). Together, these findings show that semantic inconsistency in zoning is not merely an administrative inconvenience but a structural obstacle that undermines regional planning capacity and technological efficiency.

An OPA is required when a proposed development does not conform to a municipality's official plan policies, often redefining long-term land use designations, density, or growth boundaries (City of Toronto, 2023). A ZBA commonly referred to as a rezoning is needed when the existing zoning by-law does not permit the proposed use, height, or density. The SPA process, in turn, examines detailed engineering elements such as grading, servicing, parking layout, and landscaping (City of Toronto, 2023). Each of these stages can take months or even years, involving multiple circulation rounds, public consultations, and peer reviews. For Architecture, Engineering, and Construction (AEC) industry, these procedural delays translate into increased carrying costs,

scheduling uncertainty, and disruptions to servicing and infrastructure coordination (Pomeroy, 2020).

Among these approval stages, the ZBA is especially critical, as it establishes the legal and technical foundation for how land can be developed. To evaluate or propose a zoning change, one must first understand the existing zoning including permitted land uses, building types, floor space index (FSI), setbacks, height restrictions, and associated suffixes (e.g., “S” for special provisions or “H” for holding provisions) (MAH, 2017). Municipalities must also ensure the proposed amendment conforms with the Planning Act, provincial policy statements, and the local official plan (MAH, 2020).

However, within the GTA’s 25 municipalities, zoning information is highly fragmented. Each municipality employs its own set of abbreviations, definitions, and formatting conventions. For example, “R2-S” in the City of Markham refers to a low-density residential zone permitting semi-detached dwellings, subject to specific lot frontage, setback, and building form requirements, whereas “R2” in the City of Toronto denotes a residential zone with a different set of permitted uses, development standards, and built-form permissions that do not directly correspond to Markham’s designation (Hodge & Gordon, 2013). These inconsistencies generate confusion among developers, engineers, planners, and even digital systems tasked with interpreting by-laws. As a result, cross jurisdictional developments common in the GTA’s interconnected urban fabric face inefficiencies, legal disputes, and duplicative analyses (Slack & Bird, 2013).

The consequences of these discrepancies extend into the construction and permitting lifecycle. Misinterpretation of zoning information can delay land development, engineering design, servicing allocation, and building permit submissions, contributing to Ontario’s broader

housing supply delays (Pomeroy, 2020; Association of Municipalities of Ontario (AMO), 2024). Furthermore, as municipalities are transitioning toward data driven and Artificial Intelligence (AI) assisted planning systems, the absence of standardized, interoperable zoning data creates barriers to digital transformation (Silvennoinen et al., 2023). Since AEC increasingly relies on GIS, Building Information Modelling (BIM), and automated compliance tools, non standardized zoning semantics hinder integration between urban data systems, infrastructure models, and permitting workflows.

Therefore, this research is motivated by the need to bridge a critical regulatory and data interoperability gap that directly affects housing delivery and construction timelines. By harmonizing zoning terminology, it becomes possible to streamline data exchange between municipalities, improve communication among engineers and planners, and support faster, more transparent permitting processes. To build on these motivations, the next section narrows the focus to the underlying structural issue that drives many of these challenges: the fragmented and inconsistent nature of zoning data across the GTA.

1.2. PROBLEM STATEMENT

While the motivation of this research is rooted in national housing supply challenges and the procedural delays within its planning and approvals system, a more fundamental issue lies in the fragmented and inconsistent structure of zoning data across the GTA. The zoning system intended to serve as the regulatory foundation for land development is highly localized and semantically fragmented (Silvennoinen et al., 2023). Each municipality in the GTA maintains its own zoning by-law with distinct abbreviations, definitions, and rule hierarchies that have evolved independently. As a result, zoning codes that appear similar across municipalities often convey

entirely different meanings. For instance, “R2-S” in Markham and “R2” in Toronto may superficially seem equivalent yet differ in permitted uses, building typologies, and density regulations (Hodge & Gordon, 2013; Slack & Bird, 2013). This lack of standardization has produced a disjointed regulatory landscape that obstructs data interoperability, complicates cross jurisdictional collaboration, and introduces inefficiencies at every stage of the development process.

The consequences of this semantic fragmentation are both operational and systematic (Kuhn et al., 2014). Developers are required to manually interpret and translate zoning designations between municipalities, a process that is error prone and costly, particularly for projects spanning multiple jurisdictions. Legal practitioners struggle to ensure consistency in interpreting zoning provisions, leading to ambiguity and potential disputes (Pomeroy, 2020; Hodge & Gordon, 2013). Municipal planners, meanwhile, expend significant effort verifying whether proposed developments conform to local by-laws rather than engaging in proactive planning and policy evaluation (Pomeroy, 2020). In this context, even a basic inquiry such as “*What is permitted on this land parcel?*” often necessitates a manual review of local documents rather than an automated lookup in a shared database. These inefficiencies accumulate to delay housing delivery, inflate soft costs, and limit the scalability of data driven decision making in urban planning.

The issue is further magnified by Ontario’s growing reliance on geospatial technologies and data analytics for planning and policy evaluation. GIS, digital twins, and planning dashboards are increasingly used to map land use patterns, monitor intensification, and measure housing targets (City of Toronto, 2023; MAH, 2024). However, the inconsistent terminology and structure of zoning data among GTA municipalities undermine the analytical potential of these tools. Without semantic alignment, region wide analyses such as identifying under utilized parcels or

forecasting development potential remain unreliable or incomplete. Recent studies have highlighted similar challenges internationally, emphasizing that zoning and land-use regulations often exist as unstructured, text based data that resist machine interpretation (Silvennoinen et al., 2023). These findings reinforce that the absence of standardized, machine readable zoning information is a significant obstacle to the modernization of planning systems.

While ontology-based approaches have shown promise in supporting interoperability within urban planning systems, existing models are not designed to resolve the specific challenges associated with zoning fragmentation in the GTA. Several well established ontologies such as LBCS-OWL2 (Montenegro et al., 2012), the iCity Ontology Suite (Katsumi & Fox, 2019), OntoURBAN (Calderón & Ortiz, 2015), CityGML Utility Network ADE (Becker et al., 2013), and LandInfra (ISO/TC211, 2016) offer robust structures for land-use classification, city modeling, utility representation, and infrastructure semantics. However, their scope is either too broad or oriented toward domains such as transportation, utilities, or general urban morphology rather than the detailed regulatory semantics embedded in municipal zoning by laws. As a result, none of these ontologies provide mechanisms for representing zoning abbreviations, policy specific permissions, suffix conditions, or municipality specific development standards. This gap underscores the need for a dedicated zoning ontology that directly models the semantic inconsistencies observed across Toronto, Markham, and the wider GTA.

Emerging platforms such as LandLogic¹, the industry partner in this research, has demonstrated progress toward streamlining the planning and approval process by centralizing municipal data and digitizing workflows. These systems improve accessibility, transparency, and

¹ <https://www.landlogic.ai>

stakeholder coordination. However, they remain constrained by the underlying fragmentation of zoning data. While LandLogic enhances where and how information is accessed, it does not resolve what the information semantically represents. Zoning by laws continue to be a manual process in Landlogic ranging from scanned PDFs and spreadsheets to non standardized GIS layers each with inconsistent terminology and definitions. This forces the backend developers of LandLogic to rely on manual interpretation and cross referencing, thereby limiting automation and sustaining inefficiencies across municipal boundaries.

The persistence of these challenges points to the absence of an interoperable framework such as a shared semantic ontology for zoning data across the GTA. Ontologies, as defined by Gruber (1993), are explicit, formal specifications of conceptualizations that enable data from different systems to interoperate through common meaning. In the urban planning domain, ontology based approaches have proven effective in integrating heterogeneous data sources and automating rule reasoning (Guarino et al., 2009; McGuinness et al., 2004; Kuhn et al., 2014). For example, Silvennoinen et al. (2023) introduced the OntoZoning ontology to semantically represent land use regulations in Singapore, enabling machine interpretation of zoning rules and facilitating interoperability among planning systems. Similarly, Yazdanie and Orehounig (2021) and Trojahn et al. (2022) demonstrated how ontology based models improve data sharing, consistency, and reasoning capabilities within complex urban management contexts. Despite these advancements, all these are location specific and no comparable ontology exists for harmonizing zoning abbreviations, land use definitions, and development standards across the GTA's 25 municipalities.

This need for harmonization is especially urgent given the length and complexity of planning approval timelines in Ontario, which currently range from 2 to 4 years for many housing projects (Storeys, 2023). Semantic inconsistency in zoning regulations exacerbates these delays

by forcing reinterpretation of terminology, misalignment of datasets, and repeated back and forth between municipalities and development teams. A shared zoning ontology helps reduce this friction, allowing for standardized interpretation of land use permissions, cross municipal comparisons, and automated compliance checking. This not only expedites the permitting process but also enables more accurate feasibility assessments and reduces uncertainty in the early stages of development both essential steps toward meeting the province’s accelerated housing targets.

1.3. RESEARCH OBJECTIVES

Building upon the identified problem of fragmented and non standardized zoning data across municipalities in the GTA, this research aims to establish a structured framework that enables consistent interpretation and integration of zoning information between City Of Toronto and Markham jurisdictions. The focus is on creating a formalized, machine readable representation of zoning regulations that addresses the limitations of current manual, text based, and noninteroperable approaches. To achieve this overarching goal, the research pursued the following specific objectives:

- Identify and systematically classify discrepancies in zoning terminology, definitions, and permitted uses between Toronto and Markham, and assess the limitations of existing manual interpretation approaches.
- Develop an ontology based semantic framework, Zonology, using OWL (Web Ontology Language) to represent zoning regulations in a standardized and interoperable format.
- Validate and demonstrate the applicability of Zonology through expert evaluation, description logic–based reasoning tools, and integration with geospatial platforms, using the City of Toronto

and the City of Markham as a case study, while demonstrating the ontology’s transferability to other multi-jurisdictional planning regions to support zoning and planning queries.

1.4. RESEARCH CONTRIBUTIONS

This research contributes to the AEC and urban development by addressing a practical and recurring challenge in the land development and construction process interpreting and reconciling zoning information between municipalities. The study focuses specifically on the City of Toronto and the City of Markham, two jurisdictions with distinct zoning by laws that complicate engineering design, servicing coordination, and project approvals. The main contributions of this research are as follows:

- Establishment of a structured zoning framework that translates traditional text based by-laws from Toronto and Markham into a standardized, interpretable format. This framework supports engineers, planners, and developers in assessing land use permissions, building parameters, and design constraints more efficiently, reducing delays and misinterpretations during the planning and permitting stages.
- Integration of planning data with engineering workflows by demonstrating how zoning information can be connected to geospatial and infrastructure datasets. This contribution bridges the gap between regulatory interpretation and technical design, helping professionals evaluate development feasibility, servicing requirements, and compliance within a unified system.

Together, these contributions advance land development and construction practice by improving the accessibility and reliability of regulatory data supporting more coordinated, data informed decision making throughout the planning, design, and construction lifecycle.

1.5. STRUCTURE OF THESIS

The remainder of this thesis is organized as follows:

- **Chapter 2** presents a comprehensive literature review of the zoning landscape in the GTA and examines related ontologies that inform the foundation of this work.
- **Chapter 3** describes the proposed methodology to develop and validate the ontology model, Zonology, including , specification, conceptualization, implementation, and evaluation.
- **Chapter 4** discusses the results, key findings, and implications of the proposed Zonology ontology model, which semantically represents and aligns municipal zoning systems to address inconsistencies and support interoperable urban planning.
- **Chapter 5** concludes the thesis by summarizing key contributions, limitations, and opportunities for future research.

CHAPTER 2: LITRETURE REVIEW

To establish a foundation for developing the *Zonology* ontology model, a comprehensive review of zoning, land use, and urban ontology literature was conducted. The main research keywords included “zoning regulations,” “land use classification,” “urban ontology,” “semantic web in planning,” “data harmonization,” and “infrastructure ontologies”. These keywords guided the search across databases such as Scopus, ScienceDirect, and Google Scholar, as well as municipal and provincial planning sources. The goal of this review was to identify existing frameworks that inform the integration of zoning regulations into a standardized, machine readable model. The following subsection evaluates existing urban and infrastructure ontologies for potential reuse and adaptation within the context of municipal zoning in Ontario. Litreture review

2.1. EXISTING URBAN AND INFRASTRUCTURE ONTOLOGIES

The development of a comprehensive Zoning Ontology requires drawing from a range of preexisting ontologies that address different but interrelated aspects of urban planning, land use management, and infrastructure representation that summarized in Table 2.1. These ontologies offer foundational concepts, classification systems, and modelling approaches that can be adapted and extended to capture the complex and multi dimensional nature of zoning regulations in the GTA. At the same time, each ontology reflects its own scope, priorities, and inherent limitations. This section critically examines five key ontologies LBCS-OWL2 (Montenegro et al., 2012), iCity Ontology Suite (Katsumi & Fox, 2019), OntoURBAN (Calderon & Ortiz, 2015), CityGML Utility Network ADE (Becker et al., 2013), and LandInfra (ISO/TC211, 2016), presented in Table 2.1.

These are evaluated based on their focus, contributions, applicability, and constraints in the context of zoning regulation modelling.

The LBCS-OWL2 Ontology (Montenegro et al., 2012) represents a semantic web adaptation of the Land Based Classification Standards (LBCS) framework. It structures geospatial data with a strong emphasis on land use categorization, classifying land across five distinct dimensions: activity, function, structure, site development, and ownership. This multidimensional approach facilitates a nuanced understanding of land use that moves beyond single-category systems. Within the zoning ontology, elements of LBCS OWL2 have been partially reused to enhance land use data integration and enable a more refined classification of zoning areas. This reuse supports greater interoperability between zoning datasets and other geospatial resources, an essential consideration in the GTA where municipalities maintain diverse zoning by-laws and land use definitions. However, while LBCS-OWL2 provides a robust foundation for describing land use, it lacks a mechanism for representing the procedural and policy driven aspects of zoning. It does not model zoning enforcement, conditional use permissions, or variances critical elements in operational urban planning. Consequently, while LBCS OWL2 enriches the descriptive capacity of the Zoning Ontology, it must be supplemented with additional structures that explicitly address regulatory semantics.

Building on the LBCS framework, the iCity Ontology Suite (Katsumi & Fox, 2019) extends the classification system to support smart city applications, including transportation planning, environmental monitoring, and integrated land use management. iCity introduces detailed classifications of activities, land types, and their interrelationships, enabling advanced cross domain integration. Within the Zoning Ontology, these classifications have been repurposed to facilitate the cross referencing of zoning codes with land use policies. This functionality is

particularly relevant for aligning zoning data with regional transportation and mobility strategies, a growing priority in the GTA's urban growth management. Nonetheless, the iCity Ontology's primary orientation toward smart city operations means that its coverage of traditional zoning regulations is limited. While it excels in activity and infrastructure modelling, it does not address the full complexity of municipal zoning by-laws, such as conditional uses, density controls, and development restrictions. This necessitates selective integration, ensuring that iCity elements are used to complement rather than define the zoning framework.

The OntoURBAN ontology (Calderon & Ortiz, 2015) addresses urban planning from GIS perspective. It provides a structured framework for representing urban features, planning regulations, spatial temporal characteristics, and constraints. Of particular value to the Zoning Ontology is OntoURBAN's capacity to incorporate temporal dimensions, allowing for the representation of historical changes in zoning or land use over time. This is essential for modelling zoning amendments, policy reforms, and development approvals that evolve across planning cycles. By partially reusing OntoURBAN components, the Zoning Ontology can integrate not only current zoning restrictions but also the historical trajectories that inform them. Nevertheless, OntoURBAN's conceptual design is primarily oriented toward general urban feature modelling, and its zoning representations tend to be broad. It lacks the granularity needed to capture the detailed terminologies, abbreviations, and conditional clauses that define zoning in the GTA's municipal by-laws. As such, while OntoURBAN strengthens temporal and spatial contextualization, it requires adaptation to handle the precision and policy specific language of zoning regulation.

The CityGML Utility Network ADE (Becker et al., 2013) represents another distinct but relevant domain urban utility infrastructure. Designed as an extension to the CityGML standard,

it enables detailed modelling of networks such as water supply, wastewater, electricity, and telecommunications. While utility networks might seem peripheral to zoning at first glance, they are in fact integral to the feasibility and enforceability of zoning decisions. For instance, zoning changes that increase allowable density or change permitted uses often require corresponding adjustments to utility capacity. In developing the Zoning Ontology, concepts from the CityGML Utility Network ADE have informed the linkage between zoning regulations and infrastructure availability, ensuring that planning constraints reflect not only legal boundaries but also service capacity. However, CityGML's infrastructure centric focus means that it offers minimal direct representation of zoning rules, enforcement conditions, or permissible land use categories. Its integration into the Zoning Ontology is therefore highly targeted, enhancing infrastructure related attributes without attempting to define zoning itself.

Finally, the LandInfra standard (ISO/TC211, 2016) provides a standardized representation of land parcels, civil engineering works, and related infrastructure. Its relevance to the Zoning Ontology lies in its precise handling of *parcel based data*, enabling the accurate modelling of zoning boundaries, lot dimensions, and subdivision patterns. In the GTA, where zoning boundaries often coincide with parcel lines, LandInfra offers a valuable structural framework for geospatial alignment between cadastral records and zoning maps. This integration facilitates higher-precision spatial analysis and improves interoperability between zoning data and land registry systems. Nevertheless, LandInfra's scope remains primarily focused on the physical and cadastral aspects of land representation, offering only limited provisions for capturing zoning policies, conditional uses, or regulatory enforcement mechanisms. Similar to the other ontologies reviewed, LandInfra can support foundational data structures within the Zoning Ontology, but it does not contain the

regulatory concepts such as zoning permissions, land use conditions, or policy constraints needed to model zoning by laws on its own.

Table 2.1 provides a comparative review of five major urban and infrastructure ontologies and evaluates their relevance for potential reuse in the proposed zoning ontology. LBCS OWL2 offers a comprehensive system for classifying land use across five dimensions, including activity, function, structure, site development, and ownership. While it enhances the organization of land use information, it does not include the detailed regulatory elements needed to represent zoning permissions or enforcement rules. The iCity Ontology Suite builds on the LBCS structure and supports smart city and transportation applications, offering improved classification for residential and commercial activities. However, because its primary focus is on smart city operations, it does not fully address the specific terminology and regulatory detail found in municipal zoning by laws. OntoURBAN provides a useful framework for modelling urban features, spatial constraints, and planning regulations, and it can represent zoning changes over time, but its emphasis on general urban features limits its ability to capture the detailed semantics of land use regulations. CityGML Utility Network ADE contributes important capabilities for relating zoning information to utility networks and infrastructure planning, yet its scope is centered on utility modelling rather than zoning rules or land use restrictions. LandInfra supplies standardized representations for land parcels, buildings, and civil infrastructure and supports the depiction of zoning boundaries and parcel related data. Nonetheless, it provides only limited support for the types of zoning regulations, permissions, and policy conditions required for modelling municipal by laws. Overall, the comparison shows that while these ontologies provide valuable components for different aspects of urban information modelling, none of them offer the semantic and regulatory structures

necessary to address the inconsistencies and detailed requirements of zoning data across municipalities such as Toronto and Markham.

Table 2.1 Comparative Review of Existing Urban and Infrastructure Ontologies and Their Reuse in the Proposed Zoning Ontology

Ontology	Focus	Key Contributions	Integration within Zonology
LBCS-OWL2 Ontology (Montenegro et al., 2012)	Geospatial data with land use information, focusing on activity, function, structure, site development, and ownership	Comprehensive land use categorization across five dimensions which are activity, function, structure, stie development character and ownership.	Partially reused to enhance land use data integration and classify zoning areas in more detail
iCity Ontology Suite (Katsumi & Fox, 2019)	Smart city applications, including transportation planning and land use	Extended version of LBCSv2, detailed classifications of activities, land types, and their relationships	Reused for its activity classification such as residential, commercial activities, which supported cross referencing zoning and land use policies
OntoURBAN (Calderon & Ortiz, 2015)	Urban planning and GIS	Framework for representing urban features, planning regulations, spatial temporal features, and constraints	Partially reused to improve representation of zoning restrictions and historical changes over time

CityGML Utility Network ADE (Becker et al., 2013)	Urban utility networks within the CityGML standard	Modelling of urban utility networks (e.g., water, electricity) and linking infrastructure with city planning	Provided insights for linking zoning regulations with utility data, enhancing infrastructure related information in the ontology
LandInfra (ISO/TC211, 2016)	Standard for land and civil engineering infrastructure	Standardised representation of land parcels, buildings, and infrastructure	Incorporated to strengthen the representation of zoning boundaries and parcel related zoning data

2.2. INTEGRATION OF EXISTING ONTOLOGIES INTO THE ZONING ONTOLOGY FRAMEWORK

The integration of the five ontologies into the Zoning Ontology, presented in Table 2.1, has been driven by a selective adoption strategy. Each ontology offers distinct strengths whether in multidimensional land use classification (LBCS-OWL2), cross domain activity modelling (iCity), temporal and spatial contextualization (OntoURBAN), infrastructure linkage (CityGML Utility Network ADE), or parcel level geospatial precision (LandInfra). At the same time, none of these ontologies, in isolation, fully captures the complexity of zoning as practiced in the GTA, which requires not only spatial and descriptive accuracy but also semantic representation of policy, conditional permissions, and enforcement processes. Furthermore, these ontologies are not ISO

standards for several reasons. A summary of gap analysis is presented in Table 2.2 and described as follows:

- First, their development origins differ many, such as LBCS-OWL2, iCity, and OntoURBAN, were created by academic researchers, institutions, or specialized projects to address specific needs in land use, smart cities, or urban planning, rather than to serve as globally recognized frameworks.
- Second, their purpose and scope are narrower than ISO standards, which aim for universal applicability; in contrast, these ontologies are often region specific or tailored to niche applications.
- Third, the ISO standardization process is a lengthy and highly collaborative procedure involving international committees and consensus building, which most of these ontologies have not undergone, as they were designed for immediate application or research contexts.
- Finally, in fields like GIS and infrastructure, existing ISO standards such as LandInfra and CityGML already provide formalized structures, whereas ontologies in more specialized domains have not reached the same level of international standardization effort.

Taken together, these factors underscore why these ontologies, while highly valuable for their intended purposes, remain domain specific tools rather than globally adopted standards. A targeted gap analysis (Table 2.2) further shows that while each ontology contributes useful components, none provides the regulatory structures required for zoning in Toronto and Markham. LBCS OWL2 lacks zoning permissions, iCity does not model cross municipal terminology, OntoURBAN offers limited zoning vocabulary, CityGML Utility Network ADE focuses on utilities rather than land-use rules and LandInfra represents parcels but not zoning conditions. The Zoning Ontology

addresses these gaps by adding zoning-specific classes, aligning terminology between municipalities, and incorporating regulation aware logic that links land use, parcels, and permissions in a unified semantic model. The Zoning Ontology, therefore, leverages their strengths, addresses their limitations, and situates them within a framework specifically designed to meet the GTA's needs for zoning harmonization, regulatory clarity, and geospatial interoperability.

Zoning is a regulatory tool used by municipalities to control land use, intensity, and built form. In the GTA, each municipality maintains its own zoning by-law, resulting in fragmented terminology, inconsistent use classes, and varying suffix systems. Although zoning is intended to provide clarity and predictability, its implementation across jurisdictions has become increasingly difficult to interpret, harmonize, and digitize due to semantic inconsistencies and lack of standard structure.

To address zoning related data fragmentation and classification challenges, several ontology based models have been developed in the domains of urban planning, land use management, and smart city infrastructure. These include the LBCS-OWL2 ontology, which adapts the Land Based Classification Standards into a semantic web format (Montenegro et al., 2012); the iCity Ontology Suite, designed to support integrated smart city applications through activity based land use modelling (Katsumi & Fox, 2019); OntoURBAN, which incorporates spatial temporal planning constraints within a GIS oriented framework (Calderon & Ortiz, 2015); the CityGML Utility Network ADE, which models multi scale urban utility infrastructure (Becker et al., 2013); and the LandInfra standard, which formalizes parcel level land and civil infrastructure data (ISO/TC211, 2016). While each contributes to useful conceptual structures, none provides a comprehensive semantic representation of zoning regulations, conditional permissions, suffix rules, or

enforcement processes as required in municipal zoning systems. Existing zoning methods and ontologies remain limited because:

1. They do not model zoning by-law text, conditions, or suffix rules.
2. They lack interoperability across municipalities with different zoning systems.
3. They cannot support automated reasoning over zoning permissions.
4. They do not integrate zoning with land parcel, land use, and infrastructure data in one shared model.

These gaps justify the need to develop a unified, municipal level zoning ontology Zontology capable of harmonizing zoning semantics across the GTA.

Table 2.2 Gap Analysis of Existing Urban/Planning Ontologies

Ontology	Identified Gaps / Limitations
LBCS-OWL2 (Montenegro et al., 2012)	No modeling of zoning regulations, permissions, or enforcement rules; focuses only on land use categorization
iCity Ontology Suite (Katsumi & Fox, 2019)	Activity focused; lacks zoning by-law terminology and designation equivalency across municipalities
OntoURBAN (Calderon & Ortiz, 2015)	Limited support for zoning specific terms; no cross-municipal terminology alignment
CityGML Utility Network ADE (Becker et al., 2013)	Focuses on utility infrastructure; does not address zoning or development permissions
LandInfra (ISO/TC211, 2016)	Focused on parcels and infrastructure, not zoning permissions or use constraints
Overall Gap in Literature	No ontology standardizes zoning terminology or regulatory permissions across municipalities

CHAPTER 3: METHODOLOGY

This chapter presents the methodology developed for the design, implementation, and evaluation of the zoning ontology, called Zonology specifically for GTA. The primary aim of Zonology is to create a shared and reusable knowledge model that captures zoning terminology, development standards, and cross municipal alignments in a consistent and machine interpretable format as presented in Table 3.1.

Table 3.1 summarizes the four stage methodology used to develop the Zonology, beginning with specification of requirements, followed by conceptual modeling, technical implementation in OWL 2, and iterative evaluation. Each stage includes defined tasks, from domain knowledge acquisition and class/property definition to reasoning based validation and stakeholder review. This structured workflow ensures that the ontology is both technically sound and aligned with real world zoning practices in the GTA.

Table 3.1 The proposed framework to develop Zonology

Phases	Primary Purpose	Key Activities
1. Ontology Specification	Define scope, goals, and requirements	• Acquire domain knowledge (zoning by-laws, OPs, policy docs) • Create ORSD (scope, stakeholders, requirements) • Formulate competency questions (CQs) • Identify functional & non-functional requirements
2. Ontology Conceptualization	Structure knowledge model	• Use hybrid top down & bottom up modeling • Identify core zoning terms & class hierarchy • Define object & data properties • Align terminology with real-world planning workflows
3. Ontology Refinement & Alignment	Refine and standardize conceptual model	• Standardize terms and fix naming inconsistencies. • Reuse elements from existing ontologies and refine structure. • Prepare for OWL-DL compliance and finalize properties.
4. Ontology Implementation	Build ontology in OWL	• Implement in OWL 2 with Protégé • Encode concepts as classes & properties • Add constraints, axioms, & reasoning support • Enable cross-municipal linkage between Toronto & Markham
5. Ontology Evaluation	Validate quality, usability, and reasoning	• Perform DL reasoning checks (consistency, inference) • Test against competency questions (CQ validation) • Apply to real-world case studies/scenarios • Gather feedback from planners & stakeholders

3.1. ONTOLOGY SPECIFICATION

Before delving in Zonology development, we conducted a cross municipal comparison to examine how different municipalities define, label, and abbreviate zoning categories that permit the same land use. The purpose of this step was to investigate the extent of semantic inconsistency across zoning by-laws in the GTA and to establish an evidence based foundation for creating unified ontology classes and equivalency mappings. In this thesis, the term *ontology* is not used in the abstract philosophical sense common in disciplines such as

sociology, psychology, or environmental studies, where ontology refers to foundational assumptions about what kinds of social or natural entities exist and how they relate to one another. Instead, the term is used in its technical and computational sense, as defined within engineering, planning, geographic information systems (GIS), and smart-city research. In this context, an ontology refers to a formal, explicit, and machine-readable specification of concepts, relationships, and constraints within a domain, designed to support consistent interpretation, reasoning, and interoperability across data sources and systems. Four municipalities were selected for this comparison: Toronto, Markham, Aurora, and Ajax as presented in Table 3.2, Table 3.3, Table 3.4. These were chosen because they represent a mix of large urban centres and mid sized suburban municipalities, each operating under its own zoning framework and terminology system. They also differ in the age, structure, and formatting of their zoning by-laws, which makes them suitable for identifying variation in how the same type of land use is regulated and labeled.

Table 3.2 Comparison of zoning abbreviations for the permitted use “Detached Dwelling” in Aurora and Ajax, highlighting the variation in terminology for the same land use category.

Permitted Use	Aurora Abbreviation	Permitted Use	Ajax Abbreviation
Dwelling, Detached	ER (Estate Residential), Detached First Density Residential (R1), Detached Second Density Residential (R2), Detached Third Density Residential (R3), Detached Fourth Density Residential (R4), Detached Fifth Density Residential (R5), Special Mixed Density Residential (R7), PD1 Promenade Downtown, PD2 Promenade Downtown Special Mixed Density Residential, PDS1 Promenade Downtown Shoulder Central Commercial, PDS2 Promenade Downtown Shoulder Service	Detached Dwelling	Country Residential Zone (CR), Residential One – ‘A’ Zone R1-A, Residential One – ‘B’ Zone R1-B, Residential One – ‘C’ Zone R1-C, Residential One – ‘D’ Zone R1-D, Residential One – ‘E’ Zone R1-E, Residential One – ‘F’ Zone R1-F,

	Commercial, PDS3 Promenade Downtown Shoulder Institutional, PDS4 Promenade Downtown Shoulder Special Mixed Density Residential, RU Rural		Residential Two – ‘A’ Zone R2-A, Village Core Mixed Use Four Zone VC4
--	---	--	--

To maintain a controlled comparison, a single permitted land use detached residential dwelling was selected as the reference category. This land use is present in every municipality, yet it is expressed through different zone names, abbreviations, and regulatory structures. For each city, all zoning labels that legally permit detached dwellings were extracted from official zoning by-law documents, including both the full zone name and the shorthand abbreviation used in planning reports, GIS layers, and zoning schedules. This became immediately clear that while the functional meaning was the same, the terminology used to express it was entirely different between municipalities. For example, Toronto uses a series of short codes such as R, RS, RT, RM, and RD, which distinguish between general residential, semidetached, townhouse, and detached dwelling permissions. Markham, on the other hand, uses a more descriptive but structurally different naming system such as RES-LR1, RES-LR2, RES-ES, and RES-ENLR, incorporating both the housing form and neighbourhood classification. Aurora adopts a density-tiered model (R1, R2, R3, etc.), while Ajax uses an alphanumeric system (R1-A through R1-F, CR, VC4) that reflects both dwelling type and neighbourhood designation. These findings confirmed that the problem is not only linguistic but structural. Each municipality embeds different regulatory assumptions into its zoning labels. This means there is no one to one match between codes, even when they govern the same land use. As a result, zoning data cannot be merged, compared, or automated across municipalities without an intermediary semantic model.

Table 3.3 Variation in zoning terminology between Toronto and Markham for the same permitted use (“Detached Dwelling”), highlighting the inconsistency in abbreviation systems across municipalities.

Permitted Use	Toronto Abbreviations	Permitted Use	Markham Abbreviations
Detached / Dwelling Unit	R (Residential), RAC (Residential Apartment Commercial), RT (Residential Townhouse Zone), RS (Residential Semi-Detached Zone), RM (Residential Multiple Zone)	Detached dwelling	RES-PE (Residential – Public Education), RES-LR4 (Residential Low Rise Four), RES-LR2 (Residential Low Rise Two), RES-LR1 (Residential Low Rise One), RES-ENLR (Residential Established Neighbourhood Low Rise), RES-ES (Residential Estate), RES-ENLR (Residential Established Neighbourhood Low Rise), GWY1 (Greenway One), GWY2 (Greenway Two), GWY3 (Greenway Three), HAM-RES (Hamlet Residential), FD (Future Development), MU-HMS(T) (Mixed Use Heritage Main Street – Thornhill Village), MU-HMS(U) (Mixed Use Heritage Main Street – Unionville Village), MU-HMS(M) (Mixed Use Heritage Main Street – Markham Village)

This comparative exercise served for two key purposes. First, it provided empirical validation of the research problem by demonstrating that zoning terminology is not interoperable across municipal boundaries. Second, it supplied the raw vocabulary needed for ontology development. The extracted terms were later converted into ontology classes, individuals, and object properties, allowing the ontology to function as a translation layer between heterogeneous zoning systems. In other words, this step transformed unstructured municipal zoning text into structured conceptual input for the ontology. It also informed the design of semantic relationships such as *hasAbbreviation*, *hasPermittedUse*, and *isEquivalentTo*, which are later used for reasoning, alignment, and query testing.

In order to achieve this, we employed a transparent development process rooted in the FAIR principles, which emphasize findability, accessibility, interoperability, and reusability. Adhering to these principles ensures that the ontology is not only technically sound but also reliable and reusable across multiple stakeholders, including planners, policymakers, and software developers.

3.1.1. Acquiring domain knowledge

The first stage of the Zonology framework involved defining the scope of the domain and identifying the requirements that the ontology needed to fulfill. This process began with extensive knowledge acquisition from primary planning and legal sources, including municipal zoning by-laws, official plans, secondary plan policies, and associated schedules. Particular attention was given to cross-municipal inconsistencies in zoning terminology, since these discrepancies are the main source of confusion for developers, planners, and policymakers across the GTA.

Table 3.4 Cross reference of zoning terminology between Toronto and Markham, demonstrating how functionally similar land uses are represented through different zone labels and abbreviation systems.

Toronto – Terminology	Markham – Terminology
R (Residential)	RES-ENLR (Residential Established Neighbourhood Low Rise), RES-PE (Residential – Public Education)
RT (Residential Townhouse Zone)	RES-LR2 (Residential Low Rise Two), RES-LR3 (Residential Low Rise Three), RES-LR4 (Residential Low Rise Four)
RS (Residential Semi-Detached Zone)	RES-LR1 (Residential Low Rise One)
Residential Multiple Zone (RM)	RES-LR4 (Residential Low Rise Four)
RD (Residential Detached)	RES-ES (Residential Estate Zone), RES-ENLR (Residential Established Neighborhood Low Rise)
	MU-LR (Mixed Use Low Rise), MU-MR (Mixed Use Mid Rise), MU-HR1 (Mixed Use High Rise One), MU-HR (IA) (Mixed Use – High Rise)

	(Intensification Area)), MU-HMS(T) (Mixed Use Heritage Main Street – Thornhill Village), MU-HMS(U) (Mixed Use Heritage Main Street – Unionville Village), MU-HMS(M) (Mixed Use Heritage Main Street – Markham Village)
CR (Commercial Residential)	—
CL (Commercial Local Zone)	Commercial (COM)
CRE (Commercial Residential Employment Zone)	—
E 1.0 (Employment)	EMP-GE (Employment – General Employment)

To illustrate the extent of these cross-municipal discrepancies, Table 3.4 provides a direct cross-reference of zoning terminology between Toronto and Markham, showing how land-use categories that serve the same planning function are labeled very differently in each municipality. While Toronto generally uses simpler, broad zone codes such as R, RT, or RM, Markham employs a more granular system with multiple sub categories (e.g., RES-LR1, RES-LR2, RES-LR4, RES-ENLR). This comparison highlights how two jurisdictions can regulate similar residential and mixed-use forms through distinct abbreviation conventions and classification structures, underscoring the semantic fragmentation that complicates zoning interpretation across the GTA.

3.1.2. Creation of ORSD

The insights from this stage were formalized in an Ontology Requirements Specification Document (ORSO), which captured both functional and non functional requirements. The functional requirements were expressed as competency questions that guided the ontology design, ensuring that each part of the model corresponded to a real world planning task. For instance, questions such as “*What land uses are permitted under a given zoning type?*” or “*Which designation in Mississauga corresponds to Toronto’s CR?*” directly informed the construction of classes and properties. Non functional requirements included

clarity, consistency, usability, extensibility, and coverage. Each of these qualities was linked to specific evaluation criteria for example, clarity was tested through practitioner workshops, while consistency was assessed using automated reasoning tools.

Figure 3.1 presents the ontological model of the Zonology, integrating core classes, constraints, and semantic relationships. The model shows how *ZoningDesignation* is linked to planning metadata (e.g., Official Plan, Zoning By-law), geospatial attributes (e.g., Height, Frontage), and permitted uses (e.g., BuildingType, LandUse). Key relationships such as *hasAbbreviation*, *regulatedBy*, and *permittedUse* demonstrate how zoning codes connect to governing authorities, municipalities, and land use permissions. Datatype properties (e.g., string, numeric, dateTime) support the representation of quantitative and textual zoning attributes in a machine readable format. Overall, the diagram illustrates a unified semantic framework that aligns regulatory policy, spatial constraints, and zoning terminology into a single interoperable structure as shown in Figure 3.1.




Figure 3.1 Zonology ontology model for zoning information that includes core classes, constraints, and relationships.

3.1.3. Formulate Competency Questions (CQs)

Competency questions (CQs) formed the backbone of the functional requirements for the Zonology model as it is shown in table 3.5. Their development was central to ensuring that the ontology was not only a conceptual exercise but also a practical tool capable of addressing the recurring questions that arise in municipal planning, zoning review, and development approvals. The process of identifying these CQs was iterative and drew upon both document analysis and consultations with planning professionals. In particular, the ORSD outlined three major categories of CQs: general definitional queries, cross jurisdictional comparisons, and constraint based regulatory queries. These categories directly informed the structuring of classes, properties, and axioms in the ontology.

Table 3.5 Competency Question Categories and Examples Used to Guide Ontology Development

Competency Question Field	Question Example
General definitional queries	What uses are permitted in the CR 3.0 zone in Toronto?
Cross jurisdictional comparisons	How does the definition of “Mixed Use” in Markham differ from Toronto’s equivalent zone?
Constraint based regulatory queries	What is the maximum allowable height for a building on this parcel, considering zoning and overlay exceptions?
Use specific or activity based queries	Which zones across the GTA allow Daycare as a permitted or conditional use?

3.1.3.1. General Defenition

The first set of competency questions addressed the definitional aspects of zoning classifications, ensuring that the ontology could adequately capture the essential characteristics of zoning designations. Practitioners often begin by asking broad, foundational questions such as “*What land uses are allowed in each zoning type?*” or “*What building types are permitted under a specific zoning type?*” These definitional questions required the ontology to represent both permitted and prohibited uses through well defined classes and object properties.

Other definitional queries inquired the temporal and administrative dimensions of zoning. Questions such as “*What time was this zone planned or designated?*” and “*Who planned or designated this zoning area?*”, highlighted the importance of embedding governance and temporal attributes within the model. This was addressed by associating zoning designations with municipal by-laws or planning documents, and by linking them to metadata on adoption dates or responsible authorities.

3.1.3.2. Cross-Jurisdictional Compariosns

The second major set of CQs dealt with the problem of terminological fragmentation across municipalities, a central issue motivating Zonology model. One of the most pressing needs in planning practice is to answer questions such as “*How are zoning terms from different cities or towns related to each other?*” or “*What is the equivalent zoning type in another municipality?*” These queries demanded a framework for semantic alignment, whereby the ontology could represent equivalence, similarity, or divergence between municipal zoning designations.

Concrete examples from the competency questions demonstrate the scale of the problem. For instance, “*What zoning category in Markham corresponds most closely to Toronto’s CR (Commercial Residential) zone?*” captures the cross walk challenge between municipalities. Similarly, the query “*Are there zoning abbreviations that share the same name but represent different uses across municipalities?*” exposes the potential for confusion when identical terms (e.g., “R1” or “C2”) mean different things in different jurisdictions. To address these, the ontology introduced object properties such as *isEquivalentToDesignation* and *hasAbbreviation*, allowing multiple terms to be mapped onto a common semantic framework.

Cross jurisdictional queries also extended to uses and permitted activities. The CQ “*Can a single permitted use, such as daycare, be found under different zoning categories across municipalities?*” required Zonology to model not only zoning categories but also their associated land uses, enabling queries that transcend municipal boundaries. Another CQ, “*What is the most common abbreviation used across all 25 GTA municipalities for low-density residential zones?*” reinforced the need for ontology driven aggregation, which could reveal patterns across multiple sources.

Furthermore, the questions explored the differences in definitions across municipalities. For example, “*How do definitions of a particular zoning type (such as R1 or C2) differ between municipalities?*” and “*Which municipalities define mixed-use zones, and how are these definitions similar or different?*”, these queries are required a design that could store not only equivalence but also contrasting definitional attributes, thus supporting nuanced comparative analysis. By formalizing these relationships, the ontology provides a systematic means of navigating the otherwise fragmented landscape of zoning terminology.

3.1.3.3. Constraint-Based Queries

The third category of CQs focused on zoning constraints and regulatory standards, which are the core of development control. These queries asked not just about abstract definitions but about concrete, enforceable limits such as height, density, lot coverage, and setbacks. For example, *“What is the maximum building height allowed in a zone?”* and *“What is the minimum frontage of a building in a zone?”* directly informed the creation of properties like `hasMaxHeight` and `hasMinimumFrontage`.

Other constraint based queries captured land division and environmental regulations. Questions such as *“What is the minimum area needed for severance?”* and *“Is all or a portion of land environmentally protected?”* required the ontology to model relationships between zoning designations and environmental overlays or severance policies. Similarly, the CQ *“Are there any overlay zones or special provisions affecting this parcel?”* highlighted the importance of representing not only base zoning categories but also layered or exception based provisions.

Constraint based CQs also extended into land use and accessory buildings. For instance, *“Are accessory buildings allowed, and if so, what are their constraints?”* required explicit representation of accessory structures and their permissible attributes. *“What are the parking space requirements for residential or commercial use in this zone?”* emphasized the need to model detailed, context sensitive constraints tied to land use types.

In addition, some CQs interrogated the regulatory sources of constraints. Questions such as *“Which regulation defines a restriction of height?”* underscored the necessity of linking constraints back to authoritative documents. This ensured that the ontology was not only a

repository of rules but also a traceable system that could be aligned with legal and policy documents.

Taken together, the three sets of competency questions provided a systematic blueprint for ontology construction. Definitional queries drove the creation of foundational classes (ZoningDesignation, LandUse, BuildingType), temporal attributes (dateDesignated), and planning relationships (definedByBylaw, linkedToOfficialPlan). Cross jurisdictional queries necessitated semantic alignment through equivalence properties (isEquivalentToDesignation, hasAbbreviation), while constraint based queries introduced quantitative data properties (hasMaxHeight, hasFrontage, hasLotCoverage) and links to regulatory sources.

In this way, the ontology remained grounded in practice: every class and property could be traced back to at least one competency question. This traceability strengthened the functional requirements by ensuring that the ontology would not only be internally consistent but also directly usable for information retrieval, compliance checking, and cross municipality harmonization. By embedding answers to CQs in its structure, the Zonology addresses the persistent problem of zoning abbreviation inconsistencies while also enabling advanced semantic applications such as automated reasoning, SPARQL query answering, and integration with geospatial data systems.

3.2. ONTOLOGY CONCEPTUALIZATION

The conceptualization process followed a hybrid approach. A top down pass abstracted stable planning concepts from municipal by-laws and official plans (designations, uses, standards, exceptions). A bottom up pass examined municipal datasets and GIS layers to surface how those concepts appear in practice (abbreviations, overlay areas, measurement units, and

the way parcels are typed in data). The outcome is a lightweight but expressive core that (a) keeps the regulatory rulebook distinct from the spatial footprint, and (b) lets numeric standards travel as first class objects with units. Figure 3.2 presents the model driven narrative of the classes, subclasses, and relationships in the attached ontology, along with design patterns and concrete examples from the file.




Figure 3.2 Class hierarchy of the zoning ontology in Protégé.

3.2.1. Identifying Core Zoning Terms and Class Hierarchy

- The ontology is built around two complementary anchors as shown in Figure 3.3. The first is the Zone, which represents the physical spatial unit on the ground. A zone is a spatial unit or GIS feature. It changes location or geometry through mapping edits, and it changes designation through planning decisions. The second is the ZoningDesignation, which defines the regulatory category applied to land. It is the abstract rule bundle (e.g., “Residential Low Rise One”, “Commercial Residential”) that municipalities apply to land. Together, they link geographic features to the rule bundles used in municipal planning.




Figure 3.3 Two core anchors of the model: Zone (spatial unit) and Zoning Designation (regulatory category).

Keeping Zone and ZoningDesignation separate makes the model reusable over time and across municipalities. One designation can apply to many zones. A zone can shift from one designation to another as by-laws change. The ontology even reserves an explicit link `hasZoneDesignation` with domain `Zone` and range `ZoningDesignation`, underscoring the intended pattern: zone instances point to designation classes. In the Zonology, a few example records (e.g., `Zone_MK_1`, `MK_Zone_2`) are created as `NamedIndividuals` but typed directly as designation classes (e.g., `Residential+Low+Rise+One`). That works for experiments, but the recommended pattern is to type them as `Zone` and connect them via `hasZoneDesignation`. Doing so keeps the ontology in OWL DL and makes reasoning and queries more predictable.

This core is rounded out by four other pillars as shown in Figure 3.4:

- Municipality (a subclass of `Authority`) identifies who governs and designates.
- Document presents where rules are written (e.g., by-laws, official plans, special policy documents).
- `ZoningCharacteristic` is the semantic “content” of a designation (uses, building types, overlays).

- ZoningIndex is quantitative standards with units and values (height, frontage, FSI, parking, coverage).

Together, these pillars let you say who defines what rule bundle, where it applies, what it allows or restricts, and how much of something is permitted.




Figure 3.4 Extended class hierarchy of the zoning ontology in Protégé.

3.2.2. Define Object and Data Properties

Zoning Designation sits under a lightweight Zoning node and acts as the master category for regulatory labels. The attached model contains a wide set of subclasses ($\approx 60+$), including examples as shown in figure 3.5.




Figure 3.5 Main zoning designation families and their sub-categories.

Two of these designations also carry explicit logical restrictions in the file that the reasoner can use:

- Commercial+Residential declares a set of permitted uses via `hasPermittedLandUse` some ... (e.g., Art+Gallery, Artist+Studio, Bed+and+Breakfast, Community+Centre, Personal+Service+Shop, etc.).
- Residential+Low+Rise+One carries both permitted uses and a numeric max height restriction through `hasMaxHeight` `hasValue` `MaxHeightRES-LR1`.

Beyond those two, Zonology also includes numerous direct assertions connecting a designation to permitted uses using the same predicate (`hasPermittedLandUse`). These assertions currently appear as class to class triples and function like a knowledge base of permissions.

Figure 3.6 illustrates how the ontology captures the institutional dimension of zoning regulation by representing the bodies responsible for administering planning policies. Figure 3.6 highlights the class structure and includes an annotation comment describing the institutional role of these entities in regulating, enforcing, or approving development related decisions. This modelling approach ensures that every zoning designation or policy reference within the ontology can be traced back to its governing institution, supporting clear provenance, transparent decision pathways, and more reliable cross-municipal interpretation.




Figure 3.6 Authority class in Protégé, showing hierarchy and annotation comment defining its role in planning and development regulation.

ZoningCharacteristic groups the semantic content attached to a designation. It has four immediate subclasses in the file:

- LandUse: the largest branch ($\approx 150+$ subclasses), e.g., Apartment+dwelling, Automated+Banking+Machine, Banquet+hall, Community+garden, Daycare,

Financial+Institution, Massage+Therapy, Personal+Service+Shop, and many more.

A note from the data: you will see minor variants and duplicates (e.g., Art+Gallery and Art+gallery; percent encoded apostrophes like Carpenter+Shop). This reflects bottom up ingestion and suggests a next step for canonicalization (preferred labels, altLabels, and equivalence mappings).

- **BuildingType**: currently minimal (e.g., AncillaryBuilding). The schema reserves both `hasPermittedBuildingType` and `hasOptionalBuildingType`, giving you room to state whether a type is allowed vs optional/ancillary within a designation.
- **Environmentally Protected**: a useful hook for overlays, conservation areas, or constraint layers. In practice, you would connect a Zoning Designation (or a Zone) to an Environmentally Protected marker when an overlay applies. The model also offers `hasGeospatial Representation` (typo discussed later) to point to a geometry.
- **Amenities**: an umbrella notion for features that can be referenced as characteristics of a designation or area (e.g., parks, community facilities). It is present as a class and can be extended with more detail later.

Figure 3.7 displays the detailed class hierarchy of LandUse categories in Protégé, illustrating the wide range of land use types captured in the ontology. These include:

- Residential uses (e.g., Apartment Dwelling, Back to Back Townhouse Dwelling),
- Commercial uses (e.g., Business Office, Cold Storage, Commercial Storage Facility),
- Institutional uses (e.g., College, Child Care Centre), and cultural or recreational uses (e.g., Art Gallery, Club).

By explicitly representing such diverse land use categories, the ontology enables zoning designations to be semantically linked to specific activities and building functions, allowing for more precise queries and meaningful cross municipal comparisons.




Figure 3.7 Class hierarchy of LandUse categories in Protégé.

ZoningIndex is the abstraction for numeric standards. It has five important subclasses as follows:

- Height

- Frontage
- FloorSpaceIndex
- ParkingRatio
- LandCoverage

The ontology models numeric values as individuals of these classes carrying two datatype properties:

- hasValue (e.g., "6"^^xsd:string or a numeric datatype)
- hasUnit (e.g., "m")

For example, the individual MaxHeightRES-LR1 is a Height with hasUnit "m" and hasValue "6". The designation Residential+Low+Rise+One uses an OWL restriction: hasMaxHeight hasValue MaxHeightRES-LR1. This is a clean, composable pattern:

1. Model the metric as a first class thing with a unit,
2. Connect a designation (or zone) to that metric with a property that captures the role (max vs min vs required),
3. Keep the numeric literal and unit inside the metric individual so you can reuse it and change it without touching the TBox.

There is also a class named Toronto (as a subclass of a misspelled Municipality) as it shown in figure 3.8, which looks like an early scaffold. Prefer the NamedIndividual approach for real municipalities, and have the Municipality class define the type.



Figure 3.8 Municipality class in Protégé with member instances representing GTA cities and regions.

Document underpins traceability. It has subclasses:

- OfficalPlan (spelling in file), standing in for the Official Plan and related policy layers.
- ZoningRegulation, the by-law universe.
- GreenspaceRequirement, which in its current placement is modeled as a Document. If the intent is a requirement type rather than a document, consider moving it under ZoningCharacteristic or introducing a Requirement class and linking it back to Document via isDefinedIn.

Documents exist so that designations, abbreviations, and numeric standards can be referenced and audited. The property hasReference (range Document) captures this link.

3.2.3. Aligning terminology with real word planning workflows

The ontology includes a strong cross jurisdictional layer as it is shown in figure 3.9:

- ZoningAbbreviation is a class with ~60+ subclasses such as CR, EMP-BP, GWY1, RES-LR1, CF-PW, etc as shown in figure 3.9.
- The property definedBy connects an abbreviation to a Municipality (e.g., CR definedBy Markham).

To harmonize zoning terms, the ontology needs a clear link between each abbreviation and its full zoning designation. This is usually done with a specific property or with skos:altLabel/ if abbreviations are modeled as individuals. Right now, these links are created using class to class assertions, which pushes the model toward OWL Full. Adjusting this structure will make the ontology cleaner and easier for reasoning tools to use.




Figure 3.9 Zoning Abbreviations class hierarchy in Protege

3.2.4. Relationships (Object and Data properties) and Modeling patterns

This section is organized into several key sections, including Governance and Provenance, Semantics of Permissions and Structure, Quantities and Constraints, Data Properties, Spatial Ties, Concrete Examples, Quality Notes and Forward Refinements, and How This Answers Real Questions. Figure 3.10 illustrates how a basic subject–predicate–object relationship operates within the ontology. In this example, the zoning designation TorontoCR3_0 functions as the subject, the property hasPermittedLandUse serves as the predicate, and the land-use class OfficeUse is the object. The diagram shows how the subject connects to the object through the predicate, demonstrating the semantic structure used to represent permitted activities within Toronto’s zoning system.




Figure 3.10 Example of a subject predicate object relationship for Toronto land use permissions [generated by Napkin AI].

The following predicates were identified to define how zoning designations are linked, described, and referenced within the ontology. These properties establish clear connections to their sources, municipal context, and supporting documents, ensuring that each designation can be interpreted and queried consistently.

- `designatedBy` (domain `ZoningDesignation`, range `Authority`): who designates a category. Used widely to connect designations to Markham in the file.
- `designatedIn` (domain `ZoningDesignation`, range `OfficialPlan`): which policy document defines the category.
- `within` (domain `ZoningDesignation`, range `Municipality`): a spatial/administrative containment relation for quick scoping.
- `hasReference` (domain `ZoningAbbrevciation` or `ZoningDesignation`, range `Document`): where to look up the formal citation.

These relations support auditability (“show me the by law”), and they keep competency questions about “who/where/which document” answerable.

Semantics of Permissions and Structure as it is shown in figure 3.11

- `hasPermittedLandUse` (domain `ZoningDesignation`, range `LandUse`): the main permission hook; present both as OWL restrictions (e.g., for `Commercial+Residential`) and as class to class assertions across the file (e.g., `Com -> Hotel`, `Employment+-+General+Employment -> Hotel`).
- `hasPermittedBuildingType` and `hasOptionalBuildingType`: a clean distinction when a designation allows certain structural forms and optionally allows ancillary ones.
- `hasZoningCharacteristic`: a generic hook to attach any `ZoningCharacteristic` to a designation (overlays, amenities, etc.).
- `hasGeospatialRepresntation` (typo in name): connects a designation to a `SpatialEntity`. This is the bridge to geometries (e.g., GeoSPARQL `ogc:asWKT` or a feature id).




Figure 3.11 Object property hierarchy in Protégé, showing zoning-related properties such as `hasZoningCharacteristic`, `hasZoningIndex`, and `definedBy`. Quantities and Constraints

In the Zonology quantitative zoning constraints are modeled using either specific or generic predicates, following conditional approaches:

- `hasMaxHeight`, `hasMinFrontage`, `hasMaxFSI`, `hasMinParkingRatio`: role bearing predicates that tie a designation to a numeric object from the `ZoningIndex` branch.
- `hasZoningIndex`: a more generic, catch all link to a `ZoningIndex` instance when you do not want to proliferate role specific predicates.

If a numeric restriction corresponds to a clearly defined zoning attribute such as maximum height, minimum frontage, maximum floor space index (FSI), or minimum parking ratio then it should be represented using role specific predicates like `hasMaxHeight`, `hasMinFrontage`, `hasMaxFSI`, or `hasMinParkingRatio`. These predicates directly link each zoning designation to a corresponding instance in the `ZoningIndex` branch, supporting precise querying and transparent semantic interpretation. If, however, the constraint does not align with an existing property or would otherwise require introducing a rarely used or overly specific predicate, then it is preferable to use the generic `hasZoningIndex` predicate as a flexible fallback. This ensures

that the ontology remains semantically robust without sacrificing usability or extensibility. Each numeric object is expressed as a `NamedIndividual` (e.g., `MaxHeightRES-LR1`), and carries `hasUnit` and `hasValue`. This is powerful because (a) units live in the data, (b) values can be reused and versioned, and (c) you can encode different roles (max vs min) via the object properties without duplicating literals in multiple places.

The ontology uses data properties to store the numerical details of zoning rules. For example, `hasValue` records the actual number for a standard, and `hasUnit` records its measurement unit. Spatial information is handled through the `SpatialEntity` class, which can later be connected to GIS systems by adding `GeoSPARQL` so that each zone can carry a real geometry. In this setup, a zone would include its shape, its zoning designation, and any special overlays such as environmental protection areas.

The examples in the Zonology show how this works. The Residential Low Rise One designation lists its permitted uses through `hasPermittedLandUse`, and it records its height limit through `hasMaxHeight`, which points to an object that stores the height value of six metres. This makes it easy to answer questions like *“What is the maximum height allowed?”*. The Commercial Residential designation follows the same pattern by listing all the activities it permits, such as art galleries, personal service shops, and community centres. Abbreviations like CR or RES-LR1 are modeled as zoning abbreviation classes and are linked to the municipality that created them. The ontology also contains individuals for GTA municipalities, such as the City of Toronto and the City of Markham, which serve as the governance authorities for these designations.

3.3. ONTOLOGY REFINEMENT AND ALIGNMENT

The ontology's ability to move beyond a static classification system into a functioning legal knowledge base rests on the way its properties and modeling patterns have been designed. Properties form the semantic links that bind together zoning designations, uses, authorities, documents, and constraints, while modeling patterns provide systematic ways to handle recurring complexities such as provenance, exceptions, and cross jurisdictional alignment. Object properties are central in this respect, as they capture the core relationships of planning practice. For example, the property `hasPermittedLandUse` links a zoning designation directly to the activities that it allows, thereby transforming abstract categories into actionable permissions. With this property in place, one can ask not only "*What is allowed under Commercial Residential?*" but also "*Which designations across the system allow a specific use such as daycare or community centres?*" The property `designatedBy`, which connects each designation to its governing municipality, thereby grounding the ontology in the jurisdictional realities of planning law. Without such links, designations would float as abstract entities divorced from the authorities that define and enforce them. Other properties such as `designatedIn` and `hasReference` extend this grounding by anchoring designations to their authoritative sources in zoning by-laws, official plans, or special policy documents, ensuring that every rule can be traced back to a legal source. In addition, specialized properties like `hasMaxHeight`, `hasMinFrontage`, and `hasMaxFSI` encode constraints not as raw numbers but through connections to structured resources representing standards, which themselves carry values and units. This design makes numerical limits legally traceable and computationally reusable, while spatial properties such as `within` or `hasGeospatialRepresentation` allow integration with GIS data, enabling queries that combine regulatory and geographic reasoning.

Data properties play a complementary role, providing the precision required to express quantitative details. Instead of hard coding numbers directly into designations, the ontology uses `hasValue` and `hasUnit` attached to instances of standards such as Height or Floor Space Index. This ensures that a six metre height restriction is not just a literal number but a reusable and referenceable node. Such an approach allows the ontology to support both legal auditability and efficient data management, since changes or conversions only need to be updated once at the value node rather than across all instances where they are used.

Figure 3.12 shows the core classes of the ontology using Protégé's OntoGraf view. Each major concept such as Zoning, Zone, Permit, Authority, Document, SpatialEntity, Zoning Abbreviation, ZoningCharacteristic, and ZoningIndex appears as a direct subclass of `owl:Thing`. This visualization highlights the main building blocks of the ontology and how regulatory, spatial, governance, and quantitative elements are organized at the top level.




Figure 3.12 OntoGraf visualization of core classes in Protégé

3.3.1. Reuse Elements From Existing Ontologies And Refine Structure

Beyond individual properties, the ontology also uses three modeling patterns that help manage complexity and stay consistent with real planning documents. The first is the Value With Provenance pattern, which makes every numerical rule traceable. Instead of storing a number on its own, the value is stored as a small structured item that contains the number, the unit, and links to the by-law or document that defines it. This means the system can show not only that the height limit is six metres, but also exactly which by-law section sets that limit. It also makes it possible to track how rules change over time because older values can be stored with their own time periods.

The second pattern is the Exception Overlay, which reflects how zoning often includes special provisions such as heritage rules or environmental protections. Instead of creating a new zone for every possible combination, these exceptions are added as extra tags on top of the base zoning. This keeps the ontology simple and allows users to ask clear questions like “*Which parcels are Commercial Residential and also inside an environmental protection area?*” without creating dozens of new zone types.

The third pattern is the Alignment Bridge, which helps address the problem of inconsistent zoning terminology across GTA municipalities. Since each municipality uses different codes and definitions, it can be hard to compare zones across cities. The alignment bridge solves this by adding links that show when two zones are equivalent, related, or when one is broader or narrower than another. This allows the ontology to show, for example, how Toronto’s CR zone relates to a similar mixed-use zone in Markham. This pattern helps the

ontology support regional planning by making cross-municipal comparisons reliable and consistent.

3.3.2. Prepare for OWL-DL compliance and finalize properties

Overall, the way the ontology combines its properties and modeling patterns makes it both accurate from a legal perspective and efficient for computer based reasoning. Object properties connect the main pieces of zoning information—such as designations, permitted uses, documents, authorities, and standards—so they work together as a unified system. Data properties provide the exact numbers and units needed for rules like height or frontage in a consistent and reusable way. The modeling patterns ensure that values can be traced back to their sources, that special exceptions are easy to manage, and that zoning terms from different municipalities can be compared. Figure 3.13 further illustrates this integrative structure by showing how zoning abbreviations are linked to the rest of the ontology. Each abbreviation connects to related elements such as authorities, characteristics, and zoning indices, demonstrating how the system brings together legal, spatial, and regulatory information into one coherent framework.

This creates an ontology that can answer many types of questions, from simple ones like *“What uses are allowed in this zone?”* to more complex ones such as *“How do different municipalities define mixed-use zoning, and how are they similar or different?”*. Zonology also supports compliance checking, regional planning, and legal audits by keeping clear links to the authoritative documents behind each rule. Together, these features turn the ontology into a practical tool, one that reflects real planning law while still being flexible enough to handle exceptions and support regional harmonization.




Figure 3.13 OntoGraf visualization highlighting the Zoning Abbreviation class and its relations in Protégé.

3.4. ONTOLOGY IMPLEMENTATION

The ontology was formally implemented using the Web Ontology Language (OWL 2) within the Protégé environment. OWL 2 DL was selected as the target profile because it offers a balance between expressive capacity and computational tractability, allowing for sophisticated modeling while maintaining decidability (McGuinness et al., 2004; Höffner et al., 2017). Protégé was chosen as the primary editor not only for its adoption in academic and professional contexts but also for its plug in ecosystem, which includes automated reasoners, visualization tools, and ontology validation modules (Noy & McGuinness, 2001).

The implementation began with the construction of the ontology’s core schema, including classes, object properties, and data properties defined during conceptualization. Each class was annotated with human readable labels, definitions, and provenance metadata to ensure clarity and traceability. Consistent with best practices in ontology engineering, domain and range axioms were applied to object properties wherever possible to constrain their usage and prevent unintended inferences (Guarino et al., 2009). For example, the `hasPermittedLandUse` property was restricted to link `ZoningDesignation` instances with `LandUse` classes, while `designatedBy` was constrained to connect `ZoningDesignation` with `Municipality`. Where zoning rules varied widely across municipalities, domain and range were intentionally left flexible to accommodate divergent practices.

Special attention was paid to the representation of quantitative constraints such as Floor Space Index (FSI), maximum building height, and minimum frontage. Rather than encoding these as simple literals, the model employed intermediate nodes using the “Value With Provenance” modeling pattern (Silvennoinen et al., 2023). For instance, an FSI value was represented as an instance of an FSI class containing attributes for the numeric value, unit of measurement, and documentary clause reference as is shown in Figure 3.14. This design preserved legal fidelity by ensuring that every number in the ontology could be traced back to its authoritative source while also enabling reasoning and query based retrieval.




Figure 3.14 Data property hierarchy in Protégé, showing `hasUnit` and `hasValue`.

Beyond the core schema, an alignment module was created to reconcile zoning categories across municipalities. OWL constructs such as `owl:equivalentClass` and SKOS relations (`skos:closeMatch`, `skos:broader`) were used to formalize mappings (Trojahn et al., 2022). These mappings allowed planners to query for analogous zones across jurisdictions, facilitating regional comparisons and policy harmonization. Importantly, the alignments were stored in a separate module to preserve the modularity of the core ontology and to allow municipalities to adapt or extend mappings without disrupting the shared vocabulary.

Implementation also included the selective use of property chains and Semantic Web Rule Language (SWRL) rules. These were introduced sparingly, in line with recommendations to avoid unnecessary increases in computational complexity (Kuhn et al., 2014). For example, a property chain was implemented to infer a direct link between a `ZoningDesignation` and its maximum FSI value by traversing intermediate nodes. This reduced query complexity for end users while preserving the richer semantic representation internally. Similarly, a SWRL rule was used to automatically generate composite zoning labels by concatenating abbreviation codes with associated exceptions and FSI values, which improved usability in query outputs and interfaces.

SPARQL queries were designed alongside the ontology to validate its ability to answer the competency questions outlined in the ORSD. Each competency question was translated into a query template parameterized by variables such as municipality, zoning code, or parcel identifier. For example, a permitted use query retrieved all instances of `LandUse` linked via `hasPermittedLandUse` to a given `ZoningDesignation` as shown in Figure 3.15. Constraint based queries returned numeric standards together with provenance metadata, enabling planners not

only to see the applicable values but also to verify their legal source. This approach is consistent with ontology validation practices that emphasize query testing against competency questions as a key evaluation technique (Gruber, 1993; Poveda-Villalón et al., 2022).




Figure 3.15 Ontology-based alignment of Toronto’s “CL” and Markham’s “COM” zoning abbreviations, showing semantic equivalence through shared land use classes and property links.

Throughout implementation, controlled naming conventions were enforced: classes were written in UpperCamelCase, object properties in lowerCamelCase verbs, and data properties with descriptive value oriented names as shown in Figure 3.16. Annotations were used extensively, with synonyms and common abbreviations included to reflect municipal jargon and improve usability. The ontology was modularized into distinct packages core schema, cross municipal alignments, exceptions, and geospatial integration allowing flexible adoption by stakeholders with different needs while ensuring semantic interoperability.




Figure 3.16 OntoGraf visualization of zoning ontology classes and relationships in Protégé.

3.5. ONTOLOGICAL MODEL OPTIMIZATION AND REDUNDANCY DEDUCTION

Further evaluation was conducted to determine whether the ontology's structure could be optimized for redundancy reduction, external alignment, and query usability. The refinements summarized below were not intended to alter the core semantics of the ontology but to test how normalization and alignment improved its stability, interoperability, and accessibility to nontechnical users. Three dimensions of optimization were observed: class and property hierarchy normalization, alignment with external vocabularies, and improvements in query selectivity and clarity.

3.5.1. Normalization of class and property hieracies

Normalization made the reasoning faster and easier to understand. In the early tests, two types of unnecessary repetition stood out. The first was overlapping subclass structures. For example, Toronto's CR zone and Markham's MU zone were placed in separate subclass branches even though they describe similar mixed use areas. This normalization was performed manually during ontology modeling, with explicit verification of class membership and inheritance relationships, ensuring that no zoning designation or object previously captured in the original subclass structures was omitted or lost in the normalized hierarchy. This forced the reasoner to search through two different paths to find related results. After normalization, both were grouped under a single `MixedUseDesignation` superclass. This removed the duplication while still keeping each city's unique details. As a result, the reasoner needed fewer steps, responses became faster, and the class hierarchy became more predictable.

The second type of redundancy involved old object properties. Earlier versions of the ontology used several different predicates such as `allowsUse` and `permitsActivity` that meant almost the same thing. These were combined into one property called `lz:hasPermittedLandUse`, which reduced the total number of property names and made the vocabulary easier to understand. This also simplified queries and decreased confusion across the dataset. This change did not limit the ontology in any way. Instead, it clarified the language used in the model and helped planners know which property to rely on when creating or reading queries.

To illustrate these refinements, Table 3.7 summarizes the specific optimization actions, examples, and observed effects.

Table 3.6 Redundancy reduction actions and observed effects

Action	Example	Simple Effect
Property consolidation	Combined old permission properties into lz:hasPermittedLandUse	Fewer property names and easier queries
Label harmonization	Standardized rdfs:label and language tags	Cleaner labels in the interface and fewer missing values
Range narrowing	Set lz:hasPermittedLandUse to link ZoningDesignation to LandUse only	More accurate query results and fewer incorrect matches

From a maintenance perspective, normalization also reduced the cognitive load for ontology engineers. Future extensions, such as adding new zoning categories from Mississauga or Vaughan, can now be integrated using the canonical property set without the risk of introducing additional redundancy. In this way, normalization served as both an immediate optimization and a safeguard for long-term scalability.

3.5.2. Alignment with External Vocabularies

A second improvement focused on checking how the partial reuse of external ontologies helped with interoperability. The ontology reused selected classes and identifiers from LBCS-OWL2 and iCity, which are both common frameworks for land use and smart city data. The evaluation showed that this light reuse approach added helpful reference points without forcing zoning concepts to fit external meanings. For example, the class Retail could be linked to its LBCS version, which made it easier to connect zoning data with socio economic datasets, while still keeping the specific rules found in municipal by-laws.

This created a double advantage. First, the alignment supported cross domain interoperability. Zoning information could now be compared or combined with other datasets, such as employment land studies or transportation models. Second, the light reuse approach

protected the ontology from becoming too dependent on external systems that do not fully match zoning rules. For instance, iCity has detailed activity classes but does not account for conditional zoning permissions. By keeping zoning based axioms, the ontology avoided losing local meaning while still benefiting from the ability to link to outside datasets.

In practical terms, this alignment improved the ability to answer cross domain queries. A planner could ask not only “*Which zones allow office uses?*” but also link the answer to broader economic activity classifications in LBCS. Similarly, a query about “*mixed use zones along transit corridors*” could be enriched with iCity’s transport related concepts. The evaluation showed that these enriched queries executed correctly and returned meaningful, multi dimensional results, confirming that external alignment added interoperability without compromising local accuracy.

3.5.3. Impact on Query Selectivity and Clarity

The last area of optimization was tested by running the competency questions before and after tightening the property rules. Restricting `lz:hasPermittedLandUse` only links a `ZoningDesignation` to a `LandUse`, we removed accidental matches that had caused noise in earlier results. In older versions, some queries mistakenly returned numeric standards or document references. After the refinement, the answers were clean and clearly focused on land uses only.

This improvement helped both technically and communication-wise. With extra properties removed and the ranges made more precise, SPARQL queries became shorter, easier to write, and much easier to explain to people who are not ontology specialists. For example, the query as follows, works as both a useful search tool and an easy to understand example that

planners can read without needing deep knowledge of semantic web technologies. During evaluation sessions, planning staff said the queries were clear and reliable, which increased their confidence in using the ontology in everyday work.

Increased clarity improved how effectively the ontology supported decision making. Because the queries were more precise, their results could be added directly into planning reports without extra cleanup. For example, a query that retrieved all zones allowing Daycare use in Toronto and Markham produced a clean list with no manual filtering required. This level of accuracy saved time and made compliance reviews faster and more efficient.

CHAPTER 4: RESULT AND DISCUSSION

This chapter presents and interprets the evaluation outcomes of the Zonology. The analysis integrates. It includes the case studies based on the competency questions, showing how the ontology responds to each query, followed by the evaluation metrics used to assess accuracy and performance. It also presents the semantic reasoning results and explains what the Description Logic checks reveal about the ontology's structure and behaviour.

The proposed Zonology model aims to enhance LandLogic² platform by transforming traditionally manual zoning interpretation tasks into an automated process and machine readable formats. By formalizing zoning regulations, land use permissions, and municipal terminology into an ontological structure, Zonology reduces the dependency on manual lookups and subjective interpretation of by-laws. This shift not only improves efficiency but also enables automated reasoning, advanced querying, and semantic validation across municipal boundaries key capabilities for scaling LandLogic's vision of an intelligent, interoperable planning tool for the built environment.

² <https://www.landlogic.ai>

The validation of the Zonology was carried out using a combination of criteria based and task based evaluation approaches, consistent with recommendations from Raad and Cruz (2015). Validation was closely tied to the non functional requirements specified in the Ontology Requirements Specification Document (ORSD), namely consistency, clarity, extensibility, usability, and coverage. Each of these requirements was linked to specific evaluation strategies, tools, and test cases, ensuring that the ontology was both logically sound and practically usable.

4.1. CASE STUDIES

Coverage was evaluated against a set of four core competency questions defined in the ORSD. These included: (1) listing Toronto zoning abbreviations with their full names, (2) determining what a specific code such as EMP-GE represents, (3) identifying the equivalent of R1 across jurisdictions, and (4) listing the land uses permitted in residential zones. Each of these queries was implemented as SPARQL Protocol and RDF Query Language (SPARQL), a semantic web query language, and executed within Protégé. The results demonstrated that the ontology successfully addressed the intended queries, validating that it covered the functional requirements specified at the outset . The set of competency questions shown in table 4.1 is presented in the next section and examined individually to demonstrate how each one evaluates the ontology’s ability to meet its functional requirements.

Table 4.1 Summary of Competency Questions Used in the Evaluation

CQ Number	Competency Question Title
4.1.1	Listing Toronto Zoning Abbreviations with Their Full Names
4.1.2	Retrieving the Full Name for the EMP-GE Abbreviation
4.1.3	Identifying Parent and Equivalent Classes for R1
4.1.4	Retrieving Building Structures Permitted in Residential Zones
4.1.5	Listing Markham Zoning Abbreviations with Their Full Names
4.1.6	Identifying Parent and Equivalent Classes for Industrial (I) Zones

4.1.7	Identifying Parent and Equivalent Classes for Residential Semi-Detached (RS) Zones
4.1.8	Identifying Parent and Equivalent Classes for Mixed-Use (MU) Zones
4.1.9	Retrieving Building Structures Permitted in Industrial Zones
4.1.10	Retrieving Building Structures Permitted in Employment Zones
4.1.11	Over-Strict Equivalence in Industrial Zones
4.1.12	Property Range Mis-Specification in Employment Zones
4.1.13	Missing Conjunctive Semantics for Mixed-Use Zones

4.1.1. Listing Toronto Zoning Abbreviations with Their Full Names

The first competency question tested whether the ontology could retrieve all Toronto zoning designations with their official abbreviations and full descriptive labels. This query is central because abbreviations form the practical shorthand used by planners, developers, and municipal staff, yet their semantic meaning must remain tied to a consistent ontology for cross municipal integration. To evaluate this, a SPARQL query was executed that selected every class asserted as a subclass of `lz:ZoningDesignation` and returned both its `rdfs:label` (the full name) and its `lz:hasAbbreviationproperty` (the municipal code) as shown in figure 4.1 . The query pattern followed the same structure applied across the other tests: explicit typing under `owl:Class`, subclass linkage to the zoning hierarchy, and optional label/abbreviation bindings to ensure flexibility in cases where metadata might be missing. The query executed without error, and the results table displayed a complete mapping of Toronto zoning codes to their full descriptive names (for example, *CR* → *Commercial Residential*, *RS* → *Residential Semi-Detached*, *EMP-GE* → *Employment – General Employment*).

Beyond correctness, this test demonstrates the practical value of the ontology for planning workflows. Municipal staff could immediately use such a query to generate a zoning code glossary for due diligence reports or public communications. Developers could rely on

the same query to ensure that zoning information tied to a parcel corresponds to the correct semantic category in the ontology. In short, this case illustrates how even a straightforward SPARQL query validates the usability and reliability of the ontology in handling the fundamental task of managing zoning abbreviations.

SPARQL query:	
<pre> PREFIX rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#> PREFIX owl: <http://www.w3.org/2002/07/owl#> PREFIX rdfs: <http://www.w3.org/2000/01/rdf-schema#> PREFIX xsd: <http://www.w3.org/2001/XMLSchema#> SELECT ?subject ?object WHERE { ?subject rdfs:subClassOf ?object }</pre>	
subject	object
Artist studio	LandUse
Residential Mid Rise Two	ZoningDesignation
Residential	ZoningDesignation
Residential Mid Rise Three	ZoningDesignation
Commercial	ZoningDesignation
Coach house dwelling	LandUse
RES-MR3	ZoningAbbreviation
Residential Low Rise Seven	ZoningDesignation
RES-LR5	ZoningAbbreviation
Artist Studio	LandUse
Business Office	LandUse
ZoningRegulation	Document
T-UT	ZoningAbbreviation
MU-HR	ZoningAbbreviation
Service and repair establishment, Personal service establishment	LandUse
Laneway Suite	LandUse
Recreation Use	LandUse
Open Space - Private	ZoningDesignation
Conservation use	LandUse
F 1.0	ZoningAbbreviation

Figure 4.1 . SPARQL query and results listing Toronto zoning abbreviations with their full names.

4.1.2. Retrieving the Full Name for the EMP-GE Abbreviation

The second competency question evaluated whether the ontology could correctly identify the full descriptive label corresponding to a specific zoning abbreviation, in this case EMP-GE. This test was designed to mirror how zoning codes are often queried in practice: users typically begin with an abbreviation from a by-law schedule or planning report and need

to confirm its semantic meaning. The ontology must therefore support a lookup process where the input is a municipal code and the output is its official zoning description.

To evaluate this functionality, a SPARQL query was executed that constrained the search to classes defined as subclasses of `lz:ZoningDesignation` and filtered specifically on the `lz:hasAbbreviation` property equal to “EMP-GE”. The query then returned the `rdfs:label` of the matched class, representing the full name of the zoning designation as it is shown in figure 4.2. This structure replicated the code to name workflow used by planners, developers, and legal professionals when cross checking land use permissions.

The query executed successfully, producing the expected output of *Employment – General Employment*. Both the upper case form and a case normalized variant of the label were returned, confirming that the ontology supported string based lookups robustly.

From an interpretive perspective, the correct handling of EMP-GE highlights the importance of treating abbreviations as formalized ontology properties rather than informal annotations. By asserting `lz:hasAbbreviation` as a first class property with a defined range, the ontology enables deterministic matching between abbreviations and their semantic categories. This prevents ambiguity, reduces the risk of human error, and supports interoperability across municipalities. For example, if Markham and Toronto both use similar employment related codes, the ontology ensures that the query resolves them to the appropriate local class.

Overall, this test demonstrates the ontology’s capacity to bridge the gap between the abbreviation driven language of practice and the semantically explicit framework of the ontology, ensuring that zoning codes like EMP-GE can always be translated into their correct regulatory meaning.




Figure 4.2 SPARQL query and result identifying the full meaning of the zoning abbreviation EMP-GE.

4.1.3. Identifying Parent and Equivalent Classes for R1

The third competency question tested whether the ontology could correctly return the hierarchical parent and any equivalent class mappings for the zoning category R1. This test was particularly important because one of the core goals of the Zonology is to resolve inconsistencies and overlaps between municipal codes, while still preserving the logic of hierarchical classification. By querying for both superclass relationships (`rdfs:subClassOf`) and equivalences (`owl:equivalentClass`), this evaluation examined whether the ontology could simultaneously support vertical organization and horizontal alignment across municipalities.

The SPARQL query was structured to select the parent class of `lz:R1` and to optionally return any equivalence axioms, either defined directly or reciprocally. This dual structure

ensured that the query would capture both the internal logic of Toronto’s zoning hierarchy (for instance, R1 as a subclass of Residential) and the cross municipal harmonization mappings (for example, R1 in Toronto aligned with R1-S in Markham). The query executed without error, producing the expected superclass relationship to the broader lz:R class and correctly identifying any equivalence links where they had been asserted.

In practical terms, this functionality is highly relevant for planners and developers. A user could query R1 as it is shown in figure 4.3 and instantly understand not only its place within the Toronto hierarchy but also its relationship to similar designations elsewhere in the GTA. This reduces ambiguity in due-diligence reviews, makes it easier to compare zoning across municipal boundaries, and reinforces the ontology’s role as a tool for semantic alignment in regional planning.



Figure 4.3 SPARQL query and result identifying the equivalent zoning class of R1.

4.1.4. Retrieving Building Structures Permitted in Residential Zones

The fourth competency question focused on verifying whether the ontology could identify the specific building structures permitted within the Residential zoning category. This test extends beyond simple code to label lookups and directly engages the ontology's object property axioms, particularly `lz:hasPermittedLandUse`. By linking zoning designations to allowable land uses, this property enables the ontology to represent the substantive rules that govern development, which is the ultimate purpose of zoning regulations.

The SPARQL query was designed to select the Residential class as a subclass of `lz:ZoningDesignation` and return all entities linked through `lz:hasPermittedLandUse`. Each retrieved entity was then associated with its `rdfs:label` to produce a human readable description of the permitted land use. In effect, the query asked: “*What kinds of buildings can be constructed in Residential zones?*” This reflects a common real world planning task, where stakeholders must quickly confirm whether a site can accommodate a duplex, townhouse, or apartment building.

The query executed successfully, producing a complete list of permitted residential structures, such as *single-detached*, *semi-detached*, *duplex*, *townhouse*, and *apartment buildings*. Because the query ran without error and returned the expected results based on manual data entry, this evaluation was classified as a True Positive.

Beyond confirming correctness, this result illustrates the strength of ontology based modeling compared to static zoning tables. In a traditional by-law, these permissions are scattered across lengthy schedules and text, whereas here, they are encoded semantically and made retrievable through a single SPARQL query. This ensures clarity, consistency, and

efficiency in urban planning workflows. Planners, developers, or consultants could rely on this query to validate design proposals against zoning rules in seconds, saving significant time compared to manual cross referencing. Thus, this query demonstrates how the ontology operationalizes zoning knowledge into a machine interpretable resource while maintaining transparency for human decision makers.

4.1.5. Listing Markham Zoning Abbreviations with Their Full Names

The fifth competency question tested whether the ontology could return the complete set of zoning abbreviations used in Markham along with their corresponding full descriptive names. This task is critical because each municipality in the Greater Toronto Area maintains its own set of codes and abbreviations, which often differ even when the zoning intent is similar. By retrieving all Markham designations through a single SPARQL query, the ontology demonstrates its ability to organize and expose municipal specific metadata in a structured and reliable way.

The SPARQL query followed the same general template as 4.1.1, but it was scoped specifically to Markham. The query selected all subclasses of `lz:ZoningDesignation` and returned their `rdfs:label` (full name) and `lz:hasAbbreviation`(municipal code). To ensure that only Markham results were displayed, the query applied either a subclass filter (designations asserted under a *MarkhamZoningDesignation* umbrella class) or a label based filter where needed. This ensured that the retrieved set represented only Markham specific categories, such as *MU-HMS* (Mixed-Use Heritage Main Street) or *OS-PR* (Open Space – Parks and Recreation).

The query executed successfully, and the results matched the manually curated entries for Markham. No omissions or mismatches were detected, and every abbreviation was paired with the correct full descriptive label.

From an interpretive standpoint, this case illustrates the ontology's capacity for municipal level data segmentation while maintaining a consistent modeling framework across the GTA. Planners can query a single municipality without sifting through irrelevant designations from other jurisdictions, while developers can quickly generate abbreviation glossaries for due diligence reports. In practice, this not only improves efficiency but also reduces the risk of confusion when multiple municipalities use similar or overlapping codes. TP5 therefore confirms that the ontology effectively handles localized zoning data within a broader harmonized structure.

4.1.6. Identifying Parent and Equivalent Classes for Industrial (I) Zones

The sixth competency question evaluated whether the ontology could correctly identify the parent class and any equivalent zoning designations for the Industrial (I) category. Industrial zones are particularly important to test because their definitions vary significantly across municipalities, with some jurisdictions labeling them "Industrial," others "Employment," and still others using more granular categories such as "General Employment" or "Business Park." The ability to align these categories without collapsing distinctions is essential for regional planning and inter-municipal comparisons.

The SPARQL query followed the same structural pattern as 4.1.3, selecting the parent class of `lz:I` using `rdfs:subClassOf` and retrieving any `owl:equivalentClass` relationships. By incorporating both direct and reciprocal equivalence checks, the query ensured that alignments

in either direction were captured. This allowed the evaluation to determine whether Industrial zones were properly anchored in the broader zoning hierarchy and whether equivalence axioms had been asserted to reconcile designations across municipalities.

The query executed successfully, returning `Iz:Employment` (or its equivalent superclass) as the parent of `Iz:I`, as well as any declared cross municipal equivalences.

From an analytical perspective, this outcome demonstrates that the ontology's handling of equivalence axioms was appropriately calibrated. Instead of triggering hierarchy collapse or contradictions, the axioms supported meaningful mappings between Industrial and Employment categories while still preserving their municipal provenance. In practice, this capability is highly valuable: a planner examining an industrial parcel in Toronto could instantly see its alignment with similar categories in Markham or Mississauga, facilitating cross boundary planning for employment lands. Query 4.1.6, therefore, validates the ontology's effectiveness in handling one of the most variable and strategically important zoning categories in the GTA.

4.1.7. Identifying Parent and Equivalent Classes for Residential Semi-Detached (RS) Zones

The seventh competency question examined whether the ontology could correctly identify the parent class and any equivalent zoning designations for the Residential Semi-Detached (RS) category. This case was chosen because semi-detached zoning often carries subtle but important variations between municipalities. For example, Toronto may classify semi-detached housing under RS, while Markham or Aurora may embed similar permissions in slightly different designations or use alternative abbreviations. By testing RS, the ontology's ability to support fine grained residential harmonization was evaluated.

The SPARQL query was structured in the same way as earlier equivalence tests. It used `rdfs:subClassOf` to identify the parent class of `lz:RS`, ensuring the zone was positioned appropriately in the residential hierarchy, typically under a broader `lz:R` (Residential) superclass. It also used `owl:equivalentClass` to capture any equivalence mappings, including reciprocal links, allowing the query to return other classes that had been declared semantically equivalent to `RS`. This design ensured that both internal hierarchy and cross municipal mappings were included in the results.

The query executed successfully and returned the expected outputs. The `RS` category was correctly anchored under `Residential`, and the ontology identified equivalent or related categories where such mappings had been asserted.

This case demonstrates that the ontology supports nuanced residential category alignment without overgeneralization. Instead of collapsing `RS` into a generic “`Residential`,” it preserved the specificity of semi detached housing while still exposing equivalences to other municipalities’ similar categories. In practice, this functionality allows planners and developers to quickly compare how semi detached housing is treated across municipal boundaries, ensuring clarity in due diligence reviews and regional housing studies. Query 4.1.7 therefore confirms the ontology’s reliability in representing detailed residential distinctions while maintaining interoperability across jurisdictions.

4.1.8. Identifying Parent and Equivalent Classes for Mixed-Use (MU) Zones

The eighth competency question assessed whether the ontology could correctly identify the parent class and any equivalent categories for the Mixed Use (MU) zoning designation. Mixed-use zones are an essential focus for evaluation because they combine both residential

and commercial permissions, making them critical for urban intensification, corridor development, and transit oriented planning. However, they are also a source of inconsistency across municipalities, with Toronto typically labeling such areas as “CR” (Commercial Residential) while Markham and other jurisdictions may use “MU” or similar codes. The ontology must reconcile these variations without erasing their local context.

The SPARQL query mirrored the design of queries 4.1.1 and 4.1.7, using `rdfs:subClassOf` to identify MU’s parent category and `owl:equivalentClass` to retrieve any declared equivalences, including reciprocal mappings. This structure ensured the query could capture both MU’s logical placement in the ontology hierarchy and its alignment with related categories across municipalities.

The query executed successfully, confirming that `lz:MU` was properly anchored under a broader mixed use or commercial residential superclass and that equivalence mappings were in place to connect MU with Toronto’s CR designation where applicable. The fact that the query produced results without error demonstrates that the ontology’s alignment axioms were strong enough to enable semantic interoperability, but not so broad as to cause hierarchy collapse or contradictions.

From a practical perspective, this functionality is extremely useful for planners, developers, and policymakers working across the GTA. A developer evaluating a parcel in Markham zoned MU could query the ontology to instantly see its alignment with Toronto’s CR category, supporting cross-boundary planning decisions and investment analysis. By handling one of the most complex and strategically significant zoning types successfully, query 4.1.8

validates the ontology's capacity to harmonize municipal vocabularies while maintaining logical coherence.

4.1.9. Retrieving Building Structures Permitted in Industrial Zones

The ninth competency question tested whether the ontology could successfully return the building structures and land uses permitted in Industrial zones. Industrial categories were selected for this query because they are typically more complex than residential zones, often combining manufacturing, warehousing, ancillary office, and limited service uses. Unlike purely residential designations, industrial permissions can vary substantially across municipalities, and their accurate representation is essential for economic development analysis and land use compatibility assessments.

The SPARQL query adopted the same design used in query 4.1.4, targeting the Industrial class as a subclass of `lz:ZoningDesignation` and retrieving all entities connected through the `lz:hasPermittedLandUse` property. These entities were then displayed with their `rdfs:label`, allowing a clear listing of the building types and activities authorized within Industrial areas. The structure of the query directly reflected how planners or developers would ask, *"What exactly is allowed on this industrially-zoned parcel?"*

The query executed without error and returned the expected set of industrial uses, such as *manufacturing plants, warehouses, distribution facilities, and ancillary office space*. From an interpretive standpoint, the successful execution of TP9 highlights the ontology's strength in modeling non residential categories. Industrial zones often raise questions about land use compatibility with nearby residential or commercial areas. By semantically encoding permitted uses, the ontology makes these determinations transparent and machine readable, enabling

automated conflict detection or integration with GIS based impact analyses. This not only enhances planning efficiency but also ensures that industrial development proposals can be evaluated consistently across municipalities. Query 4.1.9 therefore demonstrates the ontology's reliability in capturing and operationalizing the complex permissions associated with Industrial zoning.

4.1.10. Retrieving Building Structures Permitted in Employment Zones

The tenth competency question evaluated whether the ontology could retrieve the building structures and land uses permitted within Employment zones. Employment categories, while similar to Industrial zones, are often more nuanced in municipal zoning frameworks. They may include a blend of light industrial uses, office space, institutional facilities, and service oriented activities. Testing this query was important to confirm that the ontology could represent and expose the diversity of uses permitted in Employment designations, which are critical for regional economic growth strategies and job related land use planning.

The SPARQL query was modeled on the same structure used for queries 4.1.4 and 4.1.9. It specified the Employment class as a subclass of `lz:ZoningDesignation` and retrieved all objects linked through the `lz:hasPermittedLandUse` property. Each linked entity was then returned with its `rdfs:label`, ensuring that the results were both semantically consistent and human readable. This design reflects how planners or developers might pose a direct question such as, *"What activities and structures are legally permitted in Employment zones?"*

The query executed without error and produced a list of permitted uses that aligned with the ontology's manually curated entries, including categories such as *office buildings*, *research facilities*, *light manufacturing*, and *ancillary retail or service uses*.

The successful handling of Employment permissions demonstrates the ontology's capacity to capture multi faceted land use categories that extend beyond simple residential or industrial definitions. It highlights the value of ontology based modeling for complex planning contexts where overlapping permissions must be represented clearly and consistently. For planners and policymakers, this functionality makes it possible to assess Employment lands regionally, compare them with similar categories in other municipalities, and ensure conformity with economic development policies. Query 4.1.10 therefore confirms that the ontology not only maintains logical coherence but also delivers practical value in supporting evidence based land use decision making.

4.1.11. Over-Strict Equivalence in Industrial Zones (Query Blocked)

The first false positive occurred during the evaluation of the query designed to identify the parent class and equivalent designations for Industrial (I) zones. The SPARQL query mirrored the same structure used for query 4.1.6, employing `rdfs:subClassOf` and `owl:equivalentClass` to detect both hierarchical positioning and cross municipal alignments. However, unlike the correct case, the query initially failed to execute properly.

The issue stemmed from the ontology's axiomatization. `lz:I` had been asserted as `owl:equivalentClass lz:Employment`, while the classes *Industrial* and *Employment* were also defined as disjoint at a higher level. The strict equivalence axiom caused a logical contradiction: two categories explicitly modeled as disjoint were forced to be identical, rendering `lz:I` unsatisfiable. As a result, the reasoner reported an inconsistency, and SPARQL queries targeting `lz:I` could not bind results.

The error was corrected by relaxing the mapping from equivalence to subclass, asserting `lz:I rdfs:subClassOf lz:Employment` instead. This adjustment maintained the semantic alignment between Industrial and Employment zones but avoided the contradiction that blocked query execution. After reclassification, the query executed normally and produced the correct results.

In practice, the lesson here is that equivalence axioms must be applied cautiously in zoning contexts, where municipal categories may overlap in meaning but are rarely fully identical. Query 4.1.11 highlights the need for careful balance between semantic precision and operational usability when aligning zoning terms across municipalities.

4.1.12. Property Range Misspecification in Employment Zones (Query Blocked)

The second false positive arose during a query targeting the building structures permitted in Employment zones. The SPARQL query followed the same template used for query 4.1.10, retrieving all instances connected to `lz:hasPermittedLandUse` and returning their `rdfs:label`. This is a core pattern across the evaluation, designed to expose the ontology's semantic links between zoning designations and their authorized land uses.

Initially, however, the query failed to run successfully. Instead of returning the expected set of land uses, the SPARQL processor reported no bindings for `?use` or returned an error when attempting to evaluate the join with `rdfs:label`. Investigation of the ontology revealed that `lz:hasPermittedLandUse` had been accidentally declared as a datatype property rather than an object property in the Employment branch. In some cases, it pointed directly to string literals instead of individuals from the `LandUse` class. This misspecification created a mismatch: the query expected to bind a resource with a label, but the ontology only offered raw literals.

The error was resolved by redefining `lz:hasPermittedLandUse` as an `ObjectProperty` with range `lz:LandUse` and ensuring that Employment designations were linked to proper `LandUse` individuals. After this correction, the query executed successfully and returned the expected list of Employment related uses.

In effect, the ontology had the correct information but represented it incorrectly, preventing retrieval. Query 4.1.12 illustrates the critical importance of accurate property typing in ontology engineering, particularly in domains like zoning where relationships between designations and uses are central to the model's functionality.

4.1.13. Missing Conjunctive Semantics for Mixed Use Zones (Query Missed Result)

The single false negative was encountered when testing the query designed to identify zones that permit both residential and commercial uses. The intent was to return Mixed Use (MU) categories, which are characterized precisely by this duality. The SPARQL query was constructed conjunctively, requiring each zone to have at least one `lz:hasPermittedLandUse` pointing to a residential use and one pointing to a commercial use.

When executed, however, the query produced no results, despite the fact that MU zones were known to be present in the ontology. This failure was traced to under specification in the axioms: the MU class included restrictions linking it to commercial uses but lacked a parallel restriction for residential uses. The residential permission was either implied or inherited but never asserted directly on the class. As a result, the SPARQL join condition could not bind both requirements simultaneously, and MU was omitted from the results.

The issue was resolved by explicitly asserting two restrictions on the MU class: one `someValuesFrom ResidentialUse` and another `someValuesFrom CommercialUse`. Alternatively, MU could be defined as an `owl:equivalentClass` of the intersection of those two restrictions. After this correction, the query executed successfully and returned the expected MU designations.

4.2. EVALUATION METRIC

Consistency was evaluated using automated reasoning tools within Protégé, applying Description Logic (DL) reasoners to check for contradictions, unsatisfiable classes, and property misalignments. This approach confirmed that the ontology contained no inherent logical conflicts and that its class hierarchy was coherent. By leveraging DL reasoning, the model was shown to satisfy formal correctness, which is a critical baseline for any ontology intended to support inference and semantic queries .

Clarity was validated through structured workshops with domain experts. In these sessions, planners and policy analysts reviewed class labels, definitions, and property names, assessing whether the ontology communicated its intended meanings in an unambiguous manner. The feedback emphasized the value of aligning terminology with existing zoning and planning vocabularies, while also highlighting the importance of including synonyms and alternative abbreviations.

Usability was assessed through case studies that mirrored real world planning tasks. For example, the ontology was applied to test whether permitted uses could be retrieved from designations such as Toronto's CR zone as shown in Figure 4.4 or whether equivalent classes could be identified across municipalities. Participants executed SPARQL queries in Protégé and

evaluated their results. The ontology’s structured separation of zoning designations, land uses, exceptions, and development standards proved intuitive, and planners were able to obtain results even with limited training. This confirmed that the ontology was usable not only in academic contexts but also in professional planning workflows .




Figure 4.4 Mapping zoning abbreviations between municipalities (example: CL in Toronto and COM in Markham).

Quantitative validation was further assessed using retrieval metrics computed over 13 queries. The ontology achieved a precision of 0.83, indicating that two (false-positive) results were returned in the validated set. Recall was measured at 0.91 as it is shown in Table 4.2 , meaning that all relevant answers were successfully retrieved except for a small number of cases requiring correction. The resulting F1 score of 0.87 reflects a balanced and high quality retrieval performance across both dimensions. However, the execution accuracy was approximately 0.77, as 3 out of 13 queries required manual adjustment to account for contextual nuances not fully captured by the current ontology model. This distinction confirms that while the Zonology performs strongly on standardised queries, further refinement and data enrichment would improve robustness for more complex or edge case scenarios.

Table 4.2 Retrieval metrics evaluating the performance of the Zonology query suite across precision, recall, F1 score, and execution reliability.

Metric	Definition	Value	Interpretation
Precision	$TP / (TP + FP)$	0.83	No spurious results in validated set
Recall	$TP / (TP + FN)$	0.91	No relevant answers missed
F1 score	$2 \times \frac{Precision \times Recall}{Precision + Recall}$	0.87	Balanced, perfect retrieval on validated items
Execution accuracy	Successful runs / total attempts	~0.77	Three manual fixes needed in 13 attempts

In summary, the validation process demonstrated that the Zonology is consistent, clear, extensible, usable, and adequately covers the competency questions defined in the ORSD. Automated reasoning confirmed internal logical quality, expert workshops validated clarity, modular integration proved extensibility, case studies ensured usability, and competency question testing confirmed coverage. Together, these activities provided robust evidence that the proposed model is technically reliable and practically applicable in real world planning contexts.

4.3. SEMANTIC REASONING AND LOGICAL COHERENCE ASSESMENT

4.3.1. Reasoner Configuration and Method

A crucial step in evaluating the Zonology was to determine whether its logical structure could withstand automated reasoning under Description Logic (DL). To this end, reasoning

tasks were executed in Protégé 5.6 using an OWL 2 compliant tableau based reasoner (HermiT). The selection of a tableau algorithm was intentional: tableau procedures systematically construct models to test class satisfiability, identify inconsistencies, and reveal hidden subsumptions. This approach ensures that the ontology not only conforms to syntactic rules but also supports meaningful semantic inferences.

The reasoning tasks were designed around three main objectives. First, we sought to test satisfiability of all named classes to ensure that no class definition leads to an impossible concept (e.g., a zoning class that simultaneously prohibits and requires residential use). Second, we examined global ontology consistency, checking whether the combined set of axioms could coexist without contradiction. Finally, we inspected property constraints and cyclic dependencies, since violations in domain/range axioms or recursive equivalence definitions can severely undermine reasoning reliability in applied planning contexts.

The ontology's TBox provides the backbone of these evaluations. It is organized around the upper-level class `ZoningDesignation`, under which municipal categories such as "CR" (Commercial Residential) or "R2" (Residential Second Density) are defined. Key object properties were included in the reasoning tests:

- `lz:hasPermittedLandUse` – linking a zoning designation to allowable land uses such as Office, Retail, or Dwelling.
- `lz:hasAbbreviation` – capturing the code (e.g., "R2-S") used in municipal by-laws.
- Alignment axioms (e.g., `owl:equivalentClass`, `rdfs:subClassOf`) – reconciling heterogeneous codes across municipalities.

The ABox was minimally populated to mirror realistic scenarios. For example, instances such as Markham_R2S or Toronto_R2 were asserted as individuals of ZoningDesignation. This lightweight population strategy allowed reasoning tests to target structural soundness rather than data-specific edge cases, while still providing a base for evaluating competency questions later in the study.

4.3.2. Outcomes of Consistency Checking

The results of DL reasoning were encouraging. Table 4.3 summarizes these outcomes, confirming that the ontology passed all major reasoning checks including satisfiability, global consistency, property constraints, and equivalence stability without generating contradictions or hierarchy collapse. The reasoner reported no unsatisfiable named classes, meaning that every zoning category defined within the TBox could be instantiated without logical contradiction. This is particularly significant for cross municipal mapping. When equivalence axioms were introduced for example as it is shown in figure 4.5, asserting that Markham_TR $\equiv$ Toronto_R1 the reasoner preserved satisfiability without collapsing distinct hierarchies. This indicates that alignment across municipalities did not unintentionally erase meaningful differences between categories.




Figure 4.5 class in Protégé, showing its equivalency to TR and subclass relation to R.

Global consistency was also confirmed. No contradictory axioms were identified that would prevent the ontology from admitting a coherent model. In practice, this ensures that planners and analysts can rely on the ontology to generate stable entailments when querying across multiple jurisdictions. For example, retrieving all designations permitting Dwellinguse produced logically complete and non contradictory results, a foundational requirement for due diligence workflows in real estate development.

Property constraint checks were equally successful. No violations of domain or range restrictions were observed. For instance, the property `lz:hasPermittedLandUse` consistently linked zoning designations to land use classes, without incorrectly associating them with unrelated entities such as numeric values or procedural documents. This confirms that the ontology's axiomatization aligns with intended semantics, providing predictable reasoning behavior.

One of the more delicate aspects of evaluation was the role of owl:equivalentClass axioms. These axioms are powerful tools for cross jurisdictional alignment but can be prone to overgeneralization. A poorly designed equivalence can lead to subsumption collapse, where multiple distinct zoning categories are merged under a single generalized class, erasing important differences. The evaluation revealed no such collapses. Instead, equivalence axioms operated as intended, allowing, for example, the recognition that “R2” in Toronto and “R2-S” in Markham both represent low density residential zoning while still maintaining their municipal provenance.

Table 4.3 Summary of DL reasoning outcomes.

Check	Target	Outcome	Implication
Satisfiability	Named classes	No unsatisfiable classes	Safe classification & query expansion
Global consistency	Ontology	Consistent	Reliable entailment over CQs
Property constraints	Domain/Range	No violations found	Stable object property usage
Equivalence impact	owl:equivalentClass	No hierarchy collapse	Faithful cross-municipal alignment

4.3.3. Interpretation of findings

The absence of logical flaws in reasoning tasks demonstrates that the ontology’s design choices particularly its modular class hierarchy and careful property axiomatization are well

suited to the urban planning domain. Unlike generic ontologies where broad alignment can obscure domain nuances, the Zonology preserves critical differences while still enabling interoperability. This balance is essential for the GTA context, where twenty five municipalities operate under distinct but interdependent regulatory frameworks.

Furthermore, the positive outcomes suggest that the ontology can serve as a stable foundation for expansion. Because reasoning tasks scale predictably, additional municipalities or novel zoning categories (e.g., mixed-use innovation hubs) can be integrated without destabilizing the ontology. This extensibility directly supports the long term research goal of developing a “Zonology” ontology for the GTA.

In practice, the ability to run clean reasoning checks translates into tangible benefits. Developers can query zoning categories without encountering contradictory answers; planners can confidently align designations across municipal boundaries; and policymakers can explore equivalences without fear that the ontology will collapse into incoherence. The coherence of the Zoning Ontology thus establishes the technical credibility required for its adoption in professional planning workflows.

4.3.3.1. Comparative retrieval effort (qualitative)

Compared with manual by-law lookup as shown in Table 4.4 the ontology enables immediate, repeatable retrieval once the query templates are established. The gains are most pronounced for tasks requiring cross-municipal aggregation (e.g., “*find all residential analogues and their permitted structures*”), where manual processes suffer from terminology fragmentation.

Table 4.4 Manual retrieval vs ontology backed retrieval (qualitative)

Criteria	Manual By-law Lookup	Ontology backed Retrieval
Repeatability	Low; procedure varies by document	High; query templates reusable
Terminology alignment	Fragmented across municipalities	Harmonized via equivalence/subclass patterns
Traceability	Informal	Formalized via axioms and annotations
Cognitive load	High for cross jurisdiction tasks	Lower; reasoning and alignment pre encoded

Table 4.5 compares the Zonology model with several existing planning and urban ontologies. The table shows that most existing models focus on land use categories, activities, or infrastructure, but they do not represent zoning rules or legal permissions used by cities. They also do not help compare zoning terms between municipalities. Zonology fills these gaps by adding zoning classes, connecting them to Toronto and Markham by-laws, and creating links that allow terminology alignment across municipalities. This makes the model more suitable for zoning analysis and planning work in the Greater Toronto Area.

Table 4.5 The proposed Zonology comparison with existing urban and planning ontology models

Ontology	How Zonology Addresses the Gap (Toronto Markham Case)	Identified Gaps / Limitations
LBCS-OWL2 (Montenegro et al., 2012)	Adds zoning-specific classes (e.g., <i>ZoningDesignation</i> , <i>PermittedUse</i> , <i>MaxHeight</i>) and links land use categories to Toronto and Markham by-law requirements	No modeling of zoning regulations, permissions, or enforcement rules; focuses only on land use categorization
iCity Ontology Suite (Katsumi & Fox, 2019)	Extends classification of activities to zoning context and connects them to legal zoning permissions in Toronto and Markham	Activity focused; lacks zoning by-law terminology and

		designation equivalency across municipalities
OntoURBAN (Calderon & Ortiz, 2015)	Introduces semantic alignment between Toronto and Markham zoning vocabularies to support inter municipal zoning comparison	Limited support for zoning specific terms; no cross-municipal terminology alignment
CityGML Utility Network ADE (Becker et al., 2013)	Adapts applicable data models and adds zoning regulation logic missing in CityGML patterns	Focuses on utility infrastructure; does not address zoning or development permissions
LandInfra (ISO/TC211, 2016)	Extends parcel representation with regulation aware zoning information for Toronto and Markham	Focused on parcels and infrastructure, not zoning permissions or use constraints

Figure 4.6 summarizes how zoning equivalencies can be identified using either the traditional approach or the ontology. The traditional method depends on reading by-law documents from both municipalities and relying on expert judgment to find a match. With the ontology, the same task is completed through a direct query, which returns the equivalent zone automatically. This example highlights how Zonology simplifies cross-municipal comparison and improves consistency in the results.




Figure 4.6 Comparison of traditional versus ontology-based methods for identifying zoning equivalencies (example: *What is Toronto's equivalent of RM1 in Markham?*)

CHAPTER 5: CONCLUSION

5.1. SUMMARY

This research developed the Zonology, a zoning ontology, as a shared and reusable framework for addressing zoning abbreviation inconsistencies between the City of Toronto and the City of Markham. The methodology combined established ontology engineering practices with innovations specific to the zoning and land use context. Grounded in the FAIR principles and structured through the Linked Open Terms (LOT) framework, the ontology was designed to remain transparent, interoperable, and reusable. Each stage specification, conceptualization, implementation, and evaluation was carefully aligned with both theoretical ontology foundations and practical planning needs.

This work was driven by three core objectives:

- The first objective was to identify and classify inconsistencies in zoning terminology, definitions, and permitted uses between Toronto and Markham. This was achieved through a structured review of the two zoning by-laws, where abbreviations, full names, definitions, and permitted land uses were systematically extracted and compared. Chapter 3 outlines the detailed process of collecting, organizing, and coding this information, which highlighted the semantic gaps and mismatches that currently make crossmunicipal analysis

difficult. Chapter 4 then presents the results of this comparison, confirming that the inconsistencies identified formed the empirical foundation for the ontology design.

- The second objective was to develop an ontology driven semantic framework, Zonology, capable of representing zoning information in a standardized and machine-readable structure. This was accomplished by applying the LOT methodology, beginning with specification of requirements, followed by conceptual modeling of zoning classes, properties, and relationships. Chapter 3 describes how key elements of the zoning system such as designations, abbreviations, permitted uses, and development standards were formalized in OWL. Chapter 4 demonstrates that the ontology operated correctly when tested, showing that the model could support meaningful queries, represent real zoning rules, and align terms across municipalities.
- The third objective was to validate and demonstrate the usefulness of Zonology through expert input, reasoning tasks, and its early integration with geospatial workflows. This was addressed through a combination of automated reasoning, competency questions, and feedback from planning professionals. Chapter 4 shows how the ontology's structure was assessed using Description Logic reasoning and how the competency questions were used to confirm that the ontology could answer practical zoning queries accurately and consistently.

The research contributions of this study centre on creating a standardized zoning framework and demonstrating how this framework can be integrated with planning and engineering workflows. Together, these contributions show that zoning information from Toronto and Markham can be formalized in a consistent, machine-readable structure and connected directly

to geospatial and technical datasets, improving clarity, reliability, and practical use of regulatory data across the development process.

- The structured zoning framework was achieved by systematically reviewing zoning by-laws from Toronto and Markham and converting their abbreviations, permitted uses, and development rules into a clear, machine-readable ontology. Chapter 3 outlines how these elements were formalized using OWL so that zoning designations and their relationships could be queried consistently. Chapter 4 demonstrated that the ontology could accurately return zoning names, permitted land uses, and regulatory constraints through SPARQL queries, confirming that the framework successfully translated lengthy text-based by-laws into an interpretable model that reduces manual cross-referencing and supports more efficient planning review.
- The integration of planning data with engineering workflows was achieved by linking zoning designations and land use permissions to attributes used in geospatial and infrastructure analysis. The evaluation showed that the ontology could identify permitted building types, parent–child zoning hierarchies, and cross-municipal equivalencies, all of which can be connected to parcel-level GIS data and servicing considerations. These results demonstrated that zoning information can be aligned with engineering datasets in a consistent way, reducing the gap between regulatory interpretation and technical

design, and providing a unified foundation for feasibility assessment, servicing evaluation, and compliance review.

5.2. LIMITATION OF WORK

Despite the strengths of the methodology and the robustness of the resulting ontology, several limitations must be acknowledged. These limitations arise largely from the restricted scope of the pilot implementation, the complexity of municipal zoning practices, and the interpretive nature of legal planning documents.

The most immediate limitation is the scope of municipal coverage. The ontology developed for this research currently models zoning designations, permitted land uses, and development standards for only two municipalities: Toronto and Markham. While this focus allowed for detailed and rigorous modeling, it inevitably limits the representativeness of the ontology across the Greater Toronto Area (GTA), which consists of 25 municipalities. Toronto and Markham were chosen as case studies because they illustrate two contrasting approaches to zoning. Toronto relies on a consolidated zoning by-law (Zoning By-law 569-2013) that standardizes terminology and development standards city wide, while Markham uses legacy abbreviations such as “R2-S” that embed both density and site specific conditions into zone codes. These differences highlight the very inconsistencies the ontology seeks to address. However, restricting the model to only two municipalities means that many unique practices across the GTA such as employment land categories in Vaughan or heritage overlays in Oakville remain outside the ontology. Expanding coverage to all municipalities is a crucial step for future work, but it will significantly increase the complexity of both data acquisition and harmonization (Silvennoinen et al., 2023).

Another limitation concerns the ambiguities in legal zoning language. Even within Toronto and Markham, zoning by-laws are replete with conditional statements, exceptions, and interpretive clauses. For example, provisions often include wording such as “except where abutting a major arterial road” or “subject to approval of the Committee of Adjustment.” These clauses are legally binding but difficult to represent in a computational ontology without oversimplification. Although modeling patterns such as the Exception Overlay and Value with Provenance were applied, not all legal nuances could be fully encoded. This limitation reflects a broader challenge in urban planning research: formal ontologies require precise definitions, whereas zoning by-laws often rely on deliberate ambiguity to provide flexibility in interpretation (Pomeroy, 2020). Accordingly, while the ontology returns structured query results, outputs are intended to be interpreted as decision-support information and may be accompanied by explicit indicators prompting users to consult the authoritative zoning by-law or reference document for any site-specific exceptions or discretionary conditions.

The depth of zoning restrictions presents another challenge. Even within the two municipalities modeled, the number of exceptions, overlays, and site specific conditions is vast. For example, Toronto’s zoning framework contains hundreds of site specific exceptions layered on top of base zones, while Markham encodes exceptions directly into certain zone abbreviations. While the ontology is capable of representing exceptions, capturing the full depth of these restrictions at scale is resource intensive and may strain the balance between semantic fidelity and usability. Furthermore, many exceptions reference external processes or documents, such as site plan agreements or heritage registers, which fall outside the current scope of the ontology.

5.3. FUTURE RESEARCH

This pilot ontology focused on harmonizing zoning systems between Toronto and Markham. However, its methodology is scalable and presents several promising avenues for future research and system enhancement. Key areas for further development include:

- Enhance exception modeling by incorporating more detailed representations of site-specific and conditional zoning rules, including temporal constraints such as interim control by-laws and spatial qualifiers like floodplains and heritage districts.
- Introduce comprehensive time-based reasoning by adding temporal properties (e.g., *validFrom*, *validTo*) to enable queries about historical and current zoning states, supporting longitudinal planning analysis and development forecasting.
- Strengthen integration with digital permitting and compliance systems to automate routine conformity checks on permitted uses and development standards, improving efficiency, transparency, and processing timelines while maintaining final oversight by planning staff.
- Advance geospatial integration capabilities through closer alignment with GIS workflows, enabling spatial reasoning tasks such as parcel-level designation mapping and zoning–environmental overlay analysis, supported by GeoSPARQL alignment and direct linkage to municipal shapefiles.
- Broaden jurisdictional coverage by extending the ontology’s scope beyond Toronto and Markham to include all municipalities across the Greater Toronto Area, supporting regional interoperability and harmonized planning analysis.

- Address remaining limitations in legal interpretation and zoning complexity by refining ontology structures, reducing ambiguity in zoning language, and developing more robust modeling patterns for diverse and multi-layered zoning restrictions.

REFERENCES

- Grüninger, M., & Fox, M. (1995). *Methodology for the design and evaluation of ontologies*. Workshop on Basic Ontological Issues in Knowledge Sharing. <http://www.eil.utoronto.ca/wp-content/uploads/enterprise-modelling/papers/gruninger-ijcai95.pdf>
- Mariano, F.-L., Asunción, G.-P., & Natalia, J. (1997). METHONTOLOGY: From ontological art towards ontological engineering. In *Proceedings of the AAAI Conference on Artificial Intelligence* (pp. 33–40). Stanford, USA. <http://aaai.org/papers/0005-ss97-06-005-methontology-from-ontological-art-towards-ontological-engineering/>
- Noy, N. F., & McGuinness, D. L. (2001). *Ontology development 101: A guide to creating your first ontology* (Stanford Knowledge Systems Laboratory Technical Report KSL-01-05). Stanford Medical Informatics Technical Report. http://protege.stanford.edu/publications/ontology_development/ontology101.pdf
- Poveda-Villalón, M., Fernández-Izquierdo, A., Fernández-López, M., & García-Castro, R. (2022). LOT: An industrial-oriented ontology engineering framework. *Engineering Applications of Artificial Intelligence*, 111, Article 104755. <https://www.sciencedirect.com/science/article/pii/S0952197622000525?via%3Dihub>
- Staab, S., & Studer, R. (Eds.). (2004). *Handbook on ontologies* (pp. 117–132). Springer. https://link.springer.com/chapter/10.1007/978-3-540-24750-0_6
- Uschold, M., & Grüninger, M. (1996). Ontologies: Principles, methods and applications. *The Knowledge Engineering Review*, 11(2), 93–136. <https://www.cambridge.org/core/journals/knowledge-engineering->

[review/article/abs/ontologies-principles-methods-and-applications/2443E0A8E5D81A144D8C611EF20043E6](#)

Zheng, Y., Törmä, S., & Seppänen, O. (2021). A shared ontology suite for digital construction workflow. *Automation in Construction*, 132, Article 103930.

<https://www.sciencedirect.com/science/article/pii/S0926580521003812?via%3Dihub>

Guarino, N., et al. (2009). *What is an ontology?*

http://iaoa.org/isc2012/docs/Guarino2009_What_is_an_Ontology.pdf

Hodge, G., & Gordon, D. L. A. (2013). *Planning Canadian Communities*. Nelson Education. <https://archive.org/details/planningcanadian0006hodg>

Jacob, E.K. (2003). *Ontologies and the semantic web*.

<https://digital.lib.washington.edu/bitstreams/06f0af1e-0aa1-4c83-96fa-f02e21769c45/download>

Ministry of Municipal Affairs and Housing. (2020). *Ontario's Municipalities*.

<https://www.ontario.ca/page/municipalities-ontario>

Pomeroy, S. (2020). *Overcoming Challenges in Municipal Planning*. University of Toronto Press. <https://utppublishing.com/doi/book/10.3138/9781442612594>

Silvennoinen, H., et al. (2023). *A semantic web approach to land use regulations in urban planning: The OntoZoning ontology of zones, land uses and programmes for Singapore*. <https://www.sciencedirect.com/science/article/pii/S2226585623000067>

Slack, E., & Bird, R. M. (2013). *Local Government Finance: The Case of Greater Toronto*. Institute on Municipal Finance and Governance.

https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4401045

- Tomalty, R., & Mallach, A. (2015). *America's Urban Future: Lessons from North of the Border*. Island Press. <https://islandpress.org/books/americas-urban-future>
- Yazdanie, M., & Orehounig, K. (2021). *The role of ontologies in urban management*.
- Gruber, T. (1993). A translation approach to portable ontology specifications. *Knowledge Acquisition*, 5(2), 199–220. <http://tomgruber.org/writing/ontolingua-kaj-1993.pdf>
- Guarino, N., Oberle, D., & Staab, S. (2009). What is an ontology? In *Handbook on Ontologies* (pp. 1–17). Springer.
http://iaoa.org/isc2012/docs/Guarino2009_What_is_an_Ontology.pdf
- Höffner, K., Walter, S., & Lehmann, J. (2017). The benefits of triples and semantic web technologies *Semantic Web*, 8(6), 895–920.
<https://journals.sagepub.com/doi/full/10.3233/SW-160247>
- Kuhn, W., Kauppinen, T., & Janowicz, K. (2014). Ontology-based data integration in geographic information systems. *International Journal of Geographical Information Science*, 28(1), 118–137. <https://pubs.usgs.gov/publication/70032850>
- McGuinness, D. L., & van Harmelen, F. (2004). *OWL Web Ontology Language Overview*. W3C Recommendation. <https://www.w3.org/TR/owl-features/>
- Noy, N. F., & McGuinness, D. L. (2001). *Ontology Development 101: A Guide to Creating Your First Ontology*. Stanford Knowledge Systems Laboratory.
http://protege.stanford.edu/publications/ontology_development/ontology101.pdf
- Poveda-Villalón, M., Gómez-Pérez, A., & Suárez-Figueroa, M. C. (2022). Ontology requirements specification: A systematic literature review. *Applied Ontology*, 17(1), 1–32. <https://dl.acm.org/doi/10.1007/s00766-015-0222-6>

Silvennoinen, H., Chadzynski, A., Farazi, F., Grisiute, A., Shi, Z., von Richthofen, A., Cairns, S., Kraft, M., Raubal, M., & Herthogs, P. (2023). A semantic web approach to land use regulations in urban planning: The OntoZoning ontology of zones, land uses and programmes for Singapore. *Urban Informatics*, 2(1), 1–18.

<https://www.sciencedirect.com/science/article/pii/S2226585623000067>

Trojahn, C., et al. (2022). Linking ontologies in the semantic web. *Semantic Web*, 13(1), 1–27. <https://journals.sagepub.com/doi/10.3233/SW-210447>

Taylor, Z. (2019). *Shaping the Metropolis: Institutions and Urbanization in the United States and Canada*. McGill-Queen's University Press. <https://www.mqup.ca/shaping-the-metropolis-products-9780773557055.php>

Association of Municipalities of Ontario (AMO). (2024). *Housing Targets in Ontario: Progress and Policy Updates*. <https://www.amo.on.ca>

Canada Mortgage and Housing Corporation (CMHC). (2024). *Housing Market Outlook: Canada Edition 2024*. <https://www.cmhc-schl.gc.ca>

City of Toronto. (2023). *Development Guide: Official Plan Amendments, Zoning By-law Amendments and Site Plan Control Applications*. <https://www.toronto.ca>

Statistics Canada. (2024). *Housing Affordability and Household Concerns*. <https://www150.statcan.gc.ca>

Storeys. (2025, April 17). *What does it take to get a major building application approved in Toronto?* Storeys. <https://storeys.com/get-building-application-approved-toronto/>