Dissertation Thesis

Representing and Reasoning with Context-Sensitive Vague Place Descriptions

Author:
Madiha Yousaf

March 2023

A thesis submitted in partial fulfillment of the requirements for the degree of Doctor of Natural Sciences (Dr. rer. nat.)

–

Faculty for Information Systems and Applied Computer Sciences
University of Bamberg

Supervisor and Reviewer:
Prof. Dr. Diedrich Wolter

Second Reviewer:
Prof. Dr. Daniela Nicklas

Head of Examining Board:
Prof. Dr. Andreas Henrich

Date of defense: October 14, 2022

Bamberg, 2023

URN: urn:nbn:de:bvb:473-irb-591752
DOI: https://doi.org/10.20378/irb-59175

ZUSAMMENFASSUNG

Verstehen, wie Menschen kognitiv das Konzept Ort wahrnehmen und über Orte kommunizieren, ist Bestandteil interdisziplinärer Forschung und für verschiedene Anwendungsgebiete relevant. Erkenntnisse können beispielsweise zur Verbesserung der Mensch-Computer-Interakti-on im Bereich der Geo-Informationssysteme (GIS) dienen. Angesichts der allgegenwärtigen mobilen Anwendungen, deren Anwenderinnen und Anwender keine GIS-Spezialisten sind, besteht ein wachsender Bedarf an intuitiv nutzbaren ortsbezogenen Informationssystemen.

Ortsbeschreibungen werden alltäglich verwendet, um Informationen über Orte zwischen Personen zu kommunizieren. Dabei treten intuitive Konzeptualisierungen von Orten auf, die oftmals nicht mit der Modellierung in einem Geo-Informationsystem im Einklang stehen. Ortsbeschreibungen werden als qualitative Referenzsysteme betrachtet, die Verweise auf Orte und deren relative Lage zueinander enthalten, aber auch nicht-räumliche Informationen wie Merkmale eines Ortes oder dessen Bedeutung umfassen können. Die Diskrepanz zwischen intuitiver Ortskonzeptualisierung, gepaart mit der Vielfalt an möglichen Beschreibungen in natürlicher Sprache macht es sehr herausfordernd, Orte in geographischen Informationssystemen anhand einer natürlichsprachlichen Beschreibung identifizieren zu können.

Die vorliegende Dissertation ist im Rahmen des Forschungsprojekts „Con -text-Sensitive Qualitative Spatial Reasoning for Interpreting Vague Place Descriptions“ entstanden, welches von der Deutschen Forschungsgemeinschaft (DFG) gefördert wurde. Das Projekt zielte darauf ab, Prinzipien zu erforschen, die natürlichsprachliche Ortsbeschreibungen charakterisieren und ihre automatisierte Interpretation erlauben. Die vorliegende Arbeit betrachtet natürlichsprachliche Ortsbeschreibungen als Eingabe und versucht mit fünf Schritten eine Interpretation im Sinne des Auffindens des beschriebenen Ortes in einer Karte zu machen: Einbeziehung kontextsensitiver Information, Extraktion räumlicher Informationen aus Text und Modellierung, Schlussfolgerungstechniken, Abgleich mit einer geographischen Datenbank (Georeferenzierung), semantische Abfrage.

Der erste Schritt zielt darauf ab, ein Modell zur Repräsentation von vager und unsicherer Information für Information in natürlichsprachlichen Ortsbeschreibungen zu entwickeln. Das besondere Augenmerk wird darauf gelegt, Kontextinformation berücksichtigen zu können. Im zweiten Schritt wird eine Repräsentation von in Text enthaltenem Ortswissen unter Verwendung es entwickelten Modells aufgebaut. Die Repräsentation stellt bezeichnete Entitäten und deren Beziehungen untereinander explizit dar. Um mit den extrahierten Informationen umgehen zu können werden, in einem dritten Schritt Schlussfolgerungsmethoden entwickelt, die extrahierte Information verdichten und mögliche Fehlinterpretationen entdecken können. Hierzu wird insbesondere auch ontologische Information über die Orte verwendet, um eine Georeferenzierung effizient durchführen zu können. Abschließend wird ein Verfahren beschrieben, welches paraphrasierte Orts-

beschreibungen durch semantische Ähnlichkeit versucht aufzulösen.

Die wissenschaftlichen Beiträge dieser Dissertation liegen in der Entwicklung neuer Ansätze zur Interpretation von Ortsbeschreibungen, einem Spezialfall des Textverstehens: a) Eine Wissensrepräsentation für vage und kontextsensitive Information über Orte wird vorgestellt und mit dem Literatur verglichen. b) Es wird gezeigt, wie Schlussfolgerungsmethoden das Textverstehen im konkreten Anwendungsfall verbessern können. c) Ein technisches Verfahren wird entworfen, welches innerhalb mehrerer Schritte Ortsbeschreibungen interpretieren kann und dabei auch unbenannte Orte identifizieren kann. Dazu werden insbesondere Kontextualisierung und ontologisches Schlussfolgern angewendet. d) Eine Bewertungsfunktion anhand semantischer Ähnlichkeit wird vorgestellt, die Entitäten anhand semantisch ähnlicher Typbezeichnungen identifizieren kann.

Im Rahmen dieser Dissertation habe ich den automatisierten Ansatz SORS: Spatio-Ontological Reasoning System zum Interpretieren von Ortsbeschreibungen entwickelt und gebe damit neue Impulse für die interpretation natürlicher räumlicher Sprache. Die durchgeführten Experimente zeigen, dass die vorgeschlagene Methode eine Verbesserung gegenüber existierenden Ansätzen darstellt.

Insbesondere können erstmals auch unbenannte Entitäten (z.B. *Briefkasten* am Bahnhof Bamberg) georeferenziert werden. Die erstellte Implementation integriert die oben genannten Schritte mit weiteren, etwa Visuali sierung, zu einem prototypischen Gesamtsys-tem.

ABSTRACT

In philosophy and psychology, the concept of place has been extensively researched as a human spatial notion. Understanding how people cognitively perceive and communicate about the location to develop place-based information systems that help improve human-computer interactions has lately gained attraction in the field of GIScience (GIS). The growing need for location-based information systems in the face of the emergence of ubiquitous and mobile computing, which does not require certified GIS specialists, adds to the motivation.

Place descriptions are used every day to communicate and transfer information about places between individuals, and therefore they reflect how people cognitively represent and communicate about place knowledge. Place descriptions have been regarded as qualitative reference systems that describe geographic locations that comprise references to places and qualitative spatial relations describing their relative whereabouts, including non-spatial information like place characteristics and semantics.The plethora of natural language place descriptions provide various methods of representing people's understanding of place knowledge.

This doctoral thesis is set within the scope of the research project, "Context-Sensitive Qualitative Spatial Reasoning for Interpreting Vague Place Descriptions," funded by the German science foundation (DFG) aims to design a geographic information system based on locations to make the information about the human location interpretable for computers. This research employs natural language place descriptions as its primary research input and focuses on five research components as main goals towards understanding place descriptions: incorporating context-sensitive information, spatial information extraction, and modeling, reasoning, place geo-referencing, and querying.

The first task is aimed at designing a computational model to represent vague spatial knowledge and context in human-generated place descriptions. The second task focuses on comprehending and understanding the meaning of the place descriptions by creating an organized, explicit depiction of entities and their relation to one another. A flexible reasoning method is developed based on exploiting implicit information and results in inferring new relational knowledge. The third task takes advantage of the ontological type information present in the description and uses them to resolve and identify contextual dependencies to geo-reference all places identified by the system to locate them on the map. The fourth task supports the querying process by providing a query strategy that will help make matching to a geographic database efficient and effective. Additionally, providing an algorithm for queries leading to incorrect or no results using semantic similarity. Last, an algorithm is presented to determine suitable entity types to deal with paraphrased places.

The scientific contributions of this dissertation embody new approaches that have been developed to address the five tasks specifically: a) A computational model is presented which can capture vague and context-sensitive

information about places, and how it accounts for the dimensions of a place is demonstrated by comparing to a baseline model from the literature; b) A reasoning framework for spatial information extraction is presented,contributing to the longstanding goal of understanding spatial language; c) a multi-step geo-referencing approach to automatically interpret natural language place descriptions by enabling us to explicate implied information in the form of unnamed entities embedded in place descriptions, with sub-contributions including contextualization (incorporating type information), ontological reasoning (is-a relation), and inference to generate mappable data; d) A ranking method based on semantic similarity is presented, to query open street map database from entities informally described in natural language. In addition, a clustering and pruning algorithm is provided to generate semantically replaceable entity types for the spatial non-nouns in the description,improving the geo-reference process.

In this thesis, I have developed an automated approach *SORS: Spatio-Ontological Reasoning System* to tackle the problem of interpreting spatial language and provides a new impulse on how to tackle text interpretation tasks. The conducted experiments shows that the practical method performs well and shows advancement over the state of the art. The implementation integrates the above approaches into a complete processing chain and several other visualization and mapping functionalities.

DECLARATION

This is to certify that:

1. the thesis comprises only my original work,

2. acknowledgment has been made in the text to all other material used.

Bamberg, March 2023

Madiha Yousaf

PUBLICATIONS

This thesis is supported by published works during the PhD research project. The major contents of this thesis such as ideas, algorithms, and figures have appeared previously in the following publications:

- Madiha Yousaf,Tobias Schwartz, Diedrich Wolter: 2023. Automated Interpretation of Place Descriptions: Determining Entity Types for Querying OSM. KI-Künstliche Intelligenz 2023.
  DOI: https://doi.org/10.1007/s13218-022-00798-y
- Madiha Yousaf, Diedrich Wolter: November 2021.A Reasoning Model for Geo-Referencing Named and Unnamed Spatial Entities in Natural Language Place Descriptions. In Spatial Cognition and Computation: An Interdisciplinary Journal.Taylor and Francis.
  DOI: https://doi.org/10.1080/13875868.2021.2002872
- Madiha Yousaf, Diedrich Wolter:2019. How to identify appropriate key-value pairs for querying OSM. In Proceedings of the 13th Workshop on Geographic Information Retrieval (GIR '19). ACM, New York, NY, USA, Article 7, 6 pages.
  DOI: https://doi.org/10.1145/3371140.3371147
- Madiha Yousaf, Diedrich Wolter: 2018.Spatial Information Extraction from Text Using Spatio- Ontological Reasoning.GIScience 2018 : 71:1-71:6
  DOI: 10.4230/LIPIcs.GISCIENCE.2018.71
- Madiha Yousaf, Diedrich Wolter: 2018, Core Computations for Supporting Question Answering by Spatio-Ontological Reasoning, GIScience Workshop Core Computations on Spatial Information.
- Diedrich Wolter, Madiha Yousaf: 2017.Context and Vagueness in Automated Interpretation of Place Description: A Computational Model. COSIT (Workshop/Posters) 2017,Short paper,137-142
  DOI: https://doi.org/10.1007/978-3-319-63946-8_27
- Madiha Yousaf, Diedrich Wolter: 2017.Spatial Information Extraction from Natural language Place Description for Incorporating Contextual Variables.GIR 2017: 2:1-2:2
  DOI: https://doi.org/10.1145/3155902.3155910

We have seen that computer programming is an art,
because it applies accumulated knowledge to the world,
because it requires skill and ingenuity, and especially
because it produces objects of beauty.

— Donald E. Knuth [204]

ACKNOWLEDGEMENTS

First of all, I wish to express my greatest gratitude to my supervisor Prof. Dr. Diedrich Wolter for offering me the opportunity of working on Ph.D. research, as well as for his continuous support throughout these years. I'd also like to express my gratitude to him for his patience and encouragement. I began this project with great enthusiasm but little experience. He not only taught me how to conduct proper academic research, which is necessary for becoming a competent researcher, but he also gave me advice and guidance when I was experiencing loss or frustration. Without his assistance, I could not have completed this thesis or grown as a junior researcher.

I'd like to thank my colleague Tobias Schwartz for his help and support, as it was a pleasure to work with you. With special thanks to Prof. Dr. Daniela Nicklas and Prof. Dr. Andreas Henrich from my committee for their interesting, insightful comments and discussions that broaden my research from various perspectives.

I am grateful to my father and sisters for their unwavering financial, mental, and emotional support throughout my life, especially during my Ph.D. research. Words cannot express how grateful I am for their invaluable assistance.

Furthermore, I'd like to dedicate my doctorate to my mother, because without her, I might not have chosen academic research as a means of learning and satisfying my curiosity.

Furthermore, my heartfelt gratitude goes to my friends Ifrah Akhter and Waleeha Shoukat, who supported and encouraged me during difficult times and motivated me to strive for my goal. More importantly, they reminded me that life is more than just academic research. I'd like to thank them for all of the good times we had together.

Last but by no means least, my gratitude goes to my niece (Zobia) and nephew (Ali) who has been always able to make me smile and relax on hard days. Thank you.

Ohana means family.
Family means nobody gets left behind, or forgotten.

— Lilo & Stitch

Dedicated to my sisters for their unconditional love and support.

CONTENTS

1 INTRODUCTION 1
1.1 Importance of grasping the concept of place 1
1.2 Challenges, Tasks and Research Questions 4
1.2.1 Develop a model of context that can support text understanding 4
1.2.2 Develop a method for extracting spatial information that takes into account both implicit and explicit information 5
1.2.3 Develop an approach for recognizing both named and unnamed entities in natural language place descriptions 6
1.2.4 Develop a method to geo-reference named and unnamed entities by considering contextual dependencies 6
1.2.5 Develop an efficient and effective query method . . . 7
1.2.6 Research Questions 7
1.3 SORS: A System for Natural Language Place Descriptions . . 8
1.3.1 Approach and Significance 10
1.3.2 A Computational Model for Incorporating Context . . 11
1.3.3 Spatial Information Extraction and Modeling 11
1.3.4 Reasoning and Geo-referencing 12
1.3.5 Querying 12
1.3.6 Summary 13
1.4 Thesis Structure 13
2 UNDERSTANDING PLACE DESCRIPTIONS 15
2.1 Introduction 15
2.2 What is place and how it is defined and modeled in literature? 16
2.2.1 Existing Definitions of Place 16
2.2.2 Existing Models to Formalize the Concept of Place . . 17
2.2.3 Existing Models to Formalize Concept of Place Considering Context 19
2.2.4 Contrast Sets based on Place Models 20
2.2.5 Towards a model of place 22
2.3 Methods and techniques to identify and retrieve place related information 23
2.3.1 Named Entity Recognition (NERs) 23
2.3.2 Relation Extraction 29
2.3.3 General Limitations of RE Approaches 34
2.3.4 NLP Parsers and Open Issues 35
2.4 Geo-referencing of place 38
2.4.1 Overview of Various TR Approaches 39
2.4.2 Spatial knowledge Representations and Reasoning . . 41
2.5 Research Gaps 45
2.6 Hypothesis for Current Research 47
2.7 Summary 49
3 A COMPUTATIONAL MODEL FOR REPRESENTING CONTEXT AND VAGUENESS OCCURRING IN HUMAN-GENERATED PLACE DESCRIPTIONS 51
3.1 Motivation and Goals 51
3.2 Computational Model 54

3.2.1 Towards an Operational Definition of Context 55
3.2.2 Operationalizing Context Variables 56
3.2.3 Accounting for cognitive aspect of place 57
3.3 Summary 58
4 SPATIAL INFORMATION EXTRACTION 61
4.1 Motivation and Goals 61
4.2 Spatio-ontological Reasoning 62
4.2.1 Processing Pipeline 63
4.2.2 Example:Relation extraction using spatio-ontological reasoning 66
4.3 Discussion and Summary of the Chapter 66
4.3.1 Dealing with implicit and indirect references 67
5 GEO-REFERENCING USING SPATIO-ONTOLOGICAL REASONING TECHNIQUES 71
5.1 Motivation and Goals 71
5.2 SORS: Geo-referencing with spatio-ontological reasoning . . . 73
5.2.1 Overview and processing pipeline 74
5.3 Entity Classification 75
5.3.1 Example 78
5.4 Contextualization 78
5.4.1 Step 1: Resolving unidentified entities 78
5.5 Inference 82
5.6 Discussion and Summary of the Chapter 83
6 QUERY FORMULATION AND STRATEGY 87
6.1 Motivation and Goals 87
6.2 Existing Query Methods 89
6.3 Composing and Scheduling OSM Queries 90
6.3.1 OSM basics and processing pipeline 90
6.3.2 Contextualized Lexicon and Knowledge Base 92
6.3.3 Geometric Interpretation 92
6.3.4 Query Strategy 93
6.3.5 Limitations of presented approach 95
6.4 Identifying Appropriate OSM Key-value Pairs 96
6.4.1 Identifying Adequate OSM Tags using WordNet . . . 97
6.4.2 Semantic Similarity using WordNet 98
6.4.3 Effectiveness of identifying appropriate OSM Key-value Pairs 100
6.4.4 Summary 105
6.5 Determining Entity Types for Querying OSM 106
6.5.1 Problem Description and Approach 107
6.5.2 Inferring Spatial Knowledge from Non-Spatial Context Words 108
6.5.3 Experimental Setup and Data 113
6.5.4 Discussion of Results 114
6.5.5 Summary 116
6.6 Discussion and Summary of the chapter 116
7 EXPERIMENTS AND EVALUATION 119
7.1 Motivation and Goals 119
7.2 Evaluation Methods 121
7.2.1 Evaluation in Information Retrieval 121
7.2.2 Existing Evaluation Methodologies,Criteria and Success Measures 122
7.2.3 Evaluation in Geographic Information Systems 123

7.3 Evaluation Conducted on SORS 124
7.3.1 Experimental setup and data 124
7.3.2 Comparison to geo-taggers 127
7.3.3 Spatial Relevance of Extracted Spatial Triplets 128
7.3.4 Discussion of Results 129
7.4 Summary and Conclusion 132
8 DISCUSSION AND CONCLUSION 135
8.1 Discussion of Major Results 136
8.1.1 Incorporating Context in Automated Interpretation of Text: A Computational Model 136
8.1.2 Spatial Information Extraction 137
8.1.3 Spatio-Ontological Reasoning using Unnamed Entities and Geo-referencing 138
8.1.4 Querying 139
8.2 Summary of Contributions and Evaluation Against Hypotheses 140
8.3 Limitations and Future Work 143
8.3.1 Spatial Relation Modeling 143
8.3.2 Evaluation Data Set 143
8.3.3 Querying Spatial Databases 143
8.4 Applications of a Spatio-Ontological Reasoning System 144

BIBLIOGRAPHY 145

LIST OF FIGURES

Figure 1 Outcome Of Presented Approach SORS 10
Figure 2 Workflow showing the main tasks of this thesis and associated approaches 11
Figure 3 Interpretation of Wikipedia sentence using CoreNLP parser. 36
Figure 4 CoreNLP parser output of a sentence from a travel blog. 36
Figure 5 Excerpt from CoreNLP parser output of a multiple clause sentence from Wikipedia. 36
Figure 6 CoreNLP parser output of a sentence from Wikipedia containing indirect references. 37
Figure 7 An improved georeferencing pipeline with reasoning as Core components 49
Figure 8 Computational model of place based on name, spatial extent, and conceptualisation. 54
Figure 9 Contextual influence in interpretation of place descriptions. 55
Figure 10 Example of processing steps in generation phase of information extraction (NE: named entity, N: noun, R:relation, ID:denotes reference to objects in OSM database) 64
Figure 11 Example of generating and pruning spatial relational statements using spatio-ontological reasoning. 66
Figure 12 Methods for geo-referencing a NL place description. From the set of alternative text interpretations possible, the one which identifies most entities is selected. 74
Figure 13 Illustration of the reasoning-based interpretation procedure which comprises stages information extraction, contextualization, inference, and querying, Extracted information is given as conjunction of alternative relational statements (grouped by parentheses), new statements get inferred and inconsistent statements are pruned off. 75
Figure 14 Inference Steps Showing the Process for Generating Final Queries. 83
Figure 15 Overall processing pipeline for interpreting a place description 92
Figure 16 Constrained specific query categories for scheduling queries 94
Figure 17 Indicating all geo-referenced spatial objects for running example. 95
Figure 18 Nouns in dataset that represent spatial entities, classified into named entities recognised, unnamed spatial concepts, and nouns unknown to the lexicon. Output from text processing is compared to output from the proposed reasoning procedure, arrows indicate effect of the proposed reasoning steps to resolve nouns. . . 130

LIST OF TABLES

Table 1 Overview of Rule-based NER approaches for extracting information related to entities 25
Table 2 Overview of Machine learning NER approaches for extracting information related to entities 26
Table 3 Overview of different Hybrid NER approaches used for extracting information related to entities 27
Table 4 Overview of Limitations of NER approaches based on issues related with entities, data, language, domain and methodology 29
Table 5 Overview of Relation extraction approaches 31
Table 6 Overview of Limitations of RE approaches based on issues related with relationship identification, data,language and technical issues. 34
Table 7 Overview of spatial knowledge representations and reasoning approaches 43
Table 8 Classification of spatial nouns determined by Algorithm 6, using Stanford Core NLP for part-of-speech tagging. 78
Table 9 Resolving unidentified entities as named entities . . . 80
Table 10 Top 5 search branches for the running example, sorted by heuristic (lower is better). Interpretations corresponding to ground truth are underlined. 81
Table 11 Comparison of how the extracted spatial constraints are resolved using reasoning for sentence "St. Catherine is the largest Protestant church, located in city center at the entrance to the Zeil, the central pedestrian street for shopping." 81
Table 12 'Semantic interpretation to OSM query format for Running Example.' 93
Table 13 Values for key "leisure", sorted by their WordNet similarity to "garden" 101
Table 14 Determining tags with *Tagfinder*, position "—" indicates that the correct result is not listed. 102
Table 15 Words with multiple category associations 102
Table 16 Results for finding alternative values within single class using Algorithm 3 103
Table 17 Words with multiple category associations 103
Table 18 Results for finding alternative values within single class using Algorithm 3 104
Table 19 Table showing similarity measure between input and OSM keys 104
Table 20 Results after second evaluation 105
Table 21 Comparison of WordNet similarities for lunch 109
Table 22 Comparison of WordNet similarities for bread 109
Table 23 Annotated sentences with correct OSM tag 114
Table 24 Results 115
Table 25 Results obtained by our approach. Column Δ_{F_1} gives the improvement of performing reasoning steps. . . . 125
Table 26 Results for geo-referencing primary entities. 126

Table 27 Amount of entity types per sentence across the steps of the processing pipeline. 127
Table 29 Results of entities successfully and correctly converted to OSM tags before and after reasoning. 127
Table 30 Comparison of our system SORS with named entity recognition systems 128
Table 31 Results showing the additional spatially relevant triplets extracted by SORS with comparison to NLTK and Spacy. 129
Table 32 Table comparing which entities are recognised by named entity recognition systems 131

ACRONYMS

NL Natural Language

NLP Natural Language Processing

VGI Volunteered Geographic Information

SORS Spatio-Ontological Reasoning System

OSM Open Street Map

GIS Geographic Information System

IE Information Extraction

IR Information Retrieval

GIR Geographic Information Retrieval

GIScience Geographic Information Science

NER Named Entity Recognition

TR Toponym Resolution

RE Relation Extraction

CRF Conditional Random Fields

POS Part-of-Speech

PG Place Graph

QSR Qualitative Spatial Reasoning

SpRL Spatial Role Labelling

1 INTRODUCTION

1.1 IMPORTANCE OF GRASPING THE CONCEPT OF PLACE

Space and place are two fundamental concepts in geography [147]. A formal, unbiased, and reliable view of the geographical universe is called space. Human beings, driven by intuition, refer to space as an informal, subjective, vague, and structure-less world view, called place [82]. The concept of place existed for a long time in philosophy and psychology [108], yet it is relatively new in the domain of GIScience.

The place-based information system concept is an attempt to bridge these two extreme geographical world views. The aim is to integrate human perception of the geographical world into digital systems, allowing place to be formalized and linked to space [281]. Place functions primarily as a spatial pillar in cognitive and communicative activities and has been considered a definitive spatial reference to human, economic, and cultural geography [108]. According to Tuan [361], a place can be seen as a space filled with human meaning and experience, allowing conversations. On the other hand, according to [393], when people are talking about space, they are referring to places.

The relationship between space, place, and language has gained recent attention from researchers, addressing geography, language, and spatial humanities [63]. The driving forces are the staggering amount of unstructured data generated every day, from medical records to social media texts, news articles, and tourism websites; thus, automation will be critical to analyze text and speech data [260],[225]. The growing need for place-related information in everyday life [63] motivates place-based research and is highly acknowledged [146],[147] in the GIScience domain. The primary aim is to advance and improve interactions between humans and computers. Thus, place base systems enable geographical data and information to be collected, stored, modeled, and analyzed from language and cognitive perspectives [264]. As the place is an inherently vague and elusive concept [43], [394] and such vagueness is prominent in human cognition, perception, and natural language place descriptions [9]. It has been postulated that the ability of an automated mapping (geo-referencing) might be leveraged to create a place-based information system for modeling and utilizing human knowledge about the place [394], thereby also realizing ideas of naive geography by [104] which focuses on common-sense reasoning about geographic space.

Natural language is incredibly diverse and complex. It allows us to communicate in a seemingly endless number of ways, both verbally and in writing. It enables us to express abstract feelings, such as how we feel when we saw a particular activity or action or what we like to do at a specific place [225]. The internet has a plethora of written descriptions of places [63]. Such encapsulated data in natural language texts helps users benefit significantly from more affluent and locally valid data on the place in many circumstances [91]. Natural language processing (NLP) techniques enable the extraction of information from textual documents using information extraction (IE) techniques [83] and aid in the interpretation and perception

of the extracted data [225]. Natural language processing (NLP) evolution into natural language understanding (NLU) has a substantial variety of consequences for both companies and customers. NLU is essential because it aids in resolving linguistic ambiguity and provides data with a sound numeric structure. It helps us understand the meaning and nuances of human language in various scenarios and facilitates a computer's ability to make sense of human writings [225].
Although unsupervised and supervised learning and, specifically, deep learning is now commonly used to handle human language, there is still a need for semantic and syntactic comprehension and domain knowledge, which these existing machine learning methods do not provide. This thesis is focused on understanding how humans interpret and use language to develop tools and techniques that enable computers to understand and manipulate natural languages to perform tasks [69].

One might assume that almost all unstructured text is spatial somehow; [422],thus, to analyze and understand the geo-spatial data, natural language processing concepts provide the basis for automation. Building computer programs that understand natural language involves three significant problems: the first issue is related to place conceptualization process, the second to the interpretation and meaning of the linguistic input, and the third to world knowledge [69]. Liddy [239] and Feldman [115]suggest that in order to understand natural languages, it is critical to start at the word level – to identify the morphological structure, nature (such as part-of-speech, etc.) of the words, and then proceed to the sentence level – deciding the word order, grammar, sense of the entire sentence. Then there is the context, and then there is the overall environment or domain. These natural language understanding principles give us a foundation for deciphering and interpreting natural language place descriptions to build a place-based information system for modeling and using human knowledge about a location [107].

Thus, to reason about place effectively, we need to consider the human perspective, which is broad, mainly subjective, and disruptively vague, and any attempt to formalize it should be consistent with reasonable abstract concepts. These concepts make it easier to assert meaning from natural language and are known as context. The notion of place and its definition has been discussed comprehensively [43],[91],[395],[370],yet there is no definition capturing human intuition, leading to a fully specified computational data model.

Place descriptions have been regarded as qualitative reference systems that describe geographic locations that comprise references to places (like place-names) and qualitative spatial relations describing their relative whereabouts [371]. Current geographic information systems are based on explicit, crisp, and metric geometries, which are different from human concepts [370]. Discrepancies thus arise between how humans conceptualize space and how entities are represented in GIS.

Existing geographical information systems like web mapping, navigation, and location-based systems can provide us with spatial locations despite being vague as they focus on official place names to represent places and no contextual information. Understanding geographical information in natural language or free text presents many challenges [188], like a single entity in GIS may correspond to several places. Consider a query " the post office near the train station in Bamberg," It consists of three essential parts, a type entity (post office, train station), a spatial relationship (near, in), and a location (Bamberg). As is typical for many queries in information retrieval, it is under-specified and ambiguous. In this sentence, named entity Bamberg can be a town in Germany or a district in America named alike. Moreover, what does near mean in this context? Are we talking about the closest post office from a specific point (train station) or representing all the post offices present in a certain constrained space? In addition, train stations and post offices have potential unknown spatial whereabouts.

Thus indicating that current information systems are not capable of interpreting or reasoning with spatial relations and unnamed spatial entities such as train station or post office. The answers for place-based searches are provided by existing systems ignoring spatial objects' semantics (type nouns) and relations. As both are
context-dependent, and how humans would perceive them will result in different outputs. However, dealing with spatial entities semantics requires more information like reference directions, routes, place conceptualizations, and human activities associated with places to model the semantics better. Existing information techniques cannot extract such deep information and require manual work, driving our focus towards the spatial objects and their semantics [63].

Existing literature indicates that all existing systems are largely confined to place names (named spatial entities), they do not consider a place type like the market or any other standard form, e.g., the picnic spot. References that are not place names are more difficult to locate and typically necessitate considering contextual information [347]. Thus working well on documents or place descriptions that include place mentions for geoparsers but leaves open the case of descriptions with none or very few place mentions.

Another challenge related to automated text interpretation is related to language parsing and human conceptualization. Formulating the query or search-based inquiry information extraction from NL input is usually carried out using dependency parsers and named entity recognition systems. Parsing techniques appear reasonable where the syntactic structure of text description is supported by the parser (i.e., acyclic, linear order, and semantic heads are verbs). Moreover, parsers can extract explicitly stated information in the sentence but are unable to resolve implicit dependencies between spatial objects and relations reliably unless the sentence structure is supported. As we are dealing with Volunteered Geographic Information where NL comes in various forms of text and structure, the input sentence can not always be in the same structure.

These barriers may frustrate people when they communicate with information systems such as emergency services, web searches, travel planners, or navigation. From the above state of the art, one can conclude that all these obstacles are due to the lack of human perspective (like how humans talk about the place) while building the information systems and performing different functions on them.

Considering these challenges, the motivation is to develop a place-based information system that can capture and incorporate human-based perspective while extracting information, modeling, reasoning, and querying it. As human cognition is context-dependent, defining and capturing context within such a framework is also necessary.

With a framework designed from a people-based perspective, it would be possible to explore how this structure could enhance existing services of representing and reasoning spatial descriptions. Automated interpretation w.r.t. human knowledge and geo-referencing places on the map will undoubtedly improve the emergency and navigation systems. Besides, it will help us resolve ambiguities and improve spatial language understanding by adding context to it. Moreover, it will make place search systems more flexible instead of just named entities, and we can perform a search on type nouns and improve the query mechanism. Also, it will ease the interaction between humans and computers by providing a spatially intelligent system.

1.2 CHALLENGES, TASKS AND RESEARCH QUESTIONS

This thesis identifies five significant tasks in line with the challenges mentioned: incorporating context-sensitive information (a computational model into a place model), spatial information extraction from text and modeling, reasoning, place geo-referencing, and querying. There are several challenges relevant to each of these activities, which are discussed in this section below.

1.2.1 *Develop a model of context that can support text understanding*

Places are cognitive concepts when people started naming them, telling stories about them, and indicating what can be done at them [9]. Within the representation the main goal is to capture and model the concept of place from a cognitive perspective. One practical approach is to use the surrounding context to understand human perception [188]. Souza et al. [342] define four types of context in the geo-spatial domain: user, data, procedure, and association, while Brodaric [55] uses dimensions to reflect context, which includes origins, uses, and effects, as well as entities with specific roles. Zhong et al. [415], differentiate between visual (low-level) and semantic (high-level) context, with the latter including spatial context (which refers to the relationships between various types of geographic features such as relative position and co-occurrence) and scene context (referring to the location and situation of an object within a scene). Several studies have focussed on identifying specific contextual factors that are believed to influence how location expressions are used and perceived [347], including perspective [174],[351],[135]; frame of reference [224],[233],[71],[160]; force dynamics [135]; boundedness [351],[420]; domain [200], spatial preposition meanings based on salience,relevance, tolerance, and typicality [161] or generating new scenes [221],[185],[290] and indoor/outdoor/transitional space [210]. Stock and Hall [346] provide a summary of these factors.
The work described above is mainly concerned with prepositions, with some work focusing on verbs and adverbs, although in far less detail. The prior literature looked at the context in a more restricted sense. It did not consider several contextual factors that could impact spatial language interpretation, including the factors related to spatial entities (locatum and relatum being the most important category of objects providing context) [347]. Thus, contextual considerations will increase the efficacy of automated geo-spatial language interpretation by considering both linguistic and geographic aspects [347].

Thus, before embarking on interpretations of place descriptions, we have to commit to a definition of place by considering the concept of context. In order to define a place by the human mind and thought process, the first challenge is to understand and identify the set of contextual factors the elements influencing what the most reasonable interpretation of place description is.

Once the factors have been identified, the next task is to ensure our computational model must handle partial satisfiability of a relation and a concept and determine the most plausible interpretation by maximizing satisfaction. For example "park near the city center" should result in the closest location to the city center, while it maybe clear to disambiguate candidate parks to the city center, there is no clear semantics of "near" it remains an uncertain and vague concept. Another challenge is that our model needs to include variables representing conceptual entities and the rules that determine the mapping. The fundamental challenge is to identify those contextual factors that can work for both spatial entities and relations, providing the required information for interpretation.

1.2.2 *Develop a method for extracting spatial information that takes into account both implicit and explicit information*

The rapid development of text mining and NLP techniques makes it feasible to extract information from textual documents through information extraction techniques [83] such as name entity recognition (NER) [271], entity relation extraction [87], and spatial relationship extraction [23],[207],[244]. However, even with these tools, automated interpretation of the text is still a challenging task mainly because of language parsing and ambiguity of names and lack of models that adequately capture vague place description's semantics. The biggest challenge with the existing parser is that they provide us with limited information, the one that is explicitly stated. For example, Bamberg is a town in upper Franconia, Germany on river regnitz, the relations extracted by the parser are town in upper Franconia and Germany on the river. However, curcial relations like "Bamberg in Upper Franconia" or "Bamberg in Germany " is missing. As we are trying to incorporate context in our model, the implicit relations between the objects are pretty important, and the parser is unable to provide a densely connected set of facts that would make it effective.

In addition, for the parser to work effectively, the sentence's syntactic structure must be supported by it. As we are dealing with Volunteered Geographic Information, it is impossible to assume all the sentences in a well defined grammatically correct structure. Thus if the input sentence has a cycle or the order is not linear, the output might not be the one we expected. Moreover, concerned with spatial data, the relationships are formed in the form of triplets based on spatial relation present between the objects, while for the parser, we have a semantic head, i.e., verb, and if it is not present, the sentence fragments or objects are not connected.

Another issue arises while dealing with spatial relations semantics, requiring more information like reference directions, routes, place conceptualizations, and human activities associated with places to model the semantics better. The existing information techniques cannot extract such deep information and require manual work, driving our focus towards the spatial objects and their semantics [63]. Besides, graphs of crisp relations are not sufficient as means to represent uncertainty are required.

Once the information has been extracted, the input text is processed to form a set of declarative statements in the form of triplets like rel(place ref1; place ref2), meaning place1 and place2 are related by some spatial relation (in, near, etc.) as presented by [373]. As many researchers [41],[198],[199],[373] and [384] prefer graph-based representation, which seems naturally more suitable for representing network of interrelated place than a tabular-based database, yet exploiting context information presents a chicken-and-egg problem, though: information obtained by resolving entities is to be employed simultaneously to resolving the entities.

This motivates an approach using logic programming techniques since dependencies can be expressed in a declarative manner, abstracting from algorithmic realization as it allows us to abstract from a pre-defined order of processing.

Thus,there is a need for an approach which can deal with both implicit and explicit information present within the input sentence and can extract maximum possible information about the spatial entities and relations between them.

1.2.3 *Develop an approach for recognizing both named and unnamed entities in natural language place descriptions*

This section's focus is the role reasoning can play in interpretation of information extracted from language through reasoning.Reasoning enables us to overcome limitations in previous geo-referencing approaches, i.e., using named entity recognition (NER) systems, insufficient exploitation of context in named entity disambiguation, and relying on parsers [188]. Thus, we must develop means to assemble the pieces of information into a coherent whole. Reasoning would allows us to detect some inconsistent interpretation candidates using ontological reasoning. The idea of reasoning is similar to how humans can argue based on mental models for such scenarios [202]. The challenge for qualitative reasoning is that these models have not been used for reasoning with contextual information based on collective human knowledge. The information will be extracted into two alternative sets: relational statements and the qualitative reference to place names, vernacular places, place types, and relations. This type of knowledge can be used for qualitative spatial reasoning tasks such as inferring new entities based on existing information. In the contextualization stage, new entities are determined based on the concept of abductive reasoning, which is an inference to the best explanation by creating and testing all possible hypotheses available.

Thus, to generate all plausible interpretations based on contextual factors for unnamed and vernacular places, reasoning techniques are required to be flexible and preferable for maintaining the knowledge base's consistency.
By developing a reasoning framework for text interpretation, we aim to integrate methods focusing on NER and parsers' sub-tasks and combine their strengths.

1.2.4 *Develop a method to geo-reference named and unnamed entities by considering contextual dependencies*

Recognizing place names from texts and linking them to spatial footprints are essential steps for automatically understanding the semantics of natural language texts and are studied both in computer science and GIS [172]. Place descriptions come in various forms; they may consist of a single name or referring expression, a network of interrelated expressions, or even route instructions that guide an imaginary wanderer to a target place. Current techniques for Geo-referencing are based on toponym resolution, which relies on place name matching in gazetteers [228].

The location details often contain references to fine-grained locations, such as streets, houses, and local interest points, which are often considerably more frequent and more similar, thereby requiring a different resolution method. In particular, place definitions are versatile, vernacular, and often include references to locations that can not be contained in gazetteers, such as synonyms or forms of places (e.g., the city center). Once presented with such parallels, traditional strategies easily fail. Also, place references may refer to places that are not gazetteered at all, such as places from environments that are too fine-grained or conversational contexts that are too limited to be captured by gazetteers (e.g., the old city). Non-spatial place information such as place semantics, equipment, characteristics, and affordance can also help applications such as place searching and querying.

Considering the data present in a text description provide contextual knowledge, leads to the idea that implicit information present in the description can be exploited for tackling entity disambiguation problem. We consider the problem of identifying the geographic focus of text that does or does not explicitly mention the target country, making our problem one of inference or prediction rather than one of identification.

Therefore, although it is usually easy for a definition receiver to disambiguate location references or roughly locate locations by spatial types, computational modeling of NL spatial entities to enable automatic interpretation remains a significant and open challenge. Thus, once the unnamed entities are resolved using contextual knowledge, higher geo referencing precision is expected.

1.2.5 *Develop an efficient and effective query method*

As a knowledge base with collective information on locations that people have found worth explaining, the framework should also help the query of stored knowledge. First of all, the query results are the basis for applying spatial reasoning to validate or reject some interpretation. The corresponding challenges related to this last phase include composing the correct query format concerning the geographical database we are considering *Open street Map*. Also, the second task is the scheduling of queries to the database. Moreover, the process of reasoning can facilitate the query generation stage by applying inference concepts. In the inference stage, derived information is propagated through deduction. For For example, if some entity A is known to be located inside another entity B, which itself is contained in C, constraint propagation techniques allow us to arrive at the facts that A is also located inside C and that B and C must both be region-like entities.

In the case of country name present or no named entity, the query's scheduling is affected by the context. More importantly, the query approach should yield a result that resembles to the gist of given place description.
Consequently, it is necessary to develop a contextualized query approach for the system to be developed, which can schedule the order of queries as well as decide whether a query should be submitted or the number of results expected is too large to be manageable.

1.2.6 *Research Questions*

According to the identified major goals and challenges, the major research questions of this thesis are:

- To what extent can and should context be incorporated into a model of place to allow for sensible interpretation of natural language place descriptions? What will be the influence factors respectively that can be considered to integrate different statements? (expressed in our spatial representation)
- To what extent can a qualitative spatial representation capture context-sensitive vague place descriptions? Have appropriate models been developed so far, or is it possible to devise such models?
- Does reasoning about spatial and ontological properties aid for information extraction and how efficient and effective it is in addressing problems with natural language parsing?
- To which extent do spatial and ontological reasoning model help resolve ambiguous interpretations, in particular regarding the interpretation of unnamed entities, and how effective is this reasoning model for geo-referencing?
- What will be the query semantics for handling context-sensitive vague place descriptions and relations? How efficiently and effectively can reasoning be performed on the geographic database using such a model?

Answers to these research questions will contribute to geographical information retrieval and interpretation by clarifying the benefits of reasoning methods applied as an additional layer in understanding natural language place descriptions. Additionally, results improve information retrieval at an initial stage and open doors for automation of place-based systems.

My research aims to fill an existing gap by improving the reasoning process needed to resolve indirect referencing, expanding traditional co-reference resolution techniques, and dealing with the instances of geographic entities referred by their type. By incorporating context, this research helps in dealing with vagueness and ambiguity in spatial entities and opens the doors for dealing with unstructured data. This study's reasoning methods help us understand human search behavior and how humans search for stimuli and memory patterns. Besides, this study provides us with a simple named entity recognition system based on the type of entities that allow named entity disambiguation at the information retrieval stage.

This research aims to provide a practical approach for querying the geographical database (OSM), improving the geo-referencing quality and spatial data representation on the map.

1.3 SORS: A SYSTEM FOR NATURAL LANGUAGE PLACE DESCRIPTIONS

This Ph.D. research aims to develop an automated system SORS (spatio-ontological reasoning system) capable of capturing context-sensitive vague spatial information in human-generated place descriptions and allowing geographic databases to be queried using such descriptions. Natural language place description reflects the way people mentally represent and communicate about the place and are therefore considered target sources for this thesis. This study aims to contribute to several research areas to define and mitigate knowledge gaps, primarily capturing context information, Information extraction, and location knowledge exploitation derived from descriptions.

The outcome of this thesis is multifold. The outcomes include these novel scientific contributions, which are,

- a computational model for representing vague spatial knowledge and context occurring in human-generated place description;
- a spatio-ontological reasoning method for information extraction;
- reasoning methods for interpreting unnamed spatial entities;
- automating process to answer Open street Map queries;
- automating query techniques to access geographic database (OSM);
- an approach to deal with query problems including non-spatial data using semantic similarity and clustering and rankling algorithms

The second outcome is implementing the system, which allows capturing contextual information from place descriptions and then processing that data by extracting, reasoning, and querying it. It's a tool for further experiments in the field too.

Natural language place descriptions are common in everyday communication and could thus serve as a rich source of volunteered geographic information (VGI) [260]. This idea for basic research in spatial representation and reasoning is motivated by the observation that VGI is not restricted to geo-referenced data collected by volunteers for the explicit purpose of contributing to a geographic information system. However, there also exists a virtue of implicit spatial information contained in place descriptions. New sources of information are available to applications and research based on crowd-sourced geographic information by opening up these channels. Enabling computer systems to benefit from the latter category of implicit information is particularly relevant since text-based communication is a very natural form of information exchange, particularly considering social media platforms, Tweets, or the Wikipedia encyclopedia. As a consequence, there exist manifoldly pieces of information that can serve several applications, for example, disaster relief, urban planning, and emergency service, and many more, and it may also serve as a tool to gain new insights into social aspects of utilization and conception of spatial environments.

Below is an example description of a real-world data-set, which communicates information about a district and many places present there, including the affordances related to these places, and different directional relations are provided showing how the places are linked.

"St. Pauli is a colossal district in Hamburg. If you want to go for a drink is Hamburger Berg. One of my favorite bars is Lunacy, and I can also recommend zum Silbersack just down the street. Other right places include Drei Zimmer Wohnung and Wilde Hilde bar across the neighborhood."

As the information present in the given input has multiple sentences and is ambiguous as well as contains multiple named (St.pauli, Hamburg, Lunacy and so on) and unnamed entites(street, district, and bar). Yet, the existing named entity recognition systems are only able to recognize Hamburg as a named entity and all others are identified as a person's name or remained unrecognized . Moreover, as it is a spatial description and location names, an inference is required about their spatial extent. As a new type of data source unstructured and unsupported by current GISs and spatial databases, it motivates us to develop a context-sensitive geographic information system.

Thus, to resolve these ambiguous spatial objects and relations and deal with implicit information, we need techniques to be identified that capture the semantics of spatial objects in vague place descriptions that allow spatial reasoning to make the information explicit, considering both the context of the description well as geographic context. Finally, query techniques have to be devised that can efficiently assess the geographic context and match place descriptions to geographic entities. The output of the system presented in this thesis can be seen in the Figure 1 below.

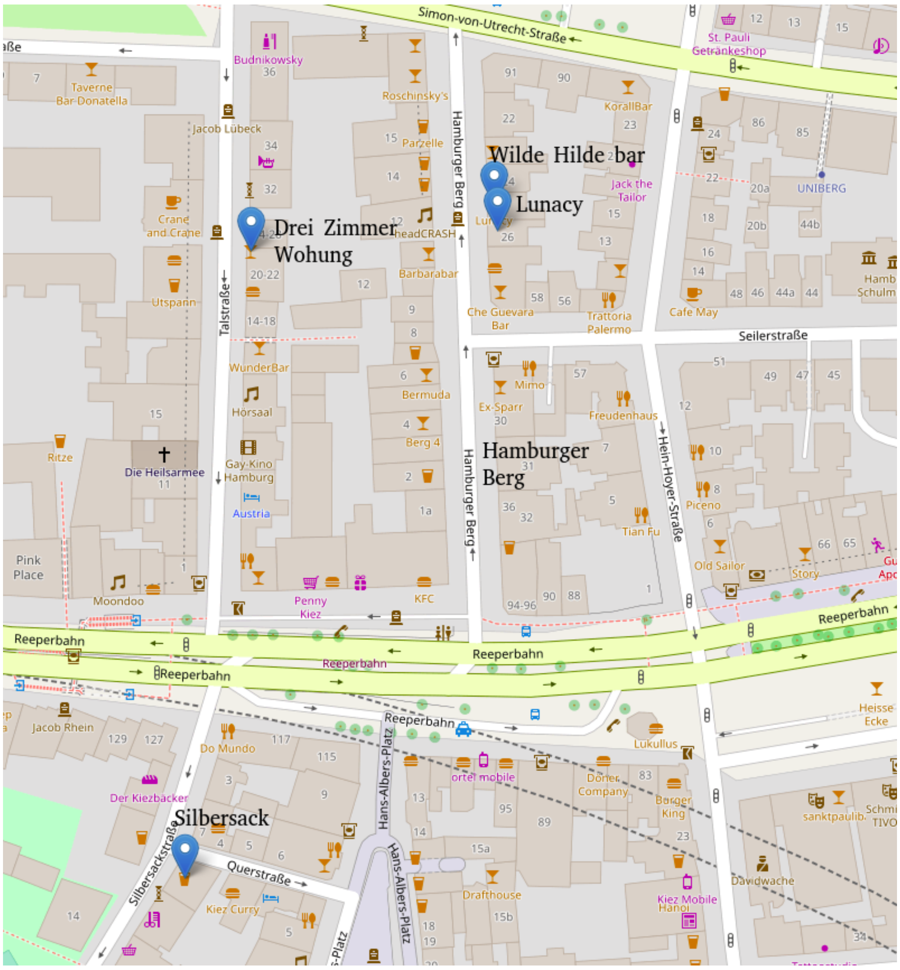

Figure 1: Outcome Of Presented Approach SORS

The idea is to delay the semantic interpretation of the relations until further constraints are obtained from a fact database. Once the entities have been mapped, we can depict relations easily.
The basic outcome of this thesis is a reasoning framework with scientific contributions include approaches proposed for the identified five major tasks from the last Section 1.2.

1.3.1 *Approach and Significance*

This thesis contributes towards identifying and incorporating context, information extraction, modeling, reasoning, and geo-referencing and querying knowledge extracted from NL place descriptions. The figure 7 below shows the work flow of this research.

All the tasks are sequentially interlinked with one another and are addressed in this work. The first phase of our dissertation is theoretical research to identify contextual factors and help define a place. After a place

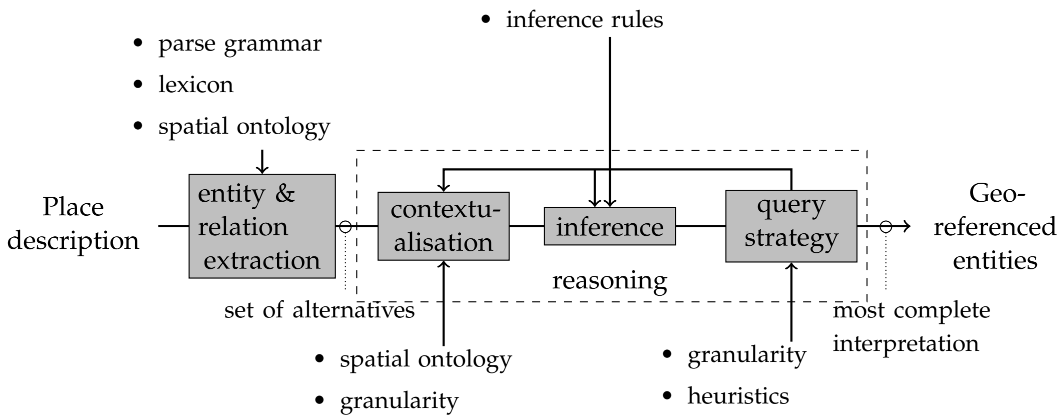

Figure 2: Workflow showing the main tasks of this thesis and associated approaches

has been defined and factors are identified, they help in the next phase of entity and relation extraction (information extraction), resolving entities based on context (contextualization), and inferring the relations to generate a set of queries (query strategy). The above presented Figure 7 highlights four significant modules, which are based on a spatial ontology, the lexicon of entity names and types, and parse grammar. The details of each module have been presented below.

1.3.2 *A Computational Model for Incorporating Context*

The first phase focuses on developing a computational model to represent vague spatial knowledge and context occurring in human-generated place descriptions. We presented computationally motivated models of place and context before integrating them into one model. We argued that places could obtain their meaning from actions that can be performed within their spatial extents and therefore a definition of space should include the conceptualization, i.e., the process which associates a particular concept with a spatial area. Thus, a place model consists of three basic elements: place name, spatial extent (i.e., boundaries), and conceptualization. Conceptualization comprises three contextual variables: environment (physical domain), human and place description.

This model allows us to tackle place description's interpretation as a constrained optimization problem using discrete and continuous variables.

1.3.3 *Spatial Information Extraction and Modeling*

The second phase focuses on information extraction of place descriptions. The already existing off-the-shelf information extraction (IE) tools like a parser, named entity recognition system (NER) are available, yet the information extracted from them is limited, ambiguous, and lacks human conceptualization. Here we presented a reasoning framework based on an exhaustive search algorithm that generates all possible relational statements based on a combination of spatial and ontological reasoning. The pruning phase eliminates the incorrect facts using a novel combination of database lookups and lightweight ontological reasoning.

Making spatial-ontological reasoning an explicit step in the interpretation also enables consideration of contextual dependencies. Clearly, an exhaustive search does not tackle the fundamental problem of language understanding, but it relies on the assumption that the most extensive set of statements that can be matched to a geographic database corresponds to the intended interpretation.

Our reasoning algorithm can prune off most invalid interpretations, whereas natural language parsing results in some wrong commitments one cannot recognize in a later processing step.

1.3.4 *Reasoning and Geo-referencing*

The third phase focuses on geo-referencing unnamed spatial entities in texts that exploit constraint-based reasoning known from logic programming and qualitative spatial reasoning with composition tables. In an iterative manner, the idea is to carry out a shallow semantic interpretation of named and unnamed entities in a text (which does not yet fix semantic relations). Then, contextualize the entities with possible interpretations (co-reference, spatial, and is-a relations) as dis-junctions that are pruned using constraint reasoning, and then formulate and run corresponding spatial queries over OSM that identify unnamed entities, providing further constraints. The search looped towards an interpretation that maximizes the number of geo-references in a sentence. The method is evaluated on automated geo-referenced example texts and shows that the reasoning component can considerably improve the interpretation of unnamed entities.

1.3.5 *Querying*

The last phase of our research is related to query generation and scheduling. As we use OSM as a geographical database, we provide a natural language interface that can be used to query OSM without using any API or ontology directly. As a query engine, we use oscar [31]. Since oscar enables us to query for key-value pairs like amenity=theatre restricted to local surroundings, we trigger a query whenever either the name or type of an entity and its spatial containment are known, i.e., given as relational expressions (e.g., William street in/of Melbourne or st. Catherine at the entrance). In the case of entities known to be countries, no further context is necessary. We refer to these kinds of queries as constrained queries.

We note that so far, we have only realized a simple named entity recognizer that introduces the option to interpret every capitalized word or phrase as a named entity. A geographic database is then consulted for potential matches related to entity type extract appropriate key-value pairs during interpretation.

In addition, this section provides us with a ranking method based on semantic similarity, which generates a set of appropriate OSM tags to query entities described in natural language descriptions. A clustering and pruning algorithm is also presented to generate semantically replaceable type nouns for non-spatial nouns present in the description, which helps in improving the geo-referencing process.

All these components are developed using python and the NLTK POS tagging library. The system is based on a lexicon made on existing OSM tags, and the implementation is in python language making it easy to access and use.

1.3.6 *Summary*

The major contributions of this thesis are summarized as follows.

- A computational model has been presented which can capture vague and context-sensitive information about places. This model allows tackling the interpretations as a constrained optimizing problem using discrete and continuous variables.
- A reasoning framework for spatial information extraction is presented to progress towards the long-standing goal of solving spatial language parsing. The framework uses spatio-ontological reasoning for generating all plausible parse options and then keep the valid ones by pruning. The framework allows the consideration of contextual dependencies.
- A new method to automatically interpret natural language place descriptions by using spatio-ontological reasoning techniques, enabling us to explicate implied information in the form of Unnamed entities embedded in place descriptions. Reasoning allows the generation of new information across three components: contextualization (incorporating type information), ontological reasoning (is-a relation), and inference to generate map-able data.
- A query strategy is presented to schedule queries to the underlying geographic database and supplement the set of relational statements accordingly. A strategy becomes necessary to avoid queries that lead to more matches in the database that can be handled.
- A ranking method based on semantic similarity is presented, which generates many suitable OSM tags to query entities described in natural language.
- In addition, a clustering and pruning algorithm generates semantically replaceable noun types for the spatial non-nouns in the description, improving the geo-reference process.
- A natural language user interface is provided where a text description is provided as an input, and the geo-referenced entities are mapped on OSM. The implementation integrates the above approaches into a complete processing chain, together with several other functionalities such as database visualization and place mapping.

1.4 THESIS STRUCTURE

Below is an overview of the remaining chapters in this thesis.

Chapter 2 gives an overview of prior efforts on how location is described and represented in current literature. In addition, it includes a description of previously suggested methodologies for identifying and extracting geographical information from textual documents and identifies and summarizes the existing research gaps.

Chapter 3 presents the proposed computational model representing spatial knowledge occurring in place descriptions and analyzes how context information shapes a place description's meaning. It outlines contextual factors that represent vague spatial knowledge and context occurring in human-generated place descriptions and will help interpret information in later tasks. This chapter addresses the first and second hypotheses.

Chapter 4 proposed a reasoning technique that focuses on grasping and interpreting the meaning of place descriptions by generating an orderly, clear portrayal of entities and their relationships with one another. The flexible reasoning approach(SORS) uses logical statements as an intermediary representation to over-generalize knowledge supplied in a given text location description. The results show that the method is reliable and resilient in dealing with implicit information and generating other geographically relevant triplets. This chapter addresses the third hypothesis.

Chapter 5 describes a SORS strategy for geo-referencing all places in natural language place descriptions, regardless of whether gazetteer names link to them. The method categorizes places into three types: named, unnamed, and unidentified, and uses abductive and deductive reasoning approaches to geo-reference them. The methodology achieves improved geo-referencing precision and recall as compared to standard NER methods. The fourth hypothesis is addressed in this chapter.

Chapter 6 initially presents a pipeline for query composition, explores the types of queries that OSM may answer, and presents a query strategy to schedule the queries. Furthermore, this chapter entails the challenges that arise while composing queries, resulting in inaccurate or no results, and employs semantic similarity to increase the query process efficacy. Finally, the chapter discusses dealing with non-spatial data and how semantic similarity might help us locate comparable concepts in the form of unnamed nouns. This chapter addresses the fifth and sixth hypotheses.

Chapter 7 focuses on identifying the impact of reasoning techniques presented in this research study. The chapter explains the various experiments conducted on SORS, provides details about the experimental setup and evaluation data and discusses how the system performs and contributes to the existing research gaps.

Chapter 8 comprehensively discusses proposed methodologies and their experimental results, relating the findings to the hypotheses. The chapter finishes with a summary of the thesis contributions, as well as suggestions for future work.

2 UNDERSTANDING PLACE DESCRIPTIONS

Based on the significant issues identified in the preceding chapter, this literature study outlines the background with related studies and highlights their relationships to this dissertation. In Section 2.2, a discussion of how the place is perceived and described in the literature. Section 2.2.3 then examines current models for a place from an information system standpoint and identifies existing limitations and research gaps. Section 2.3 discusses previously proposed approaches and methods for identifying and extracting appropriate place-related information. Named entity recognition systems, relation extraction, and NL parsers are all covered in this chapter, along with their constraints. Section examines studies on geo-referencing locations from the text. A summary of research gaps in the literature is provided in Section 2.5, which would be the focus of this thesis.
Lastly, Section 2.6 provides a set of hypotheses to cater the existing gaps.

2.1 INTRODUCTION

Volunteered geographic information (VGI) can be regarded as encompassing a broad spectrum of sources for information [399]. Data is being generated at a faster pace than ever before by people and machines. This increasing volume and variety of data being generated [270] make it more challenging to extract useful information [8]. Scientists look for meaningful trends in increasingly large datasets by gaining understanding by search.

Many fundamental problems in computer science, according to Papadimitriou [284], are search problems, described as "given an input, call it x, find a solution y such that x and y stand in a particular relation to each other that is easy to check" [284]. Many of these fundamental search problems include discovering paths in structured spaces (e.g., shortest path, Hamilton path, Clique, and the Mini-cut) [34]. Thus, the search for information has been intricately linked to a spatial dimension [357]. Due to the large volume and variety of data (i.e., structured, semi-structured, and unstructured), the search process has become more challenging. The potential opportunities and solutions are obstructed by data's unstructured nature [8], that is (1) no schema [383], [249] (2) multiple formats, [383], [251] (3) they are derived from diverse sources [383], [251], and (4) there is no standardization [383], [249], i.e., various representations. Unstructured data analytics work better when the data is transformed correctly. It is difficult to derive useful information from mixed data because of its complicated heterogeneity [8].
The importance of location in many tasks related to unstructured text search and retrieval has become increasingly apparent, as not only commercial search systems have scrambled to integrate geographic information, but research efforts have also emphasized its importance [188]. Despite this apparent growing acceptance, geography's significance in search is often reduced to location, which is only one of many possible features [188]. In [15], Allan stated that the focus of search should be instead on what is not mentioned as VGI information contains both purposefully collected geo-referenced data to implicit geographical references, for example, textual descriptions of geographic entities like "a park between the arms of the River Regnitz, near

the city Bamberg." The example demonstrates the need for vague interpretation to single out the parks since any park on the island of central Bamberg would qualify as "between two arms of a river." Besides, interpretations for "near" in the motivating example requires context. In summary, place descriptions occurring in crowd-sourced channels offer volunteered geographic information but only implicitly and using vague concepts [399]. Consequently, there are manifoldly pieces of information that can serve several applications, and it may also serve as a tool to gain new insights into social aspects of utilization and conception of spatial environments. Before pieces of implicit geographic information in text can be exploited, geographic information content must be made explicit and interpreted with respect to existing geographical knowledge – *this puts a focus on spatial knowledge understanding, representation, reasoning, and querying techniques*.

Standardizing spatial language tasks appears to be more difficult because it is more difficult to achieve an agreeable set of concepts and relationships and a domain-independent formal spatial interpretation representation [293],[208], [255],[346],[90]. Hence, we explore the spatial dimensions of and approaches to information search [22]. In this manner, the objective of this chapter is twofold: first, to explore the state-of-the-art techniques for interpreting VGI text place descriptions and to identify the limitations of existing techniques for multi-faceted unstructured data. The second is to develop different hypotheses to provide the solution for existing limitations.

The remainder of this chapter is structured as follows. Sections 2.2,2.3 and 2.4 provides an overview of the cutting edge research on understanding and interpreting, representing, reasoning, and querying with spatial information present in text descriptions. Section 2.5 discusses the limitations and highlights the research areas that require more attention in the field of interpreting crowd-sourced place descriptions.Section 2.6 will provide us with the set of hypotheses that will form the basis of this research. Section 2.7 summarizes and concludes the chapter.

2.2 WHAT IS PLACE AND HOW IT IS DEFINED AND MODELED IN LITERATURE?

This section reviews the definitions and concepts of place in establishing a conceptual framework for automatic interpretation of place descriptions in light of how humans perceive place and how it has been used in past years to improve spatial human-computer interactions and spatial reasoning. Furthermore, the current models and techniques used for representing spatial knowledge are also discussed.

2.2.1 *Existing Definitions of Place*

Since the 1970s, *place* has been described as a specific location that has acquired a set of meanings and attachments [85]. The academic definition of place is a research problem that has fascinated scientists for several years [281]. The place is considered a fundamental concept in communication and spatial cognition, closely associated with human experience [82],[361]. Relph [328] characterizes place as unique physical features of patterns, appearances, events, activities, and functions. Its unique quality is the ability to focus on human experiences, actions, and intentions in the spatial dimension. Curry [88] defines the place as a notion free from natural boundaries. More precisely, space existed long before humans were there, while a location is shaped and developed by a human mind. As a consequence, places are human creations designed to describe space.

Davies et al. [91] conclude that the idea of place is, at best, a daunting challenge for GIScience. Places appear to be almost defined as intangible in the literature, both spatially and semantically. The concept of the place is psychological as well as physical. The combination of physical form, activity, and context constitutes the sense of place [269]. According to Kuhn [214], space and place are linked to the ten core concepts (i.e., location, field, object, network, and event in the latest form) of presented spatial information. A place is a combination of materiality, meaning, and practice [85].

In [11], Agnew suggests thinking of a place related to other places instead of being bound and having isolated features. According to some researches, the concept of place is too vague to formalize, except in limited circumstances [9]. Other researchers investigate what features make up a place [369], while some identify the difficulties determining such features [9]. Tuan [361] suggested that places were essentially "centers of meanings developed out of life experiences." The place is usually defined by textual place names associated with coordinate locations without considering people's perception and cognition factors [63]. Overall, place emerges as a cognitive concept in several ways, when people name them, tell stories about them, and do things there [9].

In an attempt to improve place-based research in geographic information science(GIScience) and VGI approaches that aim to interpret place descriptions, it is necessary to capture and model the concept of place from a cognitive perspective to utilize place-related knowledge in an information system [63]. The next section 2.2.2 provides a review of existing work on attempting to capture the notion of place from an information system perspective.

2.2.2 *Existing Models to Formalize the Concept of Place*

This section lists the most commonly known approaches for formalizing the notion of place while facilitating its qualitative analysis and integration inside the GIS [282]. The straightforward models like gazetteers and toponym databases [144] dominate spatial information systems and GIS. gazetteers store names and alternates, place types in some more or less ad-hoc taxonomy, and place geo-references to polygons or points in a coordinate system. The way that a gazetteer maps space into locations, although helpful and practical, yet does not fit the way humans perceive their world [63]. With the rise of the idea of the Semantic Web [46], techniques for designing and implementing geospatial semantic web [103],[119],[180] have evolved rapidly.

Ontologies [117] are among the most common means of representing and organizing information, varying from general to specific domains. In the case of place, many developed instances may be specialized exclusively on the topic, such as Geonames [1] or general instances that define the place part of the upper-level ontology, such as schema [2], which states that *'Entities have a somewhat fixed, physical extension'*.

One of the research directions is to model semantics behind names in gazetteers. Studies have been performed to enrich gazetteers considering the on-

1 http://www.geonames.org/ontology
2 https://schema.org/Place

tologies related to formal place semantics [179],[178]. Ontological Gazetteers [179] expand the more straightforward digital gazetteers format [144] with knowledge graphs that incorporate thematic information on places of interest such as part-time relationships, temporal information, and other semantics that outline the meaning of place being considered. Several works depend on the simplicity of the formalization of the place introduced by the gazetteers. In [374], they concentrate on organizing geographical regions identified by location names based on their spatial and thematic significance. They also expand the graph-based organization of similar location names by adding new ones related to existing locations by containment or partitioning relations. An alternative option is to extend place names with aspects of spatial cognition, such as the principle of contrast: "People conceptualize a portion of space as a location if it shows wholeness against its environment" [394]. It is argued that these place representations, including enriched semantics, aim to enhance tasks like information retrieval and query expansion [128],[190]. In [43], they argue that the nature of place may be derived from an analysis of NL's semantic content.
Thus, it seems promising to extract place semantics, e.g., based on place type keywords, human activities, and events from unstructured, NL descriptions. Research in [13] focuses on mining methods for activities and services generally associated with places by creating a schema that associates types of a place with the types of services provided, based on the co-occurrence of language patterns derived from geographically-focused resources. Other studies investigate language models to semantically describe locations through word-embeddings [5],[105],
[344]. Nevertheless, most of these models mentioned above do not fully utilize the expressive power of place and cannot capture the notion of place as a cognitive aspect. The several limitations associated with the above-mentioned models are discussed below.

2.2.2.1 *Challenges and Limitations*

1. Gazetteers have proved to be well suited to many researchers working with spatial data, mainly where data representation is either an environmental or administrative phenomenon. However, they fall short of reflecting the intuitive geographical definition of the place, which plays a vital role in human experience and communication on a daily basis. This shortcoming complicates the integration of volunteered geographic information with authoritative sources and human-computer interaction [291].

2. Although ontological approaches include specifics that promote a better understanding of the place, they provide spatial information that is either ambiguous, qualitative, and linked to other locations, including relative locations and place names, or does not relate to human perception of place [282].

3. Above mentioned models for formalizing place are only able to cater part-of-relationships. They are unable to account for the adjacency and the proximity of places [291].

4. These methods relate space with basic semantics in the form of properties that do not always represent the specified human context. Primarily, human activities or how humans interact with the place is wholly excluded from the notion of place [282],[281].

All of the works mentioned above rely on place names; hence, they face the same difficulties when implemented as gazetteers. The most apparent concern at the time of modeling of the place is the conceptualization of the spatial context and how this can be encapsulated in a set of entities that allow for spatial representation and realization. Thus, to model a place and integrate it into GIS systems, it must be established what needs to be represented about the place and how it can be projected, which requires understanding the concept of context and various contextual factors related to the place. The next section 2.2.3 provides an overview of spatial context.

2.2.3 *Existing Models to Formalize Concept of Place Considering Context*

This section illustrates numerous models presented to reflect the definition of place by taking into account one or more contextual factors and will address how they enhance the conceptualization of place and explains the research gaps or limitations related to different models or frameworks.

2.2.3.1 *Fuzzy, Probabilistic and Membership based Models*

Place is a concept that induces uncertainty [63]. People often use ambiguous location references such as The Midwest or The Alps, which lack specific natural boundaries to describe their spatial extents [169],[324].

Furthermore, Cognitive regions [267] are ambiguous regions commonly used in everyday language to pose questions about location. They reflect concepts widely accepted by people, such as the "city center" [268]. According to Montello et al. [267], cognitive regions are a superset of the definition of place; therefore, depicting cognitive regions should pose similar challenges to demarcating and locating places. In order to deal with such uncertainty, some models have been proposed, including the egg-yolk model [79], rough sets [70], fuzzy sets [49], supervaluation [217], and broad boundary regions [70]. [132] uses a data-driven approach to delineate and semantically enrich cognitive areas based on various data sources such as social media, blogs, and encyclopedias. Topic modeling algorithms are used to extract thematic knowledge about a specific cognitive area category (North and South California). Data mining techniques and fuzzy membership approaches are used to incorporate spatial properties in the extracted data. The outcome is a fixed grid that shows the probability that each cell is of the same form as the cognitive region in question, as adjusted by probabilistic models.

Probabilistic and membership-based approaches, recently with web-harvesting techniques, have visually represented a vague place's footprint [63]. These methods usually rely on the set of point positions provided by people associated with a place. The degrees to which every position belongs to such a place are visualized using such representations [63]. [268] conducted a study to determine the footprint of downtown Santa Barbara by asking participants to draw the city's boundary and then aggregating the results using density coloring since people think about place boundaries in different ways. Later, data-driven approaches focused on geotagged social media content with place names or tags, such as from websites like Flicker or Instagram [132],[150],[169],[189], were suggested using techniques like kernel density estimation (KDE) and clustering.

[131] provides a useful data-driven approach, in which the authors define the semantic signature of functional regions using probabilistic models. These regions can be conceived as locations on the surface of the earth that support specific services. The list of possible facilities is based on the Foursquare platform's location categorization [3]. A delineated area's functional context is derived mainly based on the types of included points of interest's co-occurrence patterns. Consequently, a region is assigned a specific function, which is conveyed via weighted vital words (i.e., 90 percent shopping,10 percent nightlife) [282]. Other research focuses on capturing the ambiguity of locations rather than their boundaries. Locality uncertainty measurements from various sources, such as approximate size, unknown datum, and map scale, were combined using several methods [151],[387]. Based on interpolating sources such as Wikipedia articles [6], other studies derive continuous fields of places with various thematic topics (e.g., forests) over the earth's surface.

2.2.3.1.1 Limitations

Nevertheless, none of these methods have been widely adopted as they are argued to be insufficient since they concentrate on the boundary rather than the core notion of place [394]. The bottom-up nature of the last two methods does not adequately account for human understanding of place; rather, it indicates the correlation of specific space footprints with human-generated data, such as events and opinions. The extracted information is exclusively data-dependent, as any model does not frame it. Furthermore, the probabilities assigned express the statistical significance derived from unsupervised data processing, which is not always understandable from a human viewpoint [282]. Finally, more elaborate models are either solely theoretical or highly subjective and lack adequate spatial projections, while data-driven methods are not easily interpretable by humans [282].

2.2.4 *Contrast Sets based on Place Models*

Winter and Freksa [394] recommend capturing a place's cognitive and linguistic nature compared to other similar positions. The process that a place emerges in cognition is quite similar to how humans perceive an entity compared to the ground [232],[352]. Contrast is a fundamental concept in understanding and sensing that contributes to a sense of uniqueness and wholeness [350]. A landmark attracts attention in an environment owing to its prominence compared to other surrounding locations [299]. The degree of contrast in some instances is low and difficult to interpret by a person, and as a result, a location can become uncertain and vague (e.g., it is often difficult to determine the foot of a mountain or the boundary of a city). When referring to a place in various ways, the place contrast set may differ, resulting in different place footprints. Such contrast sets of places may be described directly, pre-exist, or implied as shared information between the descriptor and the receiver in place descriptions [63]. Vasardani and Winter in [367] proposed a conceptual model to understand the region specified by the 'at' preposition based on the contrast set theory of places. The results are context-sensitive, as the generated region can differ when places in a contrast set to change due to a change in context.

3 https://developer.foursquare.com/docs/resources/categories

2.2.4.0.1 Limitations

However, it is unclear what characteristics decide contrast sets of places in conversations, allowing deriving contrast sets from natural language is a significant challenge.

2.2.4.1 *Function-Based Place Models*

There are also function-based models (e.g., [283],[385]), which are usually based on web-harvested data to classify places related to the functional semantics like green or commercial areas. Blaschke and Papadakis proposed a function-based model of place in [281], which depicts place as a functionally ascribed space. This model represents place as a topological graph of spatial entities that allows for a set of functions that define functional spaces. The dimensions of spatial properties, structure, and functions are used to describe places in this model. The model introduces two specific procedures the extraction of spatial patterns (also known as spatial design) and the infusion of space with a functional meaning (also known as a functional infusion).

The approach presented in [282] leads to the formalization of place through a hybrid approach that builds on the place's function-based model. According to a comprehensive composition model, the place has been described as a system of interconnected components that allow or disable a specific functionality. The model falls in the category of works solely concerned with information representation, such as ontologies and works that are strictly data-driven. Rather than affordances, the model assumes that a place is a physical space associated with specific functions and that suitable types of functionality support a place's ability to fulfill specific human intentions and desires. Purpose, function, composition, components, and data are the five primary semantic resolution levels in the full model.

2.2.4.1.1 Limitations

However, as the function(or set of activities that can be performed) associated with a place varies from person to person, it could lead to a different perception of place properties and different contexts, making it difficult to interpret the accurate meaning.
Nevertheless, this approach assumes that each place serves specific functions that should be clearly described and generally accepted. Thus, If such information is not available, the need for alternate sources arises, which can be met by consulting experts (such as architects), conducting surveys, or using data-driven approaches.
Furthermore, various cultures can describe similar locations differently: in the Eastern world, a shopping area's intended functions can include animal trading, which is unusual in the Western world. The same is true for variations in a location's functionality at various times of the day, week, or year: a football field, for example, could host non-football activities such as concerts on certain days of the year.
Additionally, the adopted function-based view of location ignores fundamental elements that people associate with places, such as emotions, sense of place, and purpose of existence.
Furthermore, treating the spatial extent of places using crisp boundaries estimates where a place ends. However, places can also be heterogeneous areas

or spatially fragmented entities. Therefore, using crisp conceptual boundaries to treat the spatial extent of places provides an estimate of where a position ends. Places, on the other hand, may be both heterogeneous areas and spatially dispersed entities.

2.2.4.2 *Place Models based on Core Concepts*

Varsadani, Tomko, and Winter investigated if a cognitively assisted collection of place properties could be used to instruct a place constructor—in its most basic form, a natural language text parser [368]. They used Alexander's [14] 15 structural properties that defined a whole and looked at how they corresponded to properties of the different applications of place studied in GIScience. They set out to see if a subset or superset of these properties is cognitively supported in the hopes that a place constructor can then use this set to computationally represent place instances that operate as a function of place properties. With the addition of affordance, the experiment revealed people's preference for a subset of Alexander's set, i.e., the contrast, intense middle, scale, positive space, boundaries, void plus affordance.
Purves, Winter, and Kuhn argue in a recent study [291] that current principles and theory in information science can be used to capture place information at desired richness or vagueness levels. Based on a recently proposed set of Kuhn's ontology of core concepts of spatial information [16],[214], the approach bridges the gap between the vague, flexible, and socially constructed concept of place and the formalism of information science. This study's main contribution is to distinguish the concept of location from its representation in information systems. They illustrated the links between an ontology of spatial knowledge on the one hand and the social and cognitive properties of places articulated in the literature on the other. They connect a formal, theoretical model of core concepts to a theory about the nature of the place, demonstrating how the two perspectives can coexist. This conceptual framework suggests ways of deliberately including and enriching the modeling of place by conceptualizing place as a type of context, rather than, as Harrison and Dourish [156] correctly pointed out, which could result in information systems that better fit real-world information needs.

2.2.4.2.1 Limitations

A place constructor will have to create one-of-a-kind locations, necessitating allocating specific combinations of property values to each position in a text parser that relies on property values. Even though this is a sufficient prerequisite, it may not be enough to establish places with distinct identities. It is unclear if intra and inter-place spatial connections can help with place identification and become part of a place constructor[110]. Similarly, a practical implementation of the conceptual model presented in [291] is required to see how accurately it helps improve information retrieval and analysis.

2.2.5 *Towards a model of place*

After observing the shortcomings of standard geographic environment conceptualizations when it belongs to place modeling, it is clear that translating human spatial concepts into information system representations has proved to be surprisingly difficult [291]. The argument for this difficulty lies in various ways of expressing spatial information in words and geometries, along with the importance of place as a social construct [148] and the importance of context.Notice that the objective is not to develop data models for places; once what needs to be portrayed about places has been determined, appropriate data models will be available [291]. Instead, from a theoretical perspective, we are confronted with the following challenges:

1. What is an appropriate model of place, and how can it be defined to capture human intuition, leading to a fully specified computation data model? [398].
2. What is spatial context? and how it can help us in understanding the meaning of place?

2.3 METHODS AND TECHNIQUES TO IDENTIFY AND RETRIEVE PLACE RELATED INFORMATION

Geographic information retrieval (GIR) is one of the most visible initiatives in GIScience to promote geographic information search [187]. This interdisciplinary field focuses on the geographic content of text documents, using theories and processes from computational linguistics and natural language processing. The growing number of very extensive text corpora with rich Spatio-temporal information [261] adds to GIR's cogency. Information extraction (IE) procedures first emerged in the 1990s with message understanding conferences, which added various IE activities over time [183]. IE's earlier stages extract information from text using template filling, classification model-based methods, rule-based methods, and sequential labeling.
An IE method extracts structured information from natural language text and represents it. Extracted text strings, values, or tags are listed in predefined slots of user-defined structures known as templates or objects [288]. Similarly, another IE method for Web-derived unstructured text presents structured information in XML by performing data extraction, syntactic and semantic analysis, classification, and inference rules steps [312]. Furthermore, machine translation, auto-coding, indexing, and term extraction are the primary techniques used in the IE process to provide meaning to unstructured data, with auto-coding and indexing tasks assisting in identifying words from text [45].
Recognizing place names from texts and linking them to spatial footprints are essential steps for automatically understanding the semantics of NL texts. They are studied both in computer science and GIS [172]. A Place Semantics study aims to understand Place's meaning using human descriptions [171]. Cognitive geographical concepts usually refer to informal geographic information acquired by individuals through encounters with the local environment [143]. Egenhofer and Mark referred such unstructured information as naive geography [104] and understanding these geographical concepts is assumed to improve design of a GIS [338].

Named entity recognition, relation extraction, and event extraction present the sub-tasks of the GIR process from text data. In this research, we focus on named entity recognition and the relation extraction processes.

2.3.1 *Named Entity Recognition (NERs)*

NER is a technique for extracting descriptive information from identified entities. Entities can be general, such as people or places, the organization, and numeric expressions, including time, date, and monetary amounts, or domain-specific, such as proteins, chemicals, and cells. Entity identification (named entity detection) and classification (semantic classification) are NER subtasks [257]. It involves assigning each word or group of words to a set of pre-defined categories or entities, including "not an entity" [364].

The early research in NER seems to extract coarse-grained entities, however, more detailed location classes are being explored by [364],[27]. Most NER systems are based on a list of well-known locations, organizations, and individuals (also known as

gazetteers [162]). Nevertheless, building gazetteers that are credible and accurate is quite challenging [329],[258],[1]. It is crucial to have all proper names with no duplicates and correct spellings. It should also contain vernacular names and official or administrative place names [196],[189]. Besides identifying the named entities within the text, it is also necessary to resolve the semantic ambiguity related to objects with the same name that can be either a place or belong to another category, which requires context information.

They are combined with various rule-based or machine-learning techniques to attain the ideal output [3]. Rule-based NER methods use lexico-syntactic patterns and semantic constraints to define the occurrence of related entities, while learning-based methods use machine learning to extract named entities and classify them. Learning-based methods may be supervised, unsupervised (hard and soft clustering), or semi-supervised (bootstrapping). Several studies tend to show hybrid approaches outperform single approaches in terms of efficiency and accuracy. The tables below provide an overview of different NERs approaches being used to identify and extract information related to entities.

The following Table 1 illustrates several reasoning techniques employed by academics to extract named entities in different languages, whereas Table 2 discusses machine learning approaches. The hybrid methods employed in various research are discussed in Table 3.

Table 1: Overview of Rule-based NER approaches for extracting information related to entities

NER Techniques	Purpose/Applications	Reference
a local grammar combined with a gazetteer	for identifying entities with Arabic names.	[331]
set of rules and patterns for Urdu	address issues the agglutinative nature of Urdu, lack of capitalization, spelling variations	[304]
using various rules and patterns group of dictionaries	to extract named entities in Urdu	[336]
a rule-based approach containing heuristics grammatical rules	for drug-related entities	[297]
a set of dictionaries stemmed matching fuzzy matching	for extracting named entities from electronic free health records	[296]
noun phrase chunker followed by a filter based on inverse document frequency extracts candidate entities	for extracting named entities from biomedical text	[413]

Table 2: Overview of Machine learning NER approaches for extracting information related to entities

framework of conditional random fields model	to get improved tag-set for Chinese language	[411]
ontology-based semi-supervised (CRF)-based information extraction	for extracting information entities describing existing deficiencies and performed maintenance actions from bridge inspection reports	[246]
Novel features based on latent semantics	Language independent NER	[205]
framework for domain-specific NER based on a unigram probabilistically labeled pseudosentences	for geological name entity recognition in geological reports	[295]
MEta-Map chunker-based noun phrase extraction, SVM, supervised learning CRF	for medical entity recognition	[2]
Sequential labeler based on the linear Conditional Random Fields	for clustering similar tweets	[247]
Conditional Random Field	to identify named-entities from homeopathy diagnosis discussion forum	[253]
a novel kernel function for support vector machines	for sequential labeling tasks like named entity recognition	[314]
Convolutional Neural Network (CNN)	extraction of location from tweet text	[218]

Table 3: Overview of different Hybrid NER approaches used for extracting information related to entities

NER Hybrid Approaches	Purpose	Reference
maximum entropy model language-specific rules gazetteers	for the Hindi and Bengali language text	[313]
rule-based pattern extractor using link grammar parser Stanford PoS tagger semi-supervised classifier	for entity labeling	[318]
Conditional Random Fields (CRF) with dictionary	a chemical NER system	[309]
Support Vector Machine Conditional Random Fields	for biological entities	[417]
Supervised tagger TnT rule-based support vector machine (SVM)	health and tourism text documents in Malayalam language	[182]
Combination of Hidden Markov Model with manually crafted rules	for the Punjabi language	[32]
Hidden Markov Model plus rules (PoS tagging for entity detection)	to extract named entity-specific classes from the Nepali language	[335]
Combination of dictionary-based, rule-based, machine learning	to extract molecules related properties from scientific literature in biomedical domain	[106]
Combined dictionary-based approach fuzzy matching stemmed matching	generates a set of annotation from the clinical text	[296]
Support Vector Machine Hidden Markov Model linguistic pre-processing methods	to identify gene and protein from text without using external knowledge base	[25]
manual engineered rule-based predecessor lexical resources and pattern	for semantic indexing of the Turkish text	[212]
Combination of Hidden Markov Model with a gazetteer	for tourism text in Hindi	[177]
using morphological rules	to extract nouns from classical documents in the Malay language	[319]
Rule-based Conditional Random Fields Random Forest Bidirectional-LSTM approach	Named Entity Recognition for sensitive data in Portuguese language	[97]

Tables 1,2,3 survey how information retrieval for extracting and recognizing entities in a variety of disciplines, from geography to biology, is being pursued. Nevertheless, as can be observed, the primary focus is on language comprehension and identifying diverse sets of entities from other languages. Even with this level of advancement, the substantial problems associated with information retrieval, mining, and analytics must be understood.The problems identified in the literature can be divided into entity-specific and technique-related issues, which are discussed in detail in following section2.3.1.1.

2.3.1.1 *Limitation of NER Approaches*

Traditional NER methods (POS tagging, relation extraction, Toponym resolution) are insufficient for dealing with the dimensionality and heterogeneity of unstructured data [8],[188],[249],[251]. Understanding spatial referents and their applications are critical for methods beyond named entity recognition to recognize geographic references and analyze their relation with textual content more nuancedly [188]. One area of current interest includes applying machine learning techniques, specifically so-called deep learning, to several problems in natural language processing in general [188]. The potential application of such methods is, for example, discussed in the context of named entity recognition by [260] and surveyed in table 2 and 3. However, it is vital to compare basic, well-expressed rules applied to a specific domain [231] or the use of a localized gazetteer [240] to machine learning-based approaches that involve training data. This is demonstrated by several hybrid approaches 3, that used specialized rules and a general ML method. Many machine learning methods, in particular, concentrate on datasets with readily accessible training data in the form of metadata associated with documents (e.g., Wikipedia), which may not be a practical scenario for more general applications applied to unstructured text [188]. Supervised learning methods require large amounts of annotated data for training, which is a time-consuming and challenging process for large-scale data sets. Weakly supervised learning is more efficient than supervised learning because it requires reduced manual effort. However, due to data scarcity, these techniques are inefficient [235]. As a result, we believe that a critical challenge for future research in GIR, especially geo-referencing, is the reproducible publication of methods, algorithms, data sets, and results, allowing approaches to be more easily compared across corpora [188]. Current research employing end-to-end monolithic deep models fails to solve complex tasks requiring deep language comprehension and reasoning skills [175].

On the other hand, abbreviations, the open nature of vocabulary, disambiguation, and various languages and domains are critical entity-specific challenges [8],[188]. Furthermore, entity ambiguities (single entity and global entity) [380], noise (short and domain-specific text) [183], and automatic labeling [68] are complicating the identification of entities and their relationships from free-text data sets. Several factors that influence the performance of NER techniques have been identified, including noise, data diversity, variance in text perspective, and data sparsity. Similarly, all defined factors are classified into entities, data, task, domain, and language-related limitations and are shown in Table 4.

Table 4: Overview of Limitations of NER approaches based on issues related with entities, data, language, domain and methodology

Limitations	Influencing Factors	Studies
Entities-related	Ambiguous and uncertain entities	[25, 336, 380]
	labeling issues	[106, 331, 304, 336, 253, 68, 218]
	Semantics of Named Entities	[257, 2]
	Contextual relation among entities	[106, 296]
Data-related	Noise	[183]
	Data Diversity	[183, 296]
	Textual perspective variation	[3]
Language-specific	Single Language	[32, 335, 212, 331, 411, 97]
	Multiple Languages	[296, 177, 314, 205]
	Languages with poor morphology	[182, 319]
Domain-specific	Domain-specific entities	[3, 318, 309, 417] [297, 97, 413, 246], [295, 380, 2]

In conclusion, most of the work is either domain or language-specific, with the remainder having issues such as automatic labeling or being unable to grasp the context of an existing entity and provide an unambiguous response. Since this research focuses on spatial information extraction, we concentrate on entity-related limitations and how to deal with unstructured text and the information extraction process.

2.3.1.2 *How can NER be improved?*

We observe that existing NER methods are not sufficient for advanced place interpretation. This work makes the hypothesis that instead of adjusting NER techniques, performance can be boosted by putting NER in a feedback loop with the complete approach to text understanding. Specifically we are motivated to investigate the following questions:

1. To what extent does the use of contextual information in addition to the entity's name aid in improving the performance of existing approaches?
2. Will the use of basic reasoning procedures enhance NER in context of a complex system?
3. Can reasoning and exploitation of contextual information assist us in dealing with vernacular and type references in addition to named entities?

These questions lead to various hypotheses generated in section 2.6, which motivates our research in chapter 5.

2.3.2 *Relation Extraction*

Interestingly, since the understanding of the text is considered AI-complete, a common approach to extract structured data from the text to make it machine-readable or process-able [337]. One of the major subtasks in this

context is to extract semantic connections between different entities in free-form text, known as relation extraction [337]. Relation extraction extracts and classifies semantic relations from natural language text([8],[337]). In a nutshell, it enables us to derive structured information from unstructured text [8].
Much research has been done on extracting information from a natural language place description depending on the data source, type of relationships (cardinal, distance, etc.) and several machine learning techniques [357]. This section provides us with an overview of the different existing approaches used for spatial information extraction.

2.3.2.1 *Application dependent Information Extraction*

Geographic scope resolution allows for limited possible interpretations within a defined scope [405] and is one of the approaches used for interpreting place descriptions. Work in the past relied on mining is limited. For example, application-dependent geographic information from supervised language expressions is presented in([194],[234],
[353]) or from specific contexts are presented in Tschander et al. [360], which describes an artificial agent capable of following route instruction, using a conceptual-level instruction language. In contrast to that agent, we are mostly interested in interpreting describing statements like "campground south of Bamberg", rather than processing incremental instructions like "take road R123 south".

Whereas Andogah et al. suggest a system based on a collection of predefined geographical scopes [21]. Hornsby and Li [170] provided a conceptual framework that could serve as a basis for tagging and parsing text to extract movement information. The semantics of locative expressions of various languages are studied in Kracht et al. [209]. Yet, all these approaches are not sufficient and has certain limitations presented below.

2.3.2.1.1 Limitations

Nonetheless, the following disadvantages apply to application dependent information extraction approach:

1. the method is not able to identify entities which do not exist in the data sets,

2. the approach is unable to capture the surrounding context thus just providing information that already exists within the data set.

2.3.2.2 *Information extraction using Machine Learning Approaches*

For solving the problem of relation extraction, several existing methods employ carefully built features and standard classification techniques [87]. The system must correctly annotate the data by identifying a text with the desired semantic property [8].Several methods have been used to determine the relationship between identified entities,the typical survey on relation extraction in [89] distinguishes approaches into two major branches: supervised methods and semi-supervised methods([8],[337]).
Supervised methods for domain-restricted relation extraction (RE) use feature-based and kernel-based (bag of features kernel and tree kernel) techniques for RE [8].Semi-supervised approaches for open-domain systems include Dual Iterative Pattern Relation Expansion, Snowball, and KnowItAll [206].

These methods are only applicable to sentence-level RE, while RE in paragraphs and cross documents can increase accuracy [28].
The researchers suggest machine learning approaches to manage ambiguity in linguistic, spatial information and in [207] use the SemEval-2007 project (TPP) [243] to dismiss the spatial meanings of prepositions and enhance the technique of spRL.In [23] they automated the extraction of spatial triplets using a learning model to represent qualitative spatial relations via a simple class of locative expressions.This model does not cater the indirect references like "I am here at the old building".
Linguistics may use the hybrid method to derive the relationship between complex terms. The table 5 below summarizes the details of current state-of-the-art machine learning approaches for relation extraction.

Table 5: Overview of Relation extraction approaches

RE Techniques	Purpose	Reference
filtering algoritm based on deep linguistic analysis morph syntactic linguistic filters	extract relations between complex terms of Arabic language.	[222]
semantic and probabilistic approaches	for biomedical database Medline	[256]
feature-based methods combined with convolutional recurrent neural network	to capture effective hidden structures within ACE 2005 corpus	[274]
lexical-syntactic pattern matching	for open IE system to extract domain-independent terms	[402]
statistical calculus with linguistic knowledge	for extracting noun phrases and transforming into semantic relations in Arabic	[388]
the SpRL-CWW model uses a CRF model SVM model	for spatial relation and element extraction	[276]
UTD-SpRL model based on SVM	for spatial relation extraction	[89]
X-Space model proposed uses node information uses argument information	for spatial element extraction spatial relation classification	[315]
a convolutional neural network based system	to extract spatial roles and their relations	[259]
BERT-based spatial information extraction model	extraction of spatial element and relation	[333]

Table 5 provides an overview about the research work carried out in extracting spatial relations from text, focusing on the text's static and dynamic spatial relations. It helps in natural language understanding systems, such as question-answering systems, robot navigation, and understanding

geographical relations or tracking moving objects [333]. Relation extraction helps understand various languages and serves as the foundation for query formulation, whether for a search or language interpretation. The majority of RE methods are hybrid. They usually treat the task of relation extraction as a classification problem and solve it using a machine learning approaches [408].
Most RE techniques extract one-to-one relationships between entities, but many-to-many relationships can also be observed [8]. Unstructured data sets, in this regard, necessitate high computational systems to improve performance efficiency and reduce computational latency. Because of the scalability and sparsity of unstructured data, conventional approaches are ineffective [235].

2.3.2.2.1 Limitations

The above presented research uses the concepts of machine learning, however, these methods perform well in domain specific RE but require plenty of training data. The major drawbacks of different type of learning techniques are discussed below:

1. **Supervised Learning:**Annotated training data is needed for supervised approaches ([58],[275],[410]), which is expensive to generate. This limitation also makes supervised approaches difficult to expand since detecting new relation types necessitates new training data. Furthermore, supervised classifiers are biased against the text domain used for training and produce sub-optimal results when applied to other textual contents. Overfitting is a common problem in supervised relation extraction approaches (because the corresponding training data is usually limited and specific) [337].

2. **Semi-supervised Learning:** approaches ([10],[54],[110]) usually rely on bootstrap learning. Initially, they use a small dataset to learn how and when to extract additional relations and then iteratively use the extracted relations for training.
   While semi-supervised approaches minimize the need for manual efforts to generate training data, they still require an initial set of labeled pairs for each extracted relation [337]. As a result, extending such approaches to derive new relation types usually necessitates human involvement.

3. **Unsupervised Learning:** Unsupervised relation extraction techniques have also been suggested previously [334] and have since developed into sub-field of Open Information Extraction ([35],[111], [401]). Even though unsupervised methods do not require any training data, they usually yield sub-optimal results that are difficult to interpret and map to an established set of relations, schema, or ontology [112]. This last limitation is critical for many applications, such as knowledge base refinement [337].

4. **Distant Supervised Learning:** Two main challenges in distant supervision include 1) noisy labels since the sentences mentioning a pair of entities do not necessarily express a relation; moreover, different sentences mentioning the same entity pair can express different relations. 2)incompleteness of the knowledge base, which affects the training

and evaluation phase [337].The simple distant supervision method ignores these two issues in favor of a very ad-hoc implementation based on various hand-crafted features (both lexical and syntactic) [337]. Another challenge is to consider more context and background information when extracting relationships. In that case, one alternative would be to consider more than one sentence (i.e., the entire paragraph) when extracting relations, similar to the relaxed setting in [26].

2.3.2.3 *Information extraction using Granularity and Ontology concepts*

Another method used in geographic information extraction is to know the finest possible granularity level concerning a general ontology of spatial entities, which helps resolve place descriptions [405]. Richter et al. [305] consider granularity effects caused by object types. Linguistic notions are often challenging and result in under- or over-specifications, when attempting to match NL with formal spatial information frameworks [23]. As stated in [37], majority of these linguistic complexities are not considered in the computational aspects of geographical knowledge, especially when constructing formal spatial models based on the logic of spatial cognition by humans. The researchers also argue that these formal models do not understand the way people express spatial features linguistically. Formal ontologies have been claimed to offers adequate means to represent semantic commitments of a spatial language phrase [37].As a result, Bateman [36] suggests a spatial ontology based on semantics and linguistics that would enable efficient mapping between NL spatial expressions and spatial calculi, like the Generalized Upper Model Ontology (GUM) [39].
In spatial semantics, Zlatev [419] worked on cognitive-linguistic research and provided basic theoretical concepts for its understanding. Zhang et al. [412] developed a framework that addresses the automation of geospatial search features with the combined help of geo-semantic web technologies and natural language interfaces. The area of application is based on managing disasters and emergencies. To achieve this objective, efforts are mainly focused on issues concerning the extraction of geospatial features via the Semantic Web.

2.3.2.3.1 Limitations

Current methods are still unable to incorporate the contextual dependencies successfully and have certain drawbacks which includes:

1. The construction of ontology always lacks a unified system and evaluation criteria
2. as opposed to ontological information existing ontology languages do not fully support the spatial domain and manifold spatial relations to the extent required to empower spatial reasoning
3. Unable to capture vague semantics of concepts like "near"

Existing research indicates that present methods are incapable of capturing context or background information together with the semantics of spatial relations or objects. Furthermore, the presented machine learning approaches do not exploit the implicit information present in the input, making it impossible to capture new relation types between different entities and cannot decide a single answer in multi-relationship situations due to lack of contextual information. The parser also requires re-training. In addition, the new entities remain unidentified.

Thus the next section 2.3.3 outlines and concludes the inadequacies and the problems of existing RE methods.

2.3.3 *General Limitations of RE Approaches*

The crucial factors that lead to the limitations of current RE approaches are summarized in Table 6. The majority of the work in this domain is focused on spatial knowledge extraction, language interpretation, and defining relationships in medical documents. Since spatial information extraction is also a sub-goal of our research, we concentrate on the factors contributing to relationship identification, data annotation/labeling, and how spatial language can be understood from textual information.

Table 6: Overview of Limitations of RE approaches based on issues related with relationship identification, data,language and technical issues.

Limitations	Influencing Factors	Studies
Issues related with relation identification	entities and relations related to domain	[222, 256, 333]
	Identifying Semantic relation	[274, 388, 89, 381, 358]
	errors in parsing	[23]
Data-related	Dimensionality of data	[235]
	Data Diversity	[235]
	Volume	[114, 382]
Language-specific	Language Ambiguity	[206]
	No Multilingual IE	[222]
Technical Issues	labeling issue (unlabeled data)	[206, 276, 315]

2.3.3.1 *Approaches to deal with existing challenges for RE Processes*

The limitations discussed above give rise to innovative ideas for improving the present methods of extracting relations from geographical text descriptions. The research areas listed below can assist us in identifying solutions to current shortcomings.

1. Is it possible to improve the process of relation extraction by using reasoning approaches rather than formal learning and classification methods?

2. Can deeper granularity via reasoning helps in capturing new relations and confirming the existing ones by utilizing implicit knowledge present within the sentence?

3. Can reasoning techniques assist us in overcoming the requirement for a large amount of training and test data?

4. Is reasoning capable of solving the parsing issues like co-reference resolution by generating all potential relationships between distinct entities to comprehend and absorb contextual and background information?

These questions motivate our research and lead to hypothesis generation in Section 2.6 of this chapter. The above questions are answered in Chapter 4.

2.3.4 *NLP Parsers and Open Issues*

When it comes to extracting spatial information from natural language text, either classic natural language processing tools like parsing or information retrieval techniques can be applied. Several works make use of different parser and their modules (like Standford NLTK chunking or the dependency parser[4]) for relating the objects and relations between them. The rapid development of text mining and NLP techniques makes it feasible to extract information from textual documents, through information extraction (IE) techniques [83] like spatial relationship extraction ([23],[207],[244]) , entity relation [87], and event extraction [197] as well parsers to extract triplets from text have also been developed in ([23],[244],[207]).
Parsing is involved with revealing the sentence structure by constructing a so-called parse tree that reflects the grammatical structure. A parser can thus identify how distinct parts of a sentence are related, in particular identify spatial relationships expressed. The feasibility of creating such parsers has been demonstrated in previous research (for example [13] and [165]). Our main argument, text understanding based on parsing requires that a nearly complete syntactic analysis of all input sentences is achievable, which poses various constraints on possible input data [277]. Casual or incorrect grammar generates problems, and even for correctly constructed sentences, in particular related to indirect references, errors may occur.

Thus, dedicated treatment of possible parsing errors would be required. In this section, all such cases are discussed.

2.3.4.1 *Sentence structure*

Parsers pose requirements regarding word order and dependencies on verbs. Consider for example *post office near Bamberg main station*; the phrase is as a perfectly plausible place description that identifies a particular post office. However, the phrase does not contain a verb – it is not a proper sentence. Interpretation of place descriptions thus goes beyond what is addressed by contemporary parsers.

Figures 3 and 4 exemplify how sentence structure affects the working of existing parsers, using Stanford Core NLP parser as an example. The figures show the part-of-speech tagging using the so-called Penn TreeBank Tagset [354]. Let us exemplify what the tags mean by discussing both sentences. Authors of both sentences are describing a brewery and its whereabouts in the city of Bamberg, Germany. We can see in the figures nouns labelled with tags NN (noun, singular or mass noun), NNP (noun, singular), NNS (noun, plural), and NNPS (proper noun, plural). Individual words are then connected by relations, for example in Figure 3 *one* is related to *breweries*, *Schlenkerla* and *center* are related by relation *in*, *center* and *Old Town* by relation *of*. By contrast, no relation between *Schlenkerla* and *center* or *Old Town* can be found in Figure 4. Output for the second sentence does not indicate any link between the brewery and its name, which is due to the acyclic structure (brewery name at end) and non-linear order of the sentence. We may thus conclude that the parser failed to decode the sentence structure completely. We do not blame the parser since language is incredibly flexible and hard to decipher. For an automated system to interpret place descriptions, relying on correct parsing would however require additional means to correct the output of a parser.

4 https://nlp.stanford.edu/software/

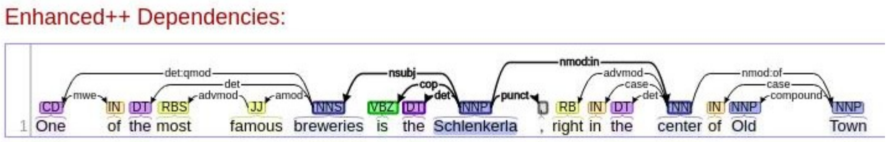

One of the most famous breweries is the Schlenkerla, right in the center of Old Town.

Figure 3: Interpretation of Wikipedia sentence using CoreNLP parser.

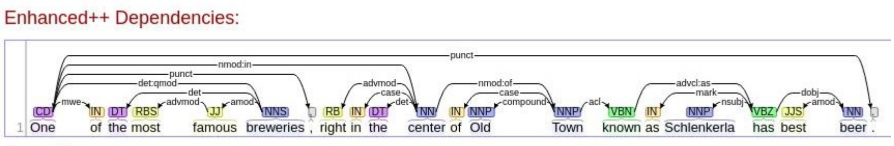

One of the most famous breweries, right in the center of Old Town known as Schlenkerla has the best beer.

Figure 4: CoreNLP parser output of a sentence from a travel blog.

2.3.4.2 *Sentence length and multiple clauses*

Sentences containing multiple clauses are known to be particularly difficult to handle by a parser due to possible dependencies between individual clauses [23]. Figure 5 shows output from parsing such type of sentence. As can be seen, the last clause is not linked to the previous one, resulting in wrong dependencies. In the sentence shown in the figure, *Ludwigstraße* is described to be one of Munich's main streets. All nouns – *Ludwigstraße, Munich, streets* – have been identified as such, marked by the appropriate tags NNP (proper noun) and NNS (proper noun, plural). While Munich and street appearing in the same clause are related by the parser, no connections between *Ludwigstraße* and *street* are made.

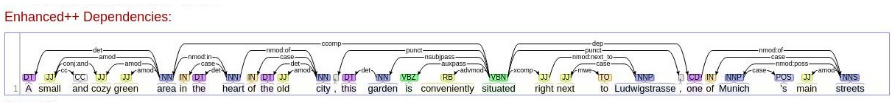

A small and cozy green area in the heart of the old city, this garden in conveniently situated next to Ludwigstrasse, one of Munich's main streets.

Figure 5: Excerpt from CoreNLP parser output of a multiple clause sentence from Wikipedia.

Similarly, the Stanford parser fails to identify the dependencies with long sentences. For example "I am[at a private residence]located on the western side [of KFC] [near shoe shop], [in Bamberg]", the dependency between private residence and Bamberg is not identified. In addition the relation between KFC and Bamberg is unidentified, Extracted triplets are:

Triplets extracted:
<I, at ,a private residence>
<a private residence, on, the western side of KFC>
<shoe shop ,in, Bamberg>

2.3.4.3 *Implicit references*

Current parsers are unable to identify the relation between two entities in a reliable and robust manner if the key phrase describing the relationship does not appear in the sentence. The example shown in the Figure 6 demonstrates how the Australian state of Victoria is wrongly identified as the city Victoria.

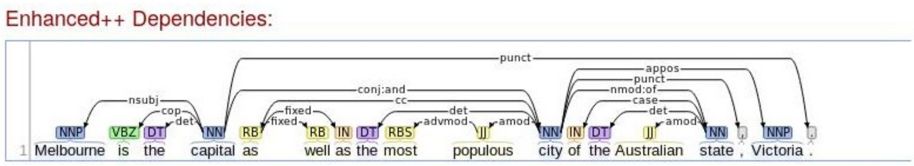

Figure 6: CoreNLP parser output of a sentence from Wikipedia containing indirect references.

2.3.4.4 *Unable to Categorize Unknown Words*

The geo-parser should ideally categorize an unknown word as a location or not, merely based on the word itself and its context as the same word can be labeled both as a place name and some other entity, also known as geo/non-geo ambiguity [18]. Geo ambiguity is a significant concern because practically all the names of places may relate to several places [339], which leads to the drawback that when a component has incorrectly categorized a text, the error is propagated across the entire pipeline, influencing the results of other components and the ultimate result [326].

2.3.4.5 *Error Propagation*

NLP pipelines typically follow a fixed order of execution of its components. This is similar to the way traditional deterministic parsers based on a shift-reduce framework work, traversing a sentence from left to right. However, humans tend to parse text by first looking for easy attachment decisions and use those to build isolated constituents [326]. With more structure forming, the knowledge can then be used to make the initially complex attachments straightforward. This method is also called easy-first parsing [142].Tratz [359] observed that errors committed in lower-level tasks propagate through the remainder of the processing pipeline.

However, as we are only interested in spatial entities which correspond to nouns in the input text and the relations holding between them, a parser which has to take the verb of a sentence as starting point may not be necessary [405]. In addition, the systems [23],[63], used parsers for extracting dependencies and identified that better results would be achieved if they are re-trained.

2.3.4.6 *Open Research Questions to deal with existing NL Parsers*

As a result of these constraints, new approaches for deriving relations and identifying spatial entities from NL text descriptions are required. The areas of research outlined below can help us find answers to existing problems.

1. Is it feasible to enhance the information extraction process by not relying on parsers and instead employing an explicit reasoning layer?

2. Can we develop an open-ended location description system that can accommodate a wide range of place descriptions, not just the structured ones?

3. Using logic and reasoning, is it feasible to categorize unknown words or unnamed entities?

The answers to these questions can be found in our overarching thesis.

2.4 GEO-REFERENCING OF PLACE

Geo-referencing is when places are associated with geographic locations, and is performed mainly by an external information source, usually a gazette-

er [63, 188]. Most current techniques for geo-referencing employ named entity recognition (NER), a technique that identifies names in a piece of text. Such systems are also called taggers or geoparsers if tailored to geographic information. They classify words by labeling every word with a category such as the name of a person/entity, noun, verb, etc. Geo-references have various essential and appealing characteristics [163] like they should be unambiguous, referring to a single location only, usually the case within a given frame of reference.

Georeferences should also be as persistent as possible over time. Finally, geo-references are often correlated with an implied level of granularity. As a result, our objective is to unambiguously classify these geographical references and, in most cases, assign spatial coordinates to them.
Moreover, in this way, we can guarantee that a geographic relation is linked to a unique location. DBpedia Spotlight[5] aims for a more fine-grained tagging that allows named geographic entities and key concepts to be distinguished. The system can, for example, identify fine-grained concepts such as waterway, but it cannot handle several everyday words like, e.g., *creek*. Simple taggers for geographic information may be realized by comparing words with entries in a gazetteer. Current NLP tools use classifiers obtained by machine learning, as spaCy[6] or Stanford Core NLP.[7] More sophisticated methods employ different techniques, see, for example, the GeoTxt system [56]. GeoTxt was mainly developed with unstructured micro-blogs like Twitter in mind, and it was found to be particularly robust against non-capitalization of place names [149], a simple syntactical change that irritates several other systems.
Place names identified by a NER system are then mapped to geographic entities by a process called toponym resolution [229]. Toponym resolution (TR) [63, 228] is a term used to describe the process of locating places mentioned in textual documents. Toponym resolution is one of the core tasks for building GIR engines and determining the geo-focus of textual documents ([19, 220, 292, 391]). Exploiting knowledge about typical co-occurrences of toponyms improves disambiguation between possible spatial interpretations, similar to how knowledge of geographical context allows the name of a city or restaurant to be disambiguated. Thus TR contains two main subtasks, namely the toponym recognition and disambiguation. Alternative names have also been used interchangeably for the two subtasks, e.g., geoparsing and geocoding [63, 188]. Toponyms are usually recognized using

5 https://www.dbpedia-spotlight.org/
6 https://spacy.io/
7 https://stanfordnlp.github.io/CoreNLP/

NLP techniques based on NER and gazetteer matching (e.g.,[230, 250]) and ignore the non-gazetteer place reference. Section 2.3.1 discusses the process of named entity recognition and existing research related to it in detail. The process of mapping each recognized place name to its accurate, unambiguous geographic location ([229, 59, 230, 20]) is necessary because toponyms are rarely unique and may have several gazetteer entries [63, 188].
There are several geocoding methods (referent disambiguation or toponym resolution). Toponym disambiguation approaches are usually performed by looking for context locations, i.e., other place names that appear in the same document, and calculating the likely hood of each candidate gazetteer entry corresponding to the place to be disambiguated. The likelihood hood is calculated as a score based on heuristics (or a combination of heuristics ([96, 226]) and information such as position or containment relationships. Additional heuristics may be linguistic or geographic properties that are well-known. Leidner [229] summarizes a detailed list of heuristics that have been used in previous TR studies.
Buscaldi [59] divides toponym disambiguation approaches into three categories: (i) map-based: those that use spatial distances to find disambiguation clues; (ii) information-based: those that use external knowledge sources like population statistics to find disambiguation clues; and (iii) data-driven or supervised: those that use machine learning techniques. Brunsting et al. [56] propose an approach that first retrieves all possible toponym interpretations by querying a spatial database and then applies a scoring function that is essentially based on geographic distance. Gelernter et al. [137] proposed a matching algorithm that can compare the tags in OpenStreetMap and Wikimapia with the place entries in a gazetteer and can add the place information not present in a gazetteer. Another way of identifying place names is by finding semantic similarities between named spatial entities using domain ontologies [238]. Also, training of end-to-end models has been considered for matching place description to entities, thus combining NER and toponym resolution [252]. Vasardani et al. [371] proposed a place graph model to represent places and relations between them. In their approach, geo-referencing starts with named entity recognition to identify locations in different gazetteers, aided by toponym resolutions [152] using a graph database. A limitation of that model is its inability to capture contextual knowledge from place descriptions.
The overview and details of existing approaches are provided in the following Section 2.4.1.

2.4.1 *Overview of Various TR Approaches*

This section provides a list of various TR approaches currently being used by researchers for geo-referencing purposes.

1. For disambiguation, map-based approaches depend primarily on the location of gazetteer entries of place names from a text and use heuristics such as minimum convex hull, minimum point-wise distance, or locations nearest to the centroid of all entries (e.g.,[18, 340, 414]) and they also work on disambiguating fine-grained locations.
(e.g., [95, 265, 279]).

2. External knowledge of locations, including population, prominence, or containment relationships, is used in knowledge-based approaches(e.g., [7, 60, 192]).

3. Machine learning approaches focus on building language models from training data that reflect the probability of seeing each of the locations from the same document associated with a disambiguation location (e.g., [133, 307, 341]).

4. More recent approaches consider other words along with place names from the same documents such as events, person names, or organization names to assist disambiguation ([66, 245, 310, 392, 278, 343]). There are also hybrid and bootstrapping methods ([4, 67, 153]).

5. Rather than relying solely on text, alternative techniques make use of spatial relationships. In comparison to text-based approaches, these methods perform disambiguation by using the coordinates of places that appear in the context of an ambiguous toponym [340, 229].

6. In [136, 345], a graph-based method for co-occurring toponyms from Wikipedia was suggested based on the construction of a weighted network.

7. Some studies are focusing on developing methods that are not based on a gazetteer. Toponyms and records have been geo-referenced using language models ([66, 310, 392]).

8. Some algorithms are less parameter-sensitive and, as a result, do not necessitate prior knowledge of the input data. [44] provides an overview of these clustering algorithms.

Despite the extensive studies related to toponym resolutions, there are areas of research that still require attention. Following section2.4.1.1 provides us with an overview of specific challenges related to place geo-referencing.

2.4.1.1 *Limitations of existing place Geo-referencing approaches*

This section highlights the challenges and limitations of georeferencing. Some of which were also proposed by Leidner and Lieberman [230] back in 2011. These include:

1. Addressing vernacular and historical references i.e., TR approaches are usually limited to gazetteer locations (i.e., officially indexed in gazetteers). There has not been as much focus and effort to develop approaches for resolving other types of places or references to places [63, 188].

2. Efficient and effective geo-parsing and geo-coding at various granularity or scale levels (i.e., working effectively with local-scale and global-scale data), i.e., spatial granularities greater than or equal to suburb- and city-level are considered, and references to unrelated locations are ignored [63, 188].

3. More intelligent ways to process spatial expression (e.g., "40 kilometers north of Kabul") are required, i.e., Using the surrounding context is a more successful strategy. Contextual information is often classified as either internal or external [63].

4. Compared to knowledge and machine learning-based disambiguation approaches, map-based approaches appear to be more robust for fine-grained position disambiguation because they only require knowledge of the locations of ambiguous candidate entries [63].

2.4.1.2 *Open Research Questions for Geo-referencing*

Based on existing research, one can conclude that TR focuses primarily on places that are gazetteered and overlook references to places that are not, resulting in numerous research topics that still need to be explored to enhance the overall effectiveness of geographic information systems.

1. To what extent can anticipating a geographical entity's type together with its name enhance toponym disambiguation?
2. What happens to the geo-referencing process if we explicitly consider contextual factors like the type of spatial entity?
3. When it comes to geo-referencing systems, can reasoning approaches help?

These concerns motivate our work of dealing with nameless entities and determining how they impact the performance of current systems.

2.4.2 *Spatial knowledge Representations and Reasoning*

When explaining the whereabouts of places and (spatial) objects, people often use qualitative spatial relations. When the descriptions are from memory, they are based on cognitive spatial representations of the world or mental models [216],[363]. People often rely on mental models to solve basic reasoning tasks, even though such models frequently provide incomplete and biased results [201]. Such qualitative spatial relations are often expressed in English by spatial prepositions, and the semantics of such prepositions in terms of spatial, temporal, and geometrical meanings have been studied in linguistics ([134, 223, 355, 421]). A parser can identify and extract spatial prepositions from the text [63].
Qualitative spatial relations have been studied in the Artificial Intelligence community for QSR ([80, 124, 242, 302, 418]), formally recognized through logical or algebraic calculi. QSR is mainly associated with the computer science field of information representation, although it is also closely related to geoinformatics, linguistics, and cognition studies. QSR is an active research area that aims to establish computationally efficient means of grasping the catalog of human common sense concepts of space and enabling automatic spatial reasoning on the concept level [301],[65].

Qualitative representations of space use symbols to describe semantically meaningful properties of spatial structures, abstracting away any information which is not considered necessary to the application context at hand. Various objectives motivate qualitative abstraction research, especially the desire to design formal models of common sense dependent on coarse concepts [389, 53] and to compile a catalog of human cognition's principles and inference trends [215, 203], which, when combined, allows for intuitive approaches towards designing intelligent systems [92] or on human-centered GIS approaches [121].
A finite collection of crisp elementary concepts emerges from qualitative abstraction. There has been a wide range of qualitative spatial representations proposed up to this stage, each concentrating on specific aspects related to specific tasks [241, 65].
A lot of research has been going on spatial knowledge representations and reasoning as people are increasingly reliant on description-based localization in query or dialog-driven geo-location services such as local search,

car navigation, emergency assistance or public transport planner services [370]. Many representations and reasoning techniques have been proposed that aim at representing the spatial knowledge underlying natural language place descriptions. Most of them represent spatial knowledge using spatial relations, which are usually restricted to one particular aspect of space.
Table 7 provides us an idea about the most common techniques present in the literature for spatial knowledge reasoning and representation. Many other formal models of relative path relations, such as the dipole model [323], the double-cross calculus [125], and others (e.g., see [323, 397]), are based on these techniques. Frank [121] differentiated three models for a cardinal direction calculus: a cone model and a half-plane model for point-like objects, and a neutral zone model for spatially extended objects. Other models, such as the internal cardinal direction model [248], are also available. [173] has also suggested a method for relative direction reasoning for various reference frames. There is also a review of existing formal representations in [127]. There have also been a few reasoning models proposed ([81, 139, 140]) which use calculi combining numerous families.
Recently, Chen et al. [64] proposed an extended place graph model to interpret place descriptions. Spatial information from text is presented in the form of nodes representing entities and labeled edges in the graph representing spatial relations. The approach focuses on spatial relation modeling to describe places relative to one another. Approximate location regions are determined based on spatial relationships. However, the interpretation of spatial relationships is context-dependent. The approach would thus benefit from a more in-depth knowledge of contextual factors.
Such proposed frameworks in the domain of QSR, particularly spatial relation models or associated calculi, are usually for objects with crisp, simple geometries and may not even be suitable for locations. Whether and how the frameworks contribute to spatial reasoning of human spatial knowledge of locations is yet to be investigated [63].

Table 7: Overview of spatial knowledge representations and reasoning approaches

Spatial knowledge representations and reasoning approaches	Purpose	Limitation
Qualitative calculi [298]	This technique provides us with a discrete representation distinguishing selected properties and efficient means of reasoning.	the real world is continuous and it is not clear from the beginning which properties matter.
Fuzzy set methods [120] (FSM)	FSM provide information about the degree to which an entity belongs to a particular class. defines several operations which can serve as the basis for topological operations.	membership functions are subjective and often lacks semantic justification.
Probabilistic methods [109]	Probabilistic methods usually deal with positional or measurement uncertainty and are best suited for modeling phenomena with measurable objectives.	Probabilistic models are based on clear semantics but may be hard to construct.
The Egg-Yolk model [79]	The Egg-Yolk model represents a vague region containing an interior part (yolk) representing definitive membership within the vague region and contained in an surrounding part (egg) which represents the maximal extent of the vague region.	Egg-Yolk model lacks numeric values makes it difficult to balance several crisping originating from distinct entities with indeterminate boundaries

However, modeling spatial relationships for computations in GISs and spatial databases remains a significant challenge due to the ambiguity and flexibility of the spatial relations in natural language ([118, 327]). Setting open questions in identifying natural language semantics aside, automated interpretations of place descriptions involves several challenging questions related to spatial knowledge representation and reasoning.The following section 2.4.2.1 discusses the limitations of the techniques mentioned above.

2.4.2.1 *Limitations*

Spatial relations are critical for both disambiguating place references and spatially anchoring places by restricting their probable locations. Spatial relations and place referents have enough information to convey locative information about places, especially when combined. One of the most significant concerns is interpreting and modeling spatial relationships from NL expressions, which is challenging for computers. Previous research has typically focused on spatial relations from specific contexts, such as spatial relations related to major cities. While such simplification often yields valuable results, the method of interpreting the query is fundamentally different from how a person would proceed. However, these models either require knowledge or are conceptual.

Therefore for spatial reasoning in NL, a more in-depth consideration of context is needed. Although consistency rules can be described logically, they do not always represent pragmatics from a linguistic perspective. Thus in NL, the use and interpretation of spatial relationships are both flexible and context-dependent.

Thus, it is essential to study the contextual factors and how they affect the meaning of NL spatial relationship expressions. Reasoning models based on logic types such as fuzzy, probabilistic, or defensible logic-based logic can also be considered to model the semantics of spatial relationships from NL expressions.

In order to generate contextualized models for answering spatial queries, Wallgrün, Klippel, and Karimzadeh
[375] propose identifying essential context features influencing the human use of spatial relation expressions. Cai [61] suggested a system for modeling geospatial data in the context of tasks and transportation resources. Yao and Thill looked at how contextual factors influenced the perception of proximity measures like near and far [403]. Most of these contextual features, such as familiarity with the location, financial and time budget, network connectivity, and personal characteristics, are challenging to extract from place descriptions. In this thesis, contextual features derived from location descriptions are preferable.
Even though no implementation is yet provided, existing literature shows that no service yet exists which could provide such an interface to users capable of respecting general context information and in specific vagueness of both relations (e.g., "near") and objects (e.g., extent of a forest) [63]. The significant point is that the understanding of vague relations is not exclusive to the relation in question but must be consistent with the entire domain theory. For instance, if we interpret "near" and "far" vaguely, we want to be sure that the interpretations are complementary; the more acceptable "near" appears to be, the less acceptable "far" should be.

Such consistent interpretation, not only consistent on the level of relations but in the context of a complete query, has been shown to improve recognition rates in the context of querying process occurrences from robot observations [211].
The use of logic axioms in relation semantics has resulted in introducing a standpoint, which is any fixed understanding of all definitions consistent with domain theory [42].

2.4.2.2 *Open Research Question*

Even though dealing with spatial relations is out of the scope of this thesis,yet the major constraint from the existing research is identified as follows:

1. There is a need for consistent interpretation of vague regions regarding contextual influences, thus it is necessary to define context and identify contextual factors that will aod in understanding and interpreting spatial information.

Despite comprehensive studies in geographic information science and geographic information retrieval, state of the art discussed above has put-forth several limitations and challenges, which require more effort and focus and will help in improving the performance and automation process for geographic information systems. Therefore, in the following Section 2.5 research gaps are concluded concerning spatial dimensions of information interpretation and querying.

2.5 RESEARCH GAPS

This chapter provided an overview of research on the geographic information system and its components, including spatial information extraction, representation and reasoning, and place-geo-referencing. Despite comprehensive work to understand crowdsourced place descriptions, significant challenges have been encountered, impeding the automation and analysis of natural language place descriptions.
Several exciting research contributions remain conceptual or are only partly automated and require some manual work. Without implementing a fully automated method, it remains hard to assess the value of contribution with respect to practical applications. It would be helpful if a diverse dataset of place descriptions at different levels of complexity would be developed in a community effort. This section summarizes the research areas and gaps that must be addressed to capture context-sensitive vague spatial information in human-generated vague place descriptions.

1. **Unstructured text:**
   Current spatial search is primarily limited to structured Spatio-temporal data, but the search should ideally be possible through vast amounts of unstructured spatial data gathered from social media and other web sources [167]. Because of recent developments in natural language processing and machine learning, subjective perceptions, feelings, and opinions can become novel search spaces, allowing for new research areas in geography and urban dynamics [34].

2. **Understanding notion of place:**
   Even though the humanistic notion of place is multidimensional and nuanced, we can not easily search for places beyond a few simple thematic dimensions (e.g., "cities with more than a million inhabitants"). Improved "palatial" models are required to incorporate the concept of place into geographic information systems. The contextual, ad-hoc, and mutable nature of the place is one of the challenges to place computing. To a large degree, the information retrieval community still neglects space and place, and much more GIScience efforts are required to make these perspectives more relevant to information search research.

3. **Need of context:**
   Contextual information is hardly exploited, i.e., places are processed individually and not related to one another. References that are not place names are most challenging to locate and typically necessitate considering contextual information [347]. Geospatial search entails the use of spatial words, which are often ambiguous and context-dependent. The definition of nearness varies according to context, and the disambiguation of the place name is challenging, especially for vernacular place names which are not present in a gazetteer. The search of the geographical domain is greatly influenced by scale. It is helpful to

organize content in hierarchies. That being said, spatial and thematic hierarchies pose an evaluation problem. In order to gather meaningful input, these hierarchies should be made more explicit. Likewise, the development of context-sensitive spatial relations has the potential to improve search methods vastly.

4. **Indirect or implicit references:**
In the retrieval of geographic information (as in other searches), queries often refer to instances of geographic entities by referring to their type. (e.g., [34], "the beach next to University of California, Santa Barbara" when referring to Goleta Beach). This method of indirect referencing necessitates spatial reasoning and geographic information, which goes beyond conventional co-reference resolution techniques. Place descriptions should therefore be interpreted by evaluating all pieces of information in the text concerning a geographical database or map.

5. **Dealing with entities other than named ones:**
Performance of current systems hinges on the availability of place names that are detectable employing named entity recognition, respectively geoparsers. Existing techniques either only perform toponym resolution for named entities, or they perform forward chaining, starting with geo-referenced named entities [64]. Approaches do not consider unnamed spatial entities or other colloquial forms of place description, e.g., the picnic spot. Thus, existing systems may identify the city of Bamberg within the phrase *post office near Bamberg main station*, possibly even *Bamberg main station*, but they cannot resolve the main place described, i.e., the post office.

6. **NL parser limitations:**
Disambiguation in toponym resolution is limited to simple scoring functions or relies on identifying geographic context, typically by explicitly mentioning either continent or country. Consideration of co-occurrence will fail to handle phrases that require the sentence structure to be considered. For example, consideration of co-occurrence fails to interpret *outside Europe, we only know of a single town named Bamberg* as Europe and Bamberg would be related, but the constraint 'Bamberg outside Europe' ignored.
The unavailability of training data challenges machine learning techniques [188], in particular considering comprehensive, unbiased coverage of the natural language, which is highly flexible. As virtually all modern geoparsers or named entity recognition systems rely on machine learning, their shortcomings complicate geo-referencing. For example, existing systems processing the phrase *The Sherlock Holmes pub in London* may tag *Sherlock Holmes* as a name or fictional character but miss the fact that a pub is described. It thus requires techniques to correct errors of geoparsers.

7. **Natural language understanding:**
More cognitive psychology research is required to illuminate the techniques and heuristics used in search behavior in physical and knowledge spaces, which will help us better understand how humans search for patterns in stimuli and memory. This knowledge may subsequently be applied to developing information systems that supplement and augment human search abilities.

8. **Quality of geo-referencing:**
   As used by Ross [188], quality is an umbrella term that captures different notions for performance and survival measures. Though open-source and commercial geo-referencing tools are available, their quality varies considerably. Practical benchmark tests and assessments are needed to support geographic information search effectively. For delivering more meaningful results, mainstream search engines need better topological and geographic knowledge bases. Furthermore, the system should be able to deal with scenarios where no query-able spatial entity exists.

Beyond these challenges, there also exist fundamental problems which are out of the scope of this research. First of all, place descriptions could also contain references to geographic locations in the form of historical dates, monuments, ethnicity, typical food, place type, and others [266]. Thus, it seems questionable whether one can truly solve the problem of interpreting place descriptions without solving natural language understanding in general. Also, existing geographic information systems are built on unambiguous, crisp, and metric geometries distinct from human conceptualization [63] and therefore offer little means to handle place knowledge [214]. It thus requires much more research to interpret place descriptions like the *picnic spot*. Based on presented challenges and limitations, a set of hypotheses for our current research is put forth in our next step, and a processing pipeline is presented to automatically interpret context-sensitive vague place descriptions.

2.6 hypothesis for current research

The idea for a basic research thesis in spatial representation and reasoning is motivated by the observation that VGI is not restricted to geo-referenced data collected by volunteers for the explicit purpose of contributing to a geographic information system; however, there also exists a virtue of implicit spatial information contained in place descriptions. To achieve this aim, techniques need to be identified that capture the semantics of entities and relations in vague place descriptions that allow spatial reasoning to make the information explicit, considering both the description and geographic context.

Nevertheless, given the current literature and its limitations, it is clear that interpreting and modeling the semantics of spatial relations necessitates a more in-depth consideration of context. Consequently, it is vital to investigate contextual variables and how they influence the interpretation of NL spatial relationship expressions. Furthermore, extracting knowledge from place descriptions using parsers is still limited in terms of accuracy and the type of information that can be extracted, necessitating manual intervention while extracting data. Indeed ML enables to improve the accuracy of place reference and spatial relation extraction at the cost of extensive training [63]. The same can be seen in [63], where interpreting and modeling spatial relationships from NL expressions is a core task, but the method needs a manual effort to obtain the desired information due to the restricted functionality of the existing parser, which impedes the automation process.

Thus, by considering all these factors, our research focuses on investigating means to factor in vague context information to improve the interpretation process for NL place descriptions by making reasoning an explicit step in interpretation. The fundamental research question we are considering for our study is: *Identifying to which extent reasoning about spatial and ontological properties of spatial entities is capable of overcoming problems with existing methods*.

For this purpose, the following hypotheses are stated:

1. **H1:** A more precise understanding of place in geographic information systems can be achieved by defining it in terms of context, i.e, by making implicit information explicit.

2. **H2:** By incorporating context, reasoning approaches to aid in automating the process of understanding place descriptions (using the presented computational model). It could improve the information extraction from unstructured text and can overcome the limitations of the existing parser.

3. **H3:** Considering the finest level of granularity possible for relation extraction aids in resolving indirect or implied information in place descriptions and can improve the parser's co-reference resolution problem.

4. **H4:** Explicit use of the type information in the description helps to interpret and geo-reference spatial entities in different languages and unnamed entities. This way we are able to improve NER performance as well.

5. **H5:** Query methods can be constructed based on semantic similarity and ranking methods.

6. **H6:** Using semantic similarity measure and clustering technique helps in finding type of place for non-spatial information present in the textual descriptions.

On a technical level, this research investigates the contribution of reasoning in interpreting the NL place descriptions. The aim is to provide a solution for existing limitations related to information retrieval and geo-referencing. Geographic information can be derived from structured and unstructured data by considering an additional layer of reasoning. Therefore, we have chosen to ignore most pieces of information provided by a parser and rely on shallow part-of-speech tagging only (classifying words as nouns, names, verbs, etc.). We are interested in learning how much information can be inferred from a given place description in conjunction with a geographic database that can be used to verify extracted information. Second, we are motivated to investigate how spatial and ontological reasoning contributes to interpreting spatial text – focusing on the entities themselves and disregarding syntactical information appears to be the more compelling and practical research methodology.

A solid geo-referencing system should be capable of dealing with the highlighted research issues. Thus, we offer a primary reasoning pipeline based on the proposed assumptions. The overall notion is based on Pouliquen et al. [289] who introduced concepts of shallow- and deep-geoparsing.

Following this motivation, a geo-referencing pipeline is proposed in Figure 7. Like the basic pipelines outlined before, it uses the TextReader to read in the plain text documents and eventually outputs a list of places with their coordinates using the geo-referencing component. The other functionality is broken down further into smaller components that allow for a more

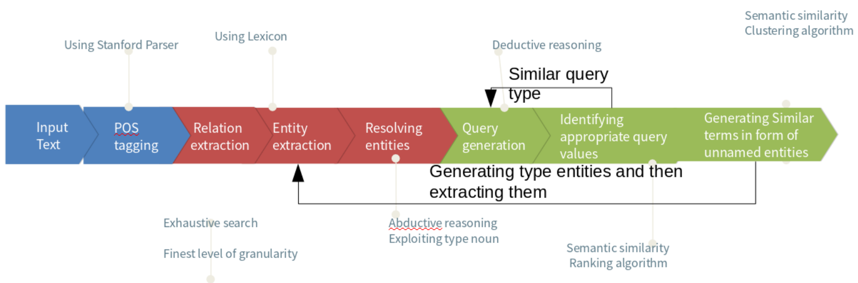

Figure 7: An improved georeferencing pipeline with reasoning as Core components

flexible and robust built. However, in this process, the entities are extracted and then first resolved by employing spatial type as the contextual variable, resulting in a spatial entity having a name and type. The triplets will be extracted, and using deduction will be inferred as probable queries. In case, of wrong or no results a ranking algorithm using semantic similarity will help us in generating similar terms related to the entity type and thus the query can be generated again. In case of non spatial entities semantic similarity based on OSM tags will provide us with terms similar to them in the form of type nouns. Technically it is also possible to execute the last two components in a loop. The last component will generate the type nouns which will then be identified as an entity and the pipeline processes normally. For suitable query tag, it will also provide a list of similar spatial types that will replace the wrong term at that stage and will execute the query again.

2.7 SUMMARY

This chapter describes relevant work that is related to the identified challenges of this thesis, i.e., incorporating context, information extraction and modeling, spatial reasoning and representation, place geo-referencing, and querying. It has been argued that Place descriptions provide a rich source of human knowledge about places, and approaches have been proposed for decoding spatial language from the text.

These presented approaches have certain limitations regarding how the human cognitive apparatus searches for information in physical space [102] and retrieves information from memory [357]. Understanding how humans search for information is arguably critical for designing better information retrieval systems, assisting in interaction design and visual analytics, and spatializing abstract spaces effectively [287].
Geographic information retrieval (GIR) has addressed computational problems such as spatial language processing and disambiguation of place-names [270, 187]. Still, the drawbacks of the existing rule-based or machine learning techniques (presented in Sections 2.5, 2.3.1.1, 2.3.3, 2.4.1.1) hinders the process of automating a geographic information system.

Based on the challenges (Section 2.3.1.1,Section 2.4.1.1,Section 6.3.5, Section 2.3.3,Section 2.2.5) various research areas Section 2.5 are identified, and on basis of them set of hypotheses are formulated and a geo-referencing pipeline has been presented, which provides the basis of this Ph.D. dissertation. At the very least, the pipeline will assist us in dealing with these four problems. First, the word refers to a real-world location that is not inside the geographical scope of the search and so is not taken into account. Second, a similar place exists, but it is not (correctly) marked on the map and therefore cannot be discovered. Third, numerous search results have been presented, and an inappropriate entry is chosen. Finally, an entry is discovered when the actual phrase is employed in a non-spatial context.

The workings of the procedures employed in the pipeline are discussed in depth in the following chapters.

3 A COMPUTATIONAL MODEL FOR REPRESENTING CONTEXT AND VAGUENESS OCCURRING IN HUMAN-GENERATED PLACE DESCRIPTIONS

This chapter is based on the content from a published paper by Wolter and Yousaf [398] in COSIT 2017. The contributions of the corresponding second author include problem conceptualization (through discussions with my supervisor), literature review, identifying knowledge to be modeled, the design and implementation of the computational model.

Places are frequently linked to uncertain spatial meanings and cannot automatically be represented by existing GIS and spatial databases, whereas the latter usually are built on explicit, crisp, and metric geometries [63],[188],[260],[9]. They argue in [91] that the idea of place is a challenge at best to the GIScience. In both spatial and semantical literature, places appear to be almost described as undefinable and remain an ill-defined concept of database records in computer science and information science, rather than a well-defined and useful term. Although the concept of place has been studied extensively in GIScience (e.g., [91],[43],[370],[393]), there is currently no widely adopted computational data model allowing reasoning and inquiring of place knowledge from a human perspective.

This chapter provides research linked to the thesis's first significant task, which is conceptualizing and modeling place as a cognitive construct. Current GISs and spatial databases have been argued to represent space but not of representing place. Existing place definitions and models have been presented from the literature, from both viewpoints, with and without contextual details. This chapter presents a computational model representing spatial knowledge occurring in place descriptions and analyzes how context information shapes a place description's meaning. It outlines a computational model to represent vague spatial knowledge and context occurring in human-generated place descriptions.
Findings from the paper [398] are re-organized in this chapter, and the presented concepts are employed for this research in Chapters 4, 5 and 6 respectively.

3.1 MOTIVATION AND GOALS

Substantial amounts of information currently available include references to places on earth like news stories, travel guides, blogs, social media, and photo captions, as well as verbal communication and search requests [188],[291]. Generally, such information has been organized as structured data. However, increased volumes of data are available for search and retrieval in the form of unstructured text [188]. In both written and verbal forms, natural language uses place names and descriptions as references to locations meaningful for a conversation. Besides, they also provide shared implicit knowledge between speaker and recipient, which encourages the development of spatially-aware search systems to facilitates the user's geographical information needs [188],[291], indicating that place is a fundamental concept in geography and plays a crucial role in nearly every aspect of human inquiry [43].

To make a spatially-aware system, the initial step is to deal with the conception, collection, analysis, and interpretation of geographical information [145],[400],[285]. In [145], a geographic information system is put forth as an approach to collect, manipulate, and represent spatial/geographic information. Place-based GIS enables digital systems to provide a human-centered representation of the geographic world by understanding the notion of human meaning [281]. Although place-based investigations of human phenomena have been conducted in the humanities and social sciences during the last years, this notion lately transgressed into GIScience, which can be regarded as a multidisciplinary and multi-paradigmatic field [50]. Place-based analysis has rapidly gained popularity in Geographic Information Science (GIScience) in recent years in an attempt to understand the locally rooted human meaning [51].

In order to use the spatial information in the same way humans do, the concept of place and how it can help in interactions between humans and computers is of fundamental importance [368, 84, 361, 51]. For many years, the word place exists in everyday language, and the concept has been studied in many different fields which includes philosophy, psychology, and geography [63]. Because of the growth of interactive systems based on natural language, understanding the cognitive aspects behind natural language place descriptions is an initial step towards formalizing the concept of a place to be used in decision making and spatial reasoning [368]. However, human spatial concepts like the concept of place can not be transferred to geometric foundations of spatial databases and information systems right away, which hinders progress in place-based information systems [291].

Many existing information systems address places as points of interest (POI), usually without scope and a predicament for commercial entities, in an ad-hoc and lousy manner [291]. A considerable number of search requests in existing location-based systems explicitly provide geographic search words such as place names or less ad-hoc taxonomy, and place georeferences to points or occasionally polygons in a coordinate system [130],[17],[361]. However, they cannot portray the innate geographical concept of a place that helps with everyday human experience and communication, and can only deal with the places and types present in the existing taxonomies. Moreover, they allow traditional mapping but leave no room for place concepts that are vague or require more information.
According to some researchers, the concept of place may be too vague to formalize, except under restricted conditions [147]. Discrepancies thus arise between how humans conceptualize space and how entities are represented in a geographic information system (GIS) [398]. These discrepancies arise because of the role of place and different ways to represent it [148] and the influence of the context indicating that translation of human spatial concepts has proven surprisingly hard [291] .

Thus, by observing the shortcomings of standard geographic environment conceptualizations with respect to place modeling and in developing a geographically intelligent system in section 2.2.5, leads to the motivation to identify what needs to be represented about the place and which parts of the human context ascribed to people's understanding of place, which includes place names and types, semantic descriptions, or even semantic relations between physical entities and places [281],[282]. The main challenge is to provide an interface to users capable of respecting general context information and in specific vagueness of both relations (e.g., "near") and objects (e.g., extent of a forest). Setting open questions in identifying natural language semantics and automated interpretations of place descriptions. Several works address the interpretation of place descriptions, motivated by either question regarding the semantics of spatial language or by prospects on applications in human-centered computing, for example, exploitation of volunteered geographic information (VGI) (e.g.,[6]).

Our work is aimed at a computational model for the latter, a formal knowledge representation, which provides the basis for reasoning algorithms, and allows natural language place descriptions to be interpreted by identifying the described place within a spatial database such as OpenStreetMap (OSM). Our overall aim is to develop an automatic system applicable to various settings and analyze how context information shapes a place description's meaning. We consider natural language descriptions aiming to communicate location information to an unknown recipient. In particular, we are interested in these research questions:

1. To which extent can and should context be incorporated into the model to allow for sensible interpretation of natural language place descriptions? What will be the influence factors/properties that can be considered while integrating different statements expressed in our spatial representations?

2. How can a spatial relation expression be interpreted? What is an appropriate semantics to represent vague knowledge inherent to relations (e.g., near) or entities themselves (e.g., extent of a mountain)?

3. How can a set of spatial entities describing a location be efficiently matched against a spatial knowledge base to identify location candidates?

We restrict our consideration to descriptions that refer to objects represented in the spatial database using spatial language, either directly ("the *campground* near mount Foo"), or as sub-region of larger entities ("*northern beach* of the lake near mount Foo"). Also, we disregard the interpretation of place names; for example, we do not aim to interpret "students district," which may stand for a city district close to a university containing affordable restaurants.

Before embarking on the interpretation of place descriptions, we have to commit to a definition of place [398]. The notion of place and how it can be defined has been discussed comprehensively [43, 91, 395, 370], yet there is no definition capturing human intuition, leading to a fully specified computational data model [398]. The contribution of this chapter is to discuss contextual influences and to propose a computational framework that can capture incomplete and context-sensitive information about places. First, we present computationally motivated models of place and context, then we outline our overall model.

The remainder of the chapter is structured as follows. Section 3.2 presents a computational model based on constrained optimization to match place descriptions to information provided by a GIS in an automated manner. Section 3.2.3 summarizes and concludes this chapter.

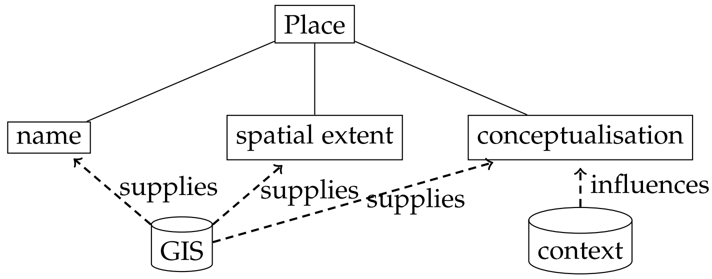

Figure 8: Computational model of place based on name, spatial extent, and conceptualisation.

3.2 COMPUTATIONAL MODEL

This section analyzes how context information shapes a place description's meaning and outlines a computational model to represent vague spatial knowledge and context occurring in human-generated place descriptions.

As seen in related work, discrepancies thus arise between how humans conceptualize space and how entities are represented in a geographic information system (GIS). A single entity in a GIS may correspond to several places (for example several barbecue places in a single park), or vice versa (for example adjacent meadows which constitute a place for playing soccer). [370] contrasts definitions of place which involve elements place name, spatial extent, and – in case of gazetteers – a type of place, often defined using some form of ontology. We argue that places can obtain their meaning from actions that can be performed within their spatial extents and therefore a definition of space should include the conceptualization, i.e., the process which associates a particular concept with a spatial area.In our prototype, we are just considering nouns that reflect the type or spatial extent of place, and we are not considering verbs. Our application-specific definition which is illustrated in in Fig. 8 contains three elements:

1. Places are designated by names, either a unique label (e.g., "Eiffel tower") or a circumscription (e.g., "picnic place"). We assume unique labels to by available from a GIS and we treat any circumscription as a free-ranging variable that can represent any type of place.

2. Places are defined by their spatial extent. Although some places have indeterminate boundaries we assume location information to be provided from a GIS. In case spatial extents are not directly provided by the GIS, either a spatial extent of some GIS entity is used as upper approximation or a single location representing a prototypical location of the place is computed.

3. Places are defined by their conceptualization. While a GIS provides information about type of objects, it may require a context-sensitive interpretation of type information. We are particularly concerned with handling spatial context variables. For example, if a specific entity is referred to as a *large lake*, then entities of the same type and similar extent are reasonable instantiations of *large lake*, too.

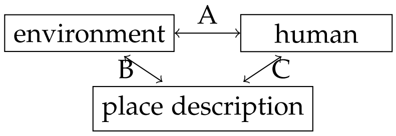

Figure 9: Contextual influence in interpretation of place descriptions.

3.2.1 *Towards an Operational Definition of Context*

To explore context variables, we must first comprehend the idea of context. Actually we need an answer to the question **"what is context?"**. Context shapes the meaning in all communications and presents a multidisciplinary topic. A lot of definitions of context exist which all refine the umbrella term *conditions and circumstances*, often differently. [40] argue that context is inherently domain- or task-specific and thus no generic definition of context can be formulated. In order to arrive at an operational definition that can be realized as a computational model we employ a bottom-up approach of identifying *context variables* which are elements influencing what the most reasonable interpretation of a place description is. We have chosen the term context variable since variables often represent unknowns: **Context is not explicitly given but must be inferred from the pieces of information available.**

We distinguish three classes of context variables in interpretation of place description: the *environment*, the *human* who generates a place description, and the *place description* as a linguistic phrase. This classification adapts a characterization of context in performing map-based tasks by [126] to a text-based task. Our linguistic context implied by the place description corresponds to their map context, yet we separate objective environmental factors from all cognitive factors. Let's have a detailed look on our classification which is depicted in Fig. 9:

environment The environment defines the physical domain in which place descriptions are interpreted. It provides input to human conceptualization (link A in the figure), that is, what mental representation is built which then serves as basis for generating the place description. Considering actions a particular entity allows a human to perform contributes to interpretation, for example a river can be followed, a hill enables going further up or down, etc. Second, entities in the environment can define reference frames which need to be identified in order to interpret a spatial relation occurring in a place description (link B in the figure). For example, "in front of X" may refer to an intrinsic front if provided by entity X provides it (e.g., a building with a designated front), yet it may also refer to a route described, i.e., it refers to a location just before reaching X. Moreover, the environment shapes the meaning of relations by providing information about scale. For example, "near Egypt" implies a different interpretation of nearness in absolute distance terms than "near the central station".

human Cognitive principles shape mental representation of the environment (link A), they also influence how places are described verbally (link B). Moreover, the intention of the human to generate the place description and her model of the recipient are important. However, we have to disregard these aspects for time being, since we assume only the place description and environment to be known.

place description The place description provides a linguistic aspect which allows us to draw conclusion about cognitive influences (link B). For example, a term such as "north-northwest" introduces a finer level of granularity in direction than "north". Also, the descriptions provides (qualitative) information about entities in the environment (link A).

3.2.2 *Operationalizing Context Variables*

We first review the requirements a computational model needs to satisfy, both general properties and those arising from context variables. With respect to general properties of the spatial knowledge model we first observe that *different aspects of spatial knowledge* (characteristics of entities, distance, etc.) can occur in a single place description and must be jointly expressible in our model. Second, several *qualitative relations and concepts need to be considered as being vague*, that is, there exists no determinate boundary of acceptability. Representing entities such as a 'city centre' also requires handling of vague concepts both in sense of conceptual and spatial boundaries. Our computational model must thus be able to handle partial satisfiability of a entity or concept and it should aim to determine the most plausible interpretation my *maximising satisfiability*. For example, if one is referring to a park near the city center, a reasonable strategy would be to look for the park closest to a location most certainly inside the city center.

With respect to context-dependent modeling, context variables influence the *mapping of GIS entity types to concepts*, for example, reference to a park for picnic or for playing football poses different space requirements on the park. Technically speaking, environmental features establish conceptual entities referred to in a place description. Our model thus needs to include variables that represent conceptual entities and represent rules that determine the mapping. The same technique is also applicable to identify *frames of reference* environmental features offer. Use of words may influence the *granularity* of concepts within a given hierarchy, that is the semantics of entity "city" depends on whether it is contrasted to either "city on river" or "city in country". Our model thus has to provide means for interpreting relation terms with spatial relations. Last but not least, environmental features and actions referred to in the place description introduce *scale* information, for example "near" depends on scale information implied by mode of transportation (e.g., driving by car vs. walking) or a reference object (e.g., "near Egypt" vs. "near the signpost").

In consideration of the requirements discussed, we propose to tackle interpretation of place description as a constrained optimization problem using discrete and continuous variables. To circumvent the long-standing (qualitative) spatial reasoning problem to reason about several types of entities in a joint manner.

This model allows us to represent variables that are either Booleans (to compose alternative interpretations) or numerical, for example to represent scale and spatial locations. Moreover, we are able to state dependencies among the distinct valuations of context variables by modelling a constraint which is part of an overall conjunctive formula representing the place description. We propose to employ the contextual factors (including entity types, affordances or activity words) in an explicit form depending upon the information provided in the place description to model the partial satisfiability of a vague entity, this allows us to optimize for highest likely possible option in a decision taken during reasoning.

The overall model we propose is based on the approach by [370] which starts by considering a place description as a set of prepositional phrases (PPs) obtained by means of parsing natural language. We extend the model by introducing an explicit step to interpret a natural language entity as a combination of name and contextual variables, which are introduced as non-spatial variables.

For example, in a natural language place description," St. Catherine is the largest protestant church, located in the city center at the entrance to the Zeil, the central pedestrian street for shopping," the presented approach allows us to extract spatial relations, entities, and contextual concepts. The place definition presented in this chapter allows extracting spatial entity like St. Catherine, Zeil, along with the type nouns (church, pedestrian street, city center), concepts like entrance, and relations from is-a to in and at. The type noun provides granularity information and identifies the concept or form of spatial object being described, and points out the mental representation of how different entities are linked by humans while describing places. On the other hand, is-a relation represents the association of different spatial entities with concepts and other entities providing the context using place descriptions.

We have identified context variables necessary for interpreting a class of natural language place descriptions and highlight the factors that a reasonable approach should consider while extracting geographic information. The following Section 3.2.3 discusses how the proposed approach accounts for the dimensions of places and covers all the essential elements required for a functional definition of place.

3.2.3 *Accounting for cognitive aspect of place*

This chapter introduces a computational model for representing spatial knowledge in place descriptions and examines how context information influences the meaning of a place description. Before arriving at an implementation, we have committed ourselves to definitions of concepts such as context or place. We have chosen a bottom-up approach of identifying a set of context variables. Here, we will discuss the model's specifics and focus on how it accounts for the dimensions of a place that has been established in the literature.

Agnew [184], [11], [98] identifies three dimensions of the place(which has been discussed in section 3.2.2) and are used by researchers (for example [291]) to account for the cognitive dimensions of place. We show here that a place conceptualization can fully capture Agnew's three dimensions in terms of objects resulting from contextual variables identified by our computational model. They are compatible with the approach to define place within three elements. Agnew also suggested that places should be considered concerning other places rather than as "bounded, isolated entities," implying that place's spatial extent of objects does not necessarily have a boundary.

In natural language, location is usually conveyed by place names (toponyms; "Germany") or place-identifying count nouns (object types; "the bar"). These two components have been shown to serve different purposes in knowledge classification and search [24] as specific or generic elements of the where facet [332]. Toponyms implicitly refer to terms with no guaranteed semantic sense [76] and are ambiguous [340]. As a result, reasoning regarding the relationships to other locations usually necessitates identifying a clear referent and awareness of their type [228], both of which can be offered by the environment and human description of place a part of the conceptualization element of our model.

On the other hand, place-like count nouns convey general knowledge about a location and its properties and are related to locale. Locale captures the properties of a location that are both perceivable and relevant to some group in its broadest context. As a result, Locale is linked to experience, affordance, and function. Additionally, several researchers focusing on the semantics of place from a formal perspective have stressed the importance of modeling place activities, mostly from the perspective of affordances (for example,[191],[213],[320] and [321]).

Agnew's triad's final component, sense of place, is crucial to place concepts. The relation that a group of people has to a place[84] or the feelings people ascribe to a specific place are two examples of a sense of place.

The sense of place is subjective and differs by culture and experience, not only in terms of the place in question but also in its relationships with other places [86]. In the presented computational model, conceptualization based on three defined contextual variables helps infer different functions, activities, events, and linguistic aspects that help us understand place as a cognitive concept.

To summarize this review of Agnew's three dimensions of location, places share specific common characteristics, even if their understanding is not universal, may be influenced by culture or language, and does not have to be permanent. This section shows how the dimensions of a place identified in the literature are consistent with our derivation of place.

3.3 SUMMARY

This chapter presents a definition of context and defines the notion of place based on context for better understanding and retrieving place information from textual descriptions of locations. The provided computational model overcomes the shortcomings of previous definitions in terms of the types of data that might be modeled and acquired how they can be modeled.

Three aspects (place name, spatial extent and conceptualization) are specified in our definition of place to aid in conceptualizing place. Context influences linking a particular concept with a specific spatial region. The presented approach has identified three context variables (environment, human, and place description), which aid in place description interpretation by clearly capturing the information contained in NL place descriptions.This provides the basis for inferring context and then, eventually, employ the resulting purely spatial description for georeferencing the place described.

Furthermore, the theoretical model presented in [291] based on the core concepts can also be integrated with our presented model as the five core concepts location, field, object, network, and event are a part of our computational model where the place with a name or type and spatial extent represents an object. In contrast, the entity in reference to the object is the location, for example, "I am at a post office near the train station," where the near train station is a location and the post office is the object we are looking for. The remaining three core concepts are covered under our conceptualization module where environment, place description, and human perception provide us with the details of events occurring at that place, explaining the

network like the road or utility lines between the cities.
The model has been found to be significantly related to the goals of this thesis for a more comprehensive understanding. Further additions are anticipated, taking into account instances in which verbs, in addition to nouns, might be used for a more precise understanding and enhancement of our system. **In summary, using the proposed model's categorization, we can manage and analyze the place data.**

The next chapters will explain how the contextual variables, especially place type, have been exploited in improving, capturing and analyzing geographic information based on the conceptualization of place as context.

4 SPATIAL INFORMATION EXTRACTION

This chapter is based on the contents of previously published papers by Yousaf, and Wolter [405], [404] in GIScience 2018, including an explicit reasoning framework. The contributions of the corresponding (first) author include problem conceptualization, literature review, methodology, implementation, experiments, and paper writing.

The tremendous volume and complexity of the unstructured data available in the form of place descriptions have changed the paradigm of computational capabilities of the conventional GIS. To process human languages using NLP techniques, several tasks are considered high-level, like machine translation, question-answering systems, information extraction, and natural language understanding. Out of all these, the information extraction process captures the most valuable and relevant information in a structured format expressed in written or verbal place descriptions. Such knowledge eases developing an automated system capable of interpreting place descriptions by inferring new spatial relationships between various entities or by identifying logical contradictions based on existing ones.
However, various information extraction techniques have been proposed concerning spatial information, which has been introduced in Section 2.3.2 of Chapter 2.To date, the data extraction process has been automated using limited knowledge and techniques that adequately capture the spatial semantics of objects and relations based on the context and implicit information available within the text. This dissertation aims to propose methods for extracting relations between spatial entities and interpreting named and unnamed spatial objects by collecting implicit contextual information and using it for reasoning, Geo-referencing, and querying. Existing methods, by comparison, are lacking in these tasks.
This chapter focuses on comprehending and understanding the meaning of the place descriptions by creating an organized, explicit depiction of entities and their relation to each other. A flexible reasoning method is developed based on exploiting implicit information and representing it in declarative statements. The aim is to utilize logical statements as an intermediate representation to over-generalize knowledge conveyed in a given text place description. The experiment results show that the approach is dependable and robust in dealing with implicit information and generating additional spatially relevant triplets.

The article's ([405],[404]) findings and experiments are re-structured and will be presented in a subsequent chapter 7, respectively.

4.1 MOTIVATION AND GOALS

Automated text interpretation is still challenging, mainly because of language parsing, ambiguous names, and human conceptualization. While the ambiguity of named entities can be tackled by considering any object with a matching name found in the database and then applying ranking techniques based on geographic scope [21], there are no secure solutions to tackle failed attempts to parse a piece of text. Ambiguity resolution can be regarded as a task of reasoning since the goal is to identify a single interpretation from a set of jointly agreeable candidates with all information given. For this purpose, the idea of how humans conceptualize a particular situation is a guide for making computers acquire spatial knowledge. We argue that many intuitively phrased sentences are ambiguous and require concept-level reasoning, i.e., reasoning based on mental representations even beyond Kuhn's core computations [322]. Traditional learning-based or rule-based techniques are insufficient to handle the volume and dimensionality of unstructured data [236].

Consequently, this chapter presents a flexible reasoning framework for extracting geographical information to interpret descriptions of places, based on the research question devised and presented in section 2.3.3.1. The chapter is motivated by the concepts identified in figure 7 for relation extraction in Chapter 2. To that end, this chapter investigates the concepts of exhaustive search and granularity and suggests novel constraint specification problems using logic programming techniques. The framework ensures exploiting implicit information present within the sentence and generating all possible relations between the spatial entities. Moreover, experiment findings indicate that the presented approach is reliable in identifying maximum spatial relevant triplets and in providing solution to parsing issues like co-reference resolution and dealing with indirect references.

Overall, this research thus investigates to which extent reasoning about spatial and ontological properties helps in spatial information extraction and is capable of overcoming problems with natural language parsing and shortcomings of existing approaches discussed in Section 2.3.2 of Chapter 2 including extensive training data sets,error propagation, dealing with contextual information and capturing semantics or inferring new relations.

Our method's main idea is to employ logical assertions as an intermediary representation that over-generalizes information in a sentence. Afterward, implausible interpretations are pruned using the presented reasoning approach (SORS).

The remaining parts of this chapter are organized as follows. The reasoning framework, including an illustration of the problem, is explained in Section 4.2. Section 4.3 presents a discussion for the developed framework and concludes this chapter.

4.2 SPATIO-ONTOLOGICAL REASONING

This section outlines an approach for information extraction from text which does not rely on natural language parsing but employs a simple part-of-speech tagging and applies spatial and ontological reasoning for interpretation.

We seek to identify named and unnamed entities in a piece of text. While geo-referencing named entities consider names and spatial relations to other entities [21], dealing with unnamed entities presents a special case that can only exploit spatial constraints and maybe type information. We can thus regard both cases jointly as tasks of ambiguity resolution.

One approach is geographic scope resolution which allows potential interpretations to be restricted to within a known scope. In another one, Richter et al. [305] consider granularity effects caused by object types. They state that knowing the finest possible level of granularity with respect to a general ontology of spatial entities helps resolve place descriptions. Both ideas can

be integrated by attuning the semantics of relations and queries to focus on results that fit a scope indicated by a geographic entity's type and location in the same text. For example, the semantics of "near" can be approximated according to objects' granularity and geographic scope.
Exploiting such context information presents a chicken-and-egg problem: information obtained by resolving entities is to be employed simultaneously to resolve the entities. As a result, this motivates an approach using logic programming techniques since dependencies can be expressed declaratively, abstracting from algorithmic realization. The declarative representation are regarded as a constraint satisfaction problem, as explained in previous Chapter 3.
A solution to the CSP is obtained when all variables are geo-referenced by matching them to a spatial database. Likewise, we employ an ontology-like representation to augment the semantic representation of words (the lexicon). However, we have chosen not to employ formal ontology techniques. First, the truth semantics of classical ontology languages are binary, i.e., entities belong to a certain class or do not. For spatial entities and concepts, such crisp classification may be hard to achieve, and concepts may vary across individuals. Instead, concepts or relations like 'near' may be more adequately represented using a semantic capturing vagueness, e.g., using Fuzzy or probabilistic models. Second, existing ontology languages do not support the spatial domain and manifold spatial relations to the extent required to empower spatial reasoning for computing likely interpretations of a locative phrase.
Moreover, we found that no freely available parser was able to resolve references in the text correctly (see examples depicted below in Figure 11). A wrongly identified reference can efficiently inhibit the correct interpretation of a sentence. As our experiments discussed further below reveal, wrongly identified references are a common problem. By contrast, an unidentified reference can potentially be inferred from context. For every relation, a term is constructed combining any word (noun or named entity) appearing before the relation with any word occurring after the relation. The designator for each relation is retrieved from the lexicon e.g., in(park, town) and of(town,Bamberg). For every noun, an ontological "is-a" relation is generated in reference to any other noun or named entity, e.g., is-a(Bamberg, town). The basic idea of our approach is thus to generate all candidate interpretations of references and then apply reasoning to single out the most likely interpretation.

4.2.1 *Processing Pipeline*

Both phases of spatio-ontological reasoning, i.e., generation and pruning phases, rely on the same sources of information:

- the OSM database as a geographic database providing information about entity names, their type with respect to the ontology, and associated geographical information
- a lexicon comprising all nouns that represent geographic entity types and all spatial relations (however my lexicon is never complete and that is why we need a query method. Moreover, we will discuss the issues of incomplete lexicon later.)

In the *generation phase* (see Fig. 15 for an example), we process a sentence as follows:

	input	Bamberg	is	a	town	north	of	Nuremberg.
1.	**POS tagging**	Bamberg:NE			town:N	north:R		Nuremberg:NE
2.	**named entity resolution**	{ID0, ID1, ...}			town:N	north:R		{ID8, ID9, ...}
3.	**ontological annotation**	{ID0, ID1, ...}			settlement	north:R		{ID8, ID9, ...}
4.	**logic program generation**	$(\mathrm{isa}(\mathrm{ID0}, '\mathit{settlement}') \land \mathrm{northOf}('\mathit{settlement}', \mathrm{ID9})) \lor$ $(\mathrm{isa}(\mathrm{ID1}, '\mathit{settlement}') \land \mathrm{northOf}('\mathit{settlement}', \mathrm{ID9})) \lor$ $\ldots \mathrm{northOf}(\mathrm{ID0}, \mathrm{ID8}) \land \ldots) \lor \ldots$						

Figure 10: Example of processing steps in generation phase of information extraction (NE: named entity, N: noun, R:relation, ID:denotes reference to objects in OSM database)

1. Perform part-of-speech (POS) tagging by applying named entity recognition using the geographic database and checking for occurrence in the lexicon. All recognized words are labeled with their category, all other words are discarded. To handle composite expressions of several nouns, (e.g., “art gallery”), nouns immediately following one another get joined and treated as a single noun. The categories of nouns considered for this step include NN, which stands for singular nouns, while NNP represents proper nouns present in the text descriptions, and NNS stands for plural nouns. In addition, IN represents the prepositions present in the text while VBZ identifies the verbs, and DT is for determinants. The following are some of the combinations that we explored for this step:
   $< \mathrm{NN} >< \mathrm{NNP} > | < \mathrm{NN} >< \mathrm{NN} > | < \mathrm{NN} >< \mathrm{NNS} > |$
   $< \mathrm{NNP} >< \mathrm{NNP} > | < \mathrm{NNS} >< \mathrm{NNP} > | < \mathrm{NNS} >< \mathrm{NNS} >$
   $< \mathrm{NN|NNP|NNS} >$
   $< \mathrm{IN} >$
   $< \mathrm{VBZ} >< \mathrm{DT} >$
   In case of ambiguities at this or any later point, all possible options are stored.

2. For all named entities, possible interpretations from the geographic database are retrieved. For example, in case of Bamberg, we would obtain an OpenStreetMap entity referring to the city of Bamberg, depicted as ID0 in Fig. 15, and to the corresponding district of Bamberg, ID1, both for Bamberg, Germany and for Bamberg, SC, USA (creating ambiguity in the extracted information).

3. For all nouns ontological type information is obtained from the lexicon. Nouns are then replaced by their ontological type. Every noun is assumed to either represent an unnamed entity (e.g., “**park** in the town of Bamberg”) or a type designator for another noun or named entity (e.g., “park in the **town** of Bamberg”).

4. Possible interpretations are determined as disjunctions by compiling interpretations of words and references of relations:
   - For every relation, a term is constructed combining any word (noun or named entity) appearing before the relation with any word occurring after the relation. The designator for each relation is retrieved from the lexicon.

- For every noun an ontological "is-a" relation is generated in reference to any other noun or named entity, e.g., is-a(Bamberg, town).

In the *pruning phase* every conjunctive term generated in the generation phase is processed individually, see also Fig. 11 for illustration. A term gets discarded if

- a single noun occurs simultaneously in a "is-a" and a spatial relation, i.e., it would represent ontological information and an unnamed entity simultaneously,
- a noun or named entity in the input is not contained in at least one relation,
- or the grouping of relations violates word order in the input sentence. We disallow for relational statements $r(w_a, w_b) \wedge r'(w_c, w_d)$ if the position in the sentence (denoted Pos()) is in crossed order, i.e., it violates $\text{Pos}(w_a) < \text{Pos}(w_c) \Rightarrow \text{Pos}(w_b) \leqslant \text{Pos}(w_d)$. For example, in "Bamberg is a town north of Nuremberg, on the river Regnitz" interpretations containing $\text{isa}('Bamberg',' river') \wedge \text{isa}('town',' Regnitz')$ get discarded.

After the pruning phase, we search for the conjunctive term which can best be satisfied. This means, for unreferenced nouns a suitable instantiation from the geographic database is searched that agrees with the relations– agreement is measured gradually and summed up. Also, relations between named entities and/or referenced nouns are tested. In case of the example,"Bamberg is a town north of Nuremberg, on the river Regnitz", we would only find for the entity representing Bamberg on river Regnitz a matching entity Nuremberg such that Bamberg is located north of Nuremberg. The ontological constraint saying Bamberg is a settlement would only be fulfilled for the city of Bamberg, not the administrative region. We thus arrive at the desired interpretation.Thus Exhaustive search (following the concept of mental model theory) contains all correct interpretations by construction, but also several statements not following from the input text and some incorrect ones as well. The reason for getting all correct interpretations is based on the following assumptions:

- I assume that by identifying all nouns and relations between them we generate all possible options,
- secondly, by identifying the type of noun using is-a relation, we make sure that the identified noun is a spatial entity,
- lastly, we generate the triplets by considering all the spatial relations between entities

4.2.2 *Example:Relation extraction using spatio-ontological reasoning*

"$\overbrace{\text{St. Catherine}}^{SC}$ <u>is the</u> largest protestant $\overbrace{\text{church}}^{C}$ located <u>in</u> the $\overbrace{\text{city center}}^{CC}$ <u>at</u> the $\overbrace{\text{entrance}}^{E}$,
<u>to</u> the $\overbrace{\text{Zeil}}^{Z}$ the central $\underbrace{\text{pedestrian street}}_{PS}$ for $\underbrace{\text{shopping}}_{SP}$."

parser output	is-a(SC, C), to(C, PS)
generation phase	is-a({SC}, C, CC, E, Z, PS, SP), in({SC, C}, {CC, E, Z, PS, SP}), at({SC, C, CC}, {E, Z, PS, SP}), to({SC, C, CC, E}, {Z, PS, SP})
pruning phase	
ontological	is-a({SC}, ~~{CC, E, Z, S, SP}~~, C),
spatial	in({SC, C}, {CC, ~~{E, Z, PS, SP}~~}),
	at({~~SC, C~~, CC}, {E, ~~Z, PS, SP~~}),
ordering	to({~~SC, C~~, CC, E}, {Z, PS, SP})

Figure 11: Example of generating and pruning spatial relational statements using spatio-ontological reasoning.

Figure 11, provides an detailed overview of the the relations extracted by the parser and by the generation method.For clarity of presentation, no entities were replaced by OpenStreetMap references and no nouns were replaced by ontological types. We write $r(\{n_1, n_2\},$
$\{n_3, n_4\})$ as shorthand notation to denote that all four interpretations
$r(n_1, n_3), r(n_2, n_3), \ldots$ are considered. As can be seen, the parser identifies that 'St.Catherine' is the name of church, but it does not make the relation between the entity church and city center explicit, as well as is unable to identify the relation between 'St.Catherine' and city center. Also, the parser commits wrongly to claiming church to be the pedestrian street. By contrast, exhaustive search introduces all correct interpretations by construction, but also several statements not following from the input text. Applying ontological reasoning one immediately identifies that only named entity 'St.Catherine' is of type church. Spatial reasoning reveals, for example, that church is not pedestrian street but 'St.Catherine' is the same church and is linked with city center and entrance.
Overall, applying "Exhaustive search," we perform a simple analysis of input text, generating many interpretation candidates (all possible declarative sets of statements against each entity present in the sentence) and then applying reasoning to discard the inconsistent ones.
The following section 4.3 will provide an insight into how the presented approach can deal with existing parsing issues as well as bits of helping us in dealing with implicit information.

4.3 DISCUSSION AND SUMMARY OF THE CHAPTER

This chapter presents a framework for identifying informative triplets by capturing the contextual dependencies. A reasoning approach, rather than a machine-learning approach chosen here for extracting triplet relations that does not depend on the parsers for triplet extraction, only uses shallow parsing, i.e., parts of the speech tagging component. Furthermore, the approach employs an ontology-like structure instead of a formal ontology. The study compares the number of ambiguities introduced by our over-generalizing method of information extraction to wrongly identified references by the parser. Also, we are interested to learn what kind of spatial and ontological reasoning is required to interpret the output of our approach.

In addition to the above example, some candidate interpretations generated by an exhaustive search that are not valid interpretations of the input text take more effort to reject. In case of *The Historical Museum of Bamberg is a museum located in the Alte Hofhaltung next to the city cathedral*, the interpretation in('museum,' 'city cathedral') cannot easily be rejected if the geographic database also includes a museum in the city cathedral. If the unintended reference is accepted during the search for the most likely interpretation, then order constraints inhibit any further connection to the named spatial entity "Alte Hofhaltung" (Old Court). So in this case, we are relying on preferring the more extensive set of jointly possible interpretations that involve in('museum,' 'Alte Hofhaltung') and next_to('Alte Hofhaltung,' 'city cathedral') over just in('museum,' 'city cathedral').

The evaluation confirms, that SORS is capable of hypothesizing relevant pieces of information. The results show that in 60 percent sentences, relation extraction using SORS provides us information not present in the sentence but generated by the algorithm (detailed experiment is discussed in Chapter 7. In many examples, these facts are not incorrect like in('Bamberg' , 'Germany').While these unintended but correct interpretation candidates did not inhibit a correct manual interpretation and hold for automated interpretation on a larger corpus.

On the other hand, by looking at the parser outputs, one can see that it provides us with limited information. In particular, relations from complex language constructs are missing. In the case of the output of Fig. 11, the relations apply to different entities, which inhibits any chaining employing reasoning. All in all, the parser cannot provide a densely connected set of facts that would make spatial or ontological reasoning effective. We carry out spatial and ontological reasoning manually and automatically and compare residual errors after processing the exhaustive search with facts extracted from the parser, and we cannot rule out all ambiguous interpretations in 25 percent of the sentences. However, we are facing wrong outputs from the parser in 50 percent of the cases. The detailed assessment is presented in the Chapter 7 of evaluation.

Furthermore, the presented over-generalization reasoning approach aids in the resolution of referent or geo ambiguity, which occurs when multiple places are related to the same name. Various approaches primarily rely on external knowledge and heuristics based on additional information about locations around the world, such as population and land surface area. Meanwhile, SORS uses the information in the provided input to determine and extract all potential spatially relevant triplets.

4.3.1 *Dealing with implicit and indirect references*

Researchers have been working on extracting information using contextual information in order to deal with unstructured text. In [294], an unsupervised context-enhanced method is proposed to detect geo-relation key phrases from web texts for extracting triplets. External semantic knowledge is introduced to relieve the influence of the sparseness of the geo-relation description terms in web texts. Individually, the contexts of geo-entities are fused with category and word semantic knowledge by determining a semantic similarity measure. The paper argues that the proposed method can efficiently enhance the ability to discover key phrases representing geo-entity relations with the sparse distribution. It can also detect new key phrases that are beneficial for generating new triplets to construct a Geo Knowledge Graph from web texts. A **limitation** of that approach is that it is unable to extract implicit phrases. For example, consider "Zijin had 49.28 tons of the gold output, and the gold produced from mining reached 20.70 tons, respectively accounting for 20.53 percent of China's total gold production". The system is unable to find the relation between China and Zijin. **However, the method proposed in this chapter focuses on spatial entities and generating all plausible interpreations of a single entity which results in hypothesizing an implicit relation between China and Zijin,which is then confirmed by identifying these in the OSM database.**

In [154], a hypothesis is considered stating that frequent simultaneous mentions of POIs would indicate their geographic proximity, i.e., spatial context is derived from similarities of travel reports. For example, bloggers would tend to visit geographically close locations and mention them in their narratives accordingly. The work presents an empirical interpretation of spatially enriched POI graphs that back up the hypothesis. Secondly, it proposes a triplet pattern and rule-based spatial relation extraction technique, which exceeds the contemporary system's performance. A significant **limitation** of this approach is that it is unable to extract spatial triplets without landmark in the current segment, also termed as indirect reference. Therefore data has been manually annotated for identifying indirect references. By contrast, the approach proposed during this research applies reasoning to hypothesize connections between pieces of spatial information extracted from text, thus automatically establishing links to (landmark-like) named entities that can be identified in the OSM database.

The following example is presented in another research [188], which requires more information than the default approach and highlights the importance of context in determining the correct location. "*London is a Canadian city located in Southwestern Ontario along with the Quebec City – Windsor Corridor. The city has a population of 366,151, according to the 2011 Canadian census. London is at the confluence of the non-navigable Thames River, approximately halfway between Toronto, Ontario, and Detroit, Michigan. The City of London is a separated municipality, politically separate from Middlesex County, though it remains the county seat*". Using SORS here provides us with a set of facts that will allow us to identify London as two different cities in two different places. The facts like <London in Ontario>, <London along Quebec city> and <London at Thames River> allow us to generate multiple queries which result in two different places with the same name.

The first query will result in London in Canada. At the same time, the third one will represent London in the UK, where at is replaced by in (using the idea presented by [372]), less specific preposition at could be interpreted as one of the more location-specific, closely related spatial prepositions "in," "on," or "by" in a more automated way, if the granularity level and type of the reference feature are identified.) Moreover, the river Thames in London provides different coordinates for London than the previous query. So in conclusion, SORS helps in dealing with both explicit and implicit information as it focuses on spatial entities and all possible relations between them.

In summary, making spatio-ontological reasoning an explicit step in the interpretation enables consideration of contextual dependencies. Clearly, exhaustive search does not tackle the fundamental problem of language understanding, but it relies on the assumption that the largest set of statements that can be matched to a geographic database corresponds to the intended interpretation. While our approach is unable to deal with negation or complex language structures, it may indeed be sufficient for typical descriptive texts. In a comparison using sentences from English Wikipedia that describes geographic entities, we see that reasoning can prune off most invalid interpretations. In contrast, natural language parsing results in some wrong commitments one is unable to recognize in a later processing step.

Before embarking on a comprehensive study to analyze this method, a comprehensive lexicon and knowledge base are prepared, and reasoning methods to be automated. Information required to build a lexicon and knowledge base is readily available using data sources such as WordNet and OpenStreetMap, yet these need to be linked on a semantic level.

The subsequent chapters discuss implementing the automated reasoning method using these sources to arrive at a spatial interpretation of the constraints. To make the approach efficient, a query strategy is presented in Chapter 6 to avoid costly queries by serializing queries and by focusing on appropriate candidate locations.

5 GEO-REFERENCING USING SPATIO-ONTOLOGICAL REASONING TECHNIQUES

This chapter is based on content from the previously published paper by Yousaf and Wolter [407] in an interdisciplinary journal- Spatial cognition and computation in 2021, which includes a reasoning method to geo-reference both named and unnamed entities. The contributions of the corresponding first author includes conceptualization of dealing with unnamed spatial entities and developing a reasoning pipeline to interpret the spatial information through discussion with my supervisor, related work, paper writing and evaluation of the put-forth approach.

Natural language Place descriptions provide spatial information in terms of spatially grounded objects and spatial relations. Such information is extracted in the form of declarative statements using Spatio-ontological reasoning in the previous chapter, which provides a basis for formulating the queries. However, that captured information is not enough for matching the spatial entities to a spatial database. Place descriptions contain name references that belong to the type of places like the museum and the exact place names. Besides, there can be indirect references to the place names or abbreviations or local names present in the descriptions, requiring more information and new methods to geo-reference the places correctly that are not toponyms.

This chapter presents an approach for automated interpretation of natural language place descriptions using an automated reasoning method that can supplement the information obtained by employing natural language processing. The approach identifies places in three different categories, i.e., named, unnamed, and unidentified, and applies abductive and deductive reasoning techniques to geo-reference. The approach takes advantage of the type nouns present in the description. It uses them to resolve and identify contextual dependencies that results in better performance for all types of entities(named or other). The presented method is capable of geo-referencing all places in natural language place descriptions, whether or not gazetteer names refer to them. Compared to traditional NER methods, the methodology explicitly takes advantage of the type noun, resulting in higher geo-referencing precision and recall.

Experiments from the paper [407] have been re-organized and will be presented in detail in Chapter 7, respectively.

5.1 MOTIVATION AND GOALS

Natural language place descriptions are common in everyday communication and could thus serve as a rich source of volunteered geographic information (VGI), given automated methods for their interpretation become available. Several works have already focused on identifying and locating place references within text [371, 186, 159, 403]. References to locations are commonly referred to as Geo-references. Geo-referencing is a common task in GIS and involves associating the information with some physical space location [188].

Place descriptions have been regarded as a qualitative reference system that describes geographic locations that comprise references to places. Geo-references can be present in several forms, a named entity like the city"Rotterdam", an unnamed entity represented by a place category like "post office", or even a vernacular form like the "picnic spot", and qualitative spatial relations are describing their relative whereabouts [64].

According to [164], Geo-references have many significant and desirable characteristics. They should not be ambiguous and refer to a single location (case within a given frame of reference). Nonetheless, there is no guarantee that there is only one city with the name of London. Place descriptions may contain places that are not referred to as official place names, like abbreviations or common names used by people to refer to a particular place, such as, *"picnic spot near Erba campus" Or UK instead of United Kingdom*. Also, people often refer to instances of geographic entities by mentioning their type instead of the actual toponym (e.g., "the beach next to the University of California, Santa Barbara" when referring to Goleta Beach) [34]. Finally, geo-references may be associated with an implicit granularity - for example, *"Calgary, Mull- postal address of the village"*. The postal address of a sparsely populated village refers to a coarser granularity than point coordinates given with a precision of 1m [188].

The process of identifying references to location helps to understand its geographical context within a particular text description [188]. Using this, we can easily recognize important locations listed in a document or a description and delete references to locations that do not apply to the core theme.

Automatically performing this task is essential, but a complex challenge for GIR and far more demanding than adding a set of coordinates to a list of well-formed addresses [409], performed on structured GIS information. While references to natural language locations can come in several ways, ranging from (e.g., "I am here") to toponyms (e.g., Bamberg) to postal addresses (e.g., Josef-kindshoven str 5, 96052, Bamberg), much of GIR 's research has concentrated on toponyms and addresses [227]. Thus, the focus is to identify these geographical references, both named and others (types and vernacular forms) unambiguously and by assigning spatial coordinates. The chapter is inspired by the notion of recognizing and dealing with unnamed entities in addition to named ones. Furthermore, it determines the extent to which this concept aids in addressing existing difficulties and gaps (identified in section 2.3.1.1 and 2.5). The processing pipeline 7 includes a module for entity extraction and resolution, which aids in dealing with geo-referencing and entity recognition.

With our research, we wish to understand how the automated interpretation of place descriptions can be achieved. We investigate a somewhat radical attempt that only performs shallow parsing of the input text, leading to an over-generalization of interpretation possibilities but avoiding wrong commitments at the parsing stage. Instead, we rely on reasoning to single out the most plausible interpretation. To this end, input text is processed to form a set of declarative statements $\text{rel}(\text{place\_ref}_1, \text{place\_ref}_2)$ using a spatial (in, near, etc.) and ontological relations (is of type) [405].

The hypothesis underlying our research is that spatial and ontological reasoning enables us to overcome limitations in previous approaches to geo-referencing. Limitations particularly addressed by our approach are a limited degree of robustness when recognising place names by means of named entity recognition (NER) systems, insufficient exploitation of context in named entity disambiguation, also called toponym resolution [229], and the need to rely on parsing natural language text. By our development of a reasoning framework for text interpretation, we aim to integrate methods focusing on sub-tasks and combine their strengths. Our approach presented in the following yields two contributions:

- An automated reasoning approach for filtering/identifying most plausible interpretations for named and unnamed entities in natural language place descriptions;
- a method for pruning off implausible interpretations by means of context-sensitive reasoning, thereby improving the performance of text understanding components such as named entity recognition.

This chapter details reasoning steps of our approach and analyses their contribution to revealing the intended interpretation of place descriptions, using a corpus of Wikipedia articles. Place descriptions are interpreted by matching words to the OpenStreetMap[1] (OSM) database.

The remainder of this chapter is structured as follows: Section 5.2 presents an overview and components of our approach. Sections 5.3–5.5 detail the spatial and ontological reasoning steps. By deriving conclusions and giving an outlook, Section 5.6 concludes the chapter.

5.2 SORS: GEO-REFERENCING WITH SPATIO-ONTOLOGICAL REASONING

Spatio-ontological reasoning follows a declarative programming approach using sets of relational statements. Each set of relational statements defines a branch in the search space of potential interpretations. In contrast to place graphs, our declarative approach can represent disjunctive information and thereby handle uncertainty. Moreover, a logic programming approach overcomes limitations of a triplet-based representation that requires all facts to be related to explicitly named entities, compare [154, 294]. Declarative statements relate distinct entities derived from the text and they can also represent relations involving more than two elements, e.g., ternary direction relations.

Our method's basic idea is to refine a set of relational statements that potentially over-generalizes information expressed in a sentence but avoids wrong parser commitments.

The focus of this chapter is interpretation of information extracted from language, and thus, the aforementioned works constitute a possible starting point. However, to empower reasoning, we require an additional semantic specification of the facts to be interpreted. In this work, we investigate whether implicit information present in the description can be exploited. The idea is to use contextual information present in form of type nouns to semantically enrich the interpretation and improve the geo-referencing process by considering unnamed types.

Therefore, we opted for standard NLP tools based on lexicons to represent semantics. A common method for representing semantics would be using ontologies [38, 47]. But in our spatio-ontological reasoning model, we are not directly relying on a formal ontology like OSMonto [78]. It is designed for interfacing OSM, since the ontological category of a word may be ambiguous, or in case of spatial relations like *near* even fuzzy. Translating ontological categories in form of *is-a* relations to logic programs gives us more

1 https://osm.org

flexibility to handle these challenges.For example, a statement is-a(castle, historic site) $\vee$ is-a(castle, building) can capture the ambiguity that in OSM castles may be represented as buildings or historic sites. While there are clear instructions for OSM how the categories differ, it is not clear from a place description which category applies. Using the appropriate ontological category is however crucial for identifying an entity in the OSM database. Using logic programming methods, we can simply search for a possible interpretation. Therefore, we employ a logic programming model in which relational statements can be composed using disjunction and conjunction.

We apply abductive and deductive reasoning methods to fill in missing information and narrow down the set of interpretations. Finally, we select the interpretation alternative, which achieves most geo-references.

5.2.1 *Overview and processing pipeline*

The processing pipeline and its required background knowledge are portrayed in Figure 12. It comprises four major modules, out of which we summarize three in the following. A detailed running example will be discussed in the context of the individual reasoning steps, yet for preview refer to Figure 13. Figure 12 provides us with an insight about the methods and output generated by each reasoning module along with the input directories which are used in different modules.

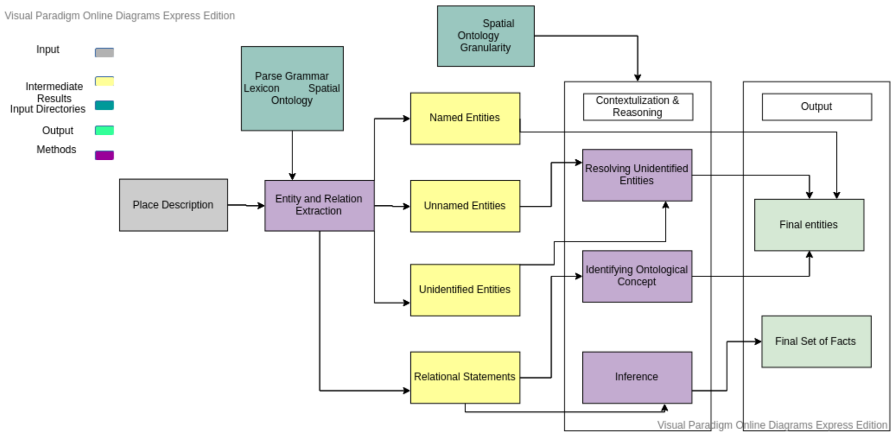

Figure 12: Methods for geo-referencing a NL place description. From the set of alternative text interpretations possible, the one which identifies most entities is selected.

The objective of the *entity classification* module is to extract information in the form of a set spatial entities, including named objects (e.g., Australia) and unnamed entities (e.g., café), as well as relations between them. Spatial and ontological relations are then exploited to resolve ambiguous named and unnamed entities. The method described below uses part-of-speech (POS) tagging, extended by a relation generation stage.

Reasoning is spread across three distinct modules repeatedly invoked until no new information is generated. The objective of the *contextualization* module is to connect statements to contextual information. It associates proximity relations like *near* with granularity information derived from the occurrence of geographic classes, for example, setting a lower threshold distance for inner-city level entities (streets, buildings, etc.) than for countries.

Contextualization performs abductive reasoning by associating (named) entities with unnamed entities if they represent the same ontological type. For example, *Danube* and *river* may be associated using an is-a relation – is a later querying step verifies the hypothesis that Danube is an instance of ontological type river. In the *inference* stage, derived information is propagated employing deduction.

Suppose, entity A is known to be located inside another entity B, which itself is contained in C. Constraint propagation techniques allow us to arrive at the facts that A is located inside C and that B and C must be region-like entities.

We employ reasoning steps in an exhaustive search manner, aiming to identify an interpretation of the input sentence, which leads to the most significant number of noun phrases matched to the OSM database.

Our approach of exhaustive search for a spatial meaning does not tackle the fundamental problem of language understanding – we assume a spatial semantics that may just not be present, for example, if considering a sentence like "Victoria and Charlotte are not close to one another". Such sentence describing the relationship of two women would be tried to interpret as to provide information about the city of Charlotte (South Carolina, USA) and the state of Victory (Australia). Nevertheless, our method can already be useful for automatic interpretation of text known to refer to spatial entities, or to gauge whether a spatial interpretation is possible.

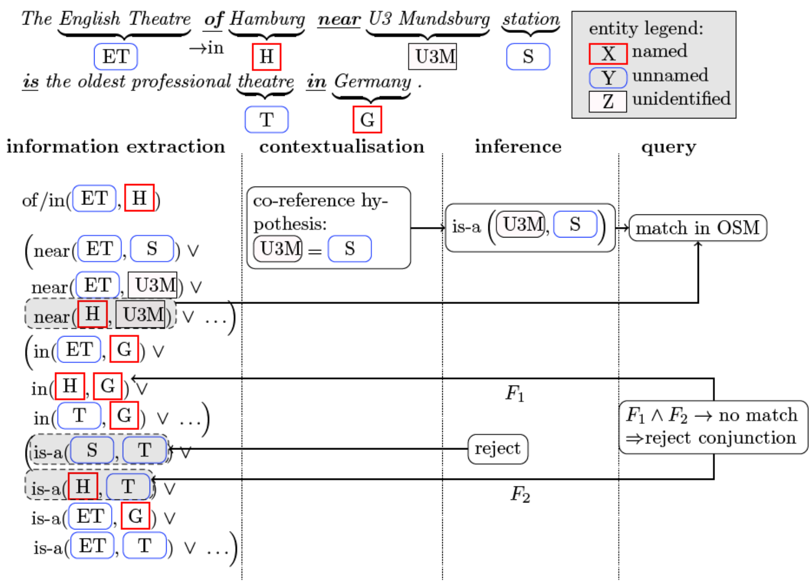

Figure 13: Illustration of the reasoning-based interpretation procedure which comprises stages information extraction, contextualization, inference, and querying, Extracted information is given as conjunction of alternative relational statements (grouped by parentheses), new statements get inferred and inconsistent statements are pruned off.

5.3 ENTITY CLASSIFICATION

The objective of this module is to translate a natural language place description into a set of spatial entities and constraints (see Figure 7). We apply the method introduced in [405], which operates as follows: Input text is processed to collect spatial information in a set of constraints Γ. Individual constraints represent assertions about spatial entities referred to in the text. If other nouns immediately follow (e.g., *art* followed by *building*), then these are treated as a single noun.[2] These entities are derived from nouns in the sentence. They constitute the set of variables V of our constraint-based approach. In the following, we make use of the following notation:

- Γ set of constraints
- V set of variables, instantiated with all nouns occurring in the input text
- Ordering relation
  1. $\prec \subseteq V \times V$ is a ordering relation of variables, $A \prec B$ means that word A occurs before B in the input text.
  2. $\succ \subseteq V \times V$ is a ordering relation of variables, $A \succ B$ means that word A occurs after B in the input text.
- N is the set of named entities, i.e., it holds that has-name$(E, x) \in \Gamma$ for all named entities $E \in N$
- *UN* is the set of unnamed entities. Unnamed entities may represent spatial concept classes. Thus, unnamed entity E may occur in form of a constraint is-a(x, E).
- *UI* is the set of unidentified entities. All elements of V are either element of N, of *Type*, or of *UI*.

Algorithm 6 shows the steps we take to extract entities and relations. The algorithm iteratively processes all words in a sentence (lines 4–18). First, nouns are grouped into named N and unnamed entities *UN*. For named entities, a 'has-name' constraint is generated (line 16). Likewise, unnamed entities are related to a normalised ontological type, e.g., words like river, creek, and waterway will all be related to the concept type 'river' (line 13). The set of ontological types and their hierarchy is modelled in our research prototype according to the OSM ontology, informally given by the set of key/value tags used, or formally by means of OSMonto [78]. We also store OSM key/value pairs per concept class for posing queries to the OSM database. Since we observed parsing to be a delicate task that is susceptible to misinterpretations even for simple English text, we refrain from tackling co-reference resolution at this stage but rely on subsequent reasoning steps to discover, for example, that *river* refers to the Danube river in the sentence *The bridge crosses the Danube river just north of a café at the river*. Line 16 generates relational statements between all possible nouns (note that W at this point is a relation such as 'in', 'near', etc.). In case of the above example sentence, there is no commitment whether *north* is a relation between entities *Danube river* and *café*, *bridge* and *river*, etc. This over-generalization helps us avoiding parse errors, which would disturb further processing.

Noun phrases neither recognised as named or unnamed entities are called unidentified entities *UI*, they can be computed by set-difference $V \setminus (N \cup UN)$. During a successful interpretation of a place description, this set would be reduced to nouns with non-geographic meanings. Figure 13, first column, shows an example output obtained from Algorithm 6. As can

2 Currently, we ignore attribute adjectives, determiners, etc. contained in the noun phrase.

Algorithm 1 Entity and relation extraction

function ProcessInput(S) ▷ input sentence
  apply part-of-speech tagging to S
  $V \leftarrow \emptyset, \Gamma \leftarrow \emptyset$
  for W is word in S **do**
    if W is noun **then**
      if W is followed by one or more nouns $W_1, W_2, \ldots W_n$ **then**
        $W \leftarrow (W, W_1, W_2, \ldots W_n)$ ▷ build compound noun
      end if
      add W to V ▷ entity we aim to geo-reference
      if W is listed in gazetteer as spatial category C **then**
        add {has-name(W, W), is-a(W, C)} to Γ, add W to N
      else if W is listed as spatial category C in lexicon **then**
        add {is-a(W, C)} to Γ, using a normalised type name; add W to UN
      end if
    else if W is a spatial relation or 'is a/is' phrase **then**
      add branch $\bigvee_{A,B \text{ are nouns in } S, A \prec W \prec B} W(A, B)$ to Γ ▷ disjunction
    end if
  end for
  return V, UI, UN, Γ
end function

be seen, the input sentence contains a compound noun phrase *U3 Mundsburg station*, which the part-of-speech tagger failed to identify as such.[3] Instead, *U3 Mundsburg* (U3M) and *station* (S) are interpreted as distinct noun phrases. In case of the relation 'near', relational statements near$(ET, U3M)$ and near$(H, U3M)$ are generated since at this stage it is unclear whether the English Theatre (ET) is related to 'U3 Mundsburg' (U3M), or Hamburg (H) is related to 'U3 Mundsburg'.

In our current implementation, we do not consider verbs except for *be* (is). However, verbs may express spatial information (e.g., a river *going round* some village or *meandering* in a valley). Whether or not spatial information from verbs is used in geo-referencing seems not to impact how subsequent interpretation components should be designed, which is the focus of our current work.

In our research prototype, we have manually implemented the OSM concept hierarchy and store OSM key/value pairs for composing queries. Mappings from nouns to ontological classes are stored in the lexicon. Noun phrases neither recognized as type specification nor as names are collected in a set of unidentified entities. During a successful interpretation of a place description, this set would be reduced to nouns with no spatial definition. Initially, unidentified entities may merge with either named entities (interpreting the noun as a name) or as entity types. In the latter case, the information is not further used in the query since no mapping to OSM types is known.

3 Even if the tagger would have identified the named entity, querying OSM would still be a challenge: *U3* designates a specific metro line, whereas *Mundsburg* is the name of the station. OSM queries only succeed if *Mundsburg* is used as name.

5.3.1 *Example*

Let us consider the following sentence as a running example:

> *St. Catherine* ***is*** *the largest protestant church, located* ***in*** *the city center* ***at*** *the entrance* ***to*** *the Zeil, the central pedestrian street for shopping.*

Table 8 shows the classification of nouns obtained by Algorithm 6, along with abbreviations we use in the following for brevity. As can be seen, the sentence is determined to contain no named entity[4], while four unidentified entities and three unnamed entities are discovered. In the table, we use the shorthand notation $r(\{A, B\}, C)$ to represent the disjunction $r(A, C) \vee r(B, C)$.

Table 8: Classification of spatial nouns determined by Algorithm 6, using Stanford Core NLP for part-of-speech tagging.

named entities N	unnamed entities UN	unidentified entities UI
(none)	church (C)	St. Catherine (SC)
	city center (CC)	entrance (E)
	pedestrian street (PS)	Zeil (Z)
		shopping (SP)
constraints extracted:		
is-a$(SC, \{C, CC, E, Z, PS, SP\})$, in$(\{SC, PS\}, \{CC, E, Z, PS, SP\})$,		
at$(\{SC, PS, CC\}, \{E, Z, PS, SP\})$, to$(\{SC, PS, CC\}, \{E, Z, PS, SP\})$		

5.4 CONTEXTUALIZATION

This module's objective is to generate hypotheses that ease the interpretation of the text, using a piece of information derived from context. Since the overall processing is performed as a search for the most plausible interpretations, invalid hypotheses generated at this stage are not fatal to the whole process. Nevertheless, a wrong hypothesis can lead to false positives in toponym resolution. Our approach consists of three independent steps, which are:

1. Resolving unidentified entities as named entities, thereby correcting misses during named entity recognition;
2. identifying ontological concepts for these entities;
3. and establishing reference to contextual factors.

5.4.1 *Step 1: Resolving unidentified entities*

Unidentified entities are transformed into named entities if, in the input text, the unidentified entity can be linked to an unnamed entity (type noun), thereby identifying their ontological type. In this step, we identify implicit and explicit relationships between spatially declared entities. The set of constraints Γ and variables V determined by entity and relation extraction serve as input for this method. Contextualization can be considered as a

4 That is clearly wrong. However, if we use available named entity recognition systems Stanford NLP or spaCy, Catherine is identified as a name while Zeil is identified as date.

form of abductive reasoning, which is *inference to the best explanation* by creating and testing all possible hypotheses available. In this step, we look for entities and relationships between them to the point where there are no new hypotheses created by Algorithm 2: it augments the set of constraints Γ and also outputs a set of *inferred entities IE*, which are considered in the evaluation in order to determine the contribution of contextualization to geo-referencing.

Algorithm 2 is responsible for identifying an unidentified entity referring to an unnamed entity, thereby changing the unidentified entity to a named entity. The motivation here is to discover relationships in phrases such as *river Elbe* or *Danube river*, etc. even if the named entities are not recognized (line 3 thus iterates over named as well ass unidentified entities). The algorithm makes use of a function $\textsc{pos}(W)$ which determines the position of a word in a sentence. For neighbouring words W, V it thus holds that $|\textsc{pos}(W) - \textsc{pos}(V)| = 1$. If there is no directly neighbouring unnamed entity present in the next or previous position of unidentified entity and a word remains unidentified thus far (line 10), any other unnamed entity in the sentence not otherwise related is proposed as concept for that word.

Algorithm 2 Generating concept class constraints

```
 1: function INFER(Γ, N, UN, UI)
 2:     IE ← ∅                                             ▷ inferred entities
 3:     for W ∈ UI ∪ N do
 4:         for V ∈ UN do
 5:             if |POS(W) − POS(V)| = 1| then
 6:                 retrieve concept is-a(V, C) from Γ   ▷ generated in first step
 7:                 add branch (is-a(W, C) ∧ has-name(W, W)) to Γ, W to IE
 8:             end if
 9:         end for
10:         if W ∉ IE then
11:             for V ∈ UN and V not related to other words do
12:                 retrieve concept is-a(V, C) from Γ
13:                 add branch (is-a(W, C) ∧ has-name(W, W)) to Γ, W to IE
14:             end for
15:         end if
16:     end for
17:     return Γ, IE
18: end function
```

Observe that the proposed algorithm also tackles the co-reference problem for special forms, i.e., it can identify that two different phrases refer to the same entity. In particular, Algorithm 2 allows phrases like *Hamburg isn't the only major town at the banks of Elbe, also Dresden is located at the river.* to be interpreted without requiring the sentence to be understood. In the given phrase, river would hypothetically be related to *Dresden*, *Elbe*, and *Hamburg*. When the second hypothesis is confirmed by querying osm and finding a match for an entity named *Elbe* of type *river*, all places mentioned can be identified. In our current system, references are only established within the chunk of input text processed. However, the approach could be extended to maintain a history of entities, as proposed in [193].

Table 9: Resolving unidentified entities as named entities

unidentified entity	hypothetical concepts
St. Catherine (*SC*)	(already associated during information extraction)
entrance (E)	city center, pedestrian street, church
Zeil (Z)	pedestrian street, city center, church
shopping (*SP*)	pedestrian street, city center, church

5.4.1.1 *Example: Resolving unidentified entities*

All four unidentified entities present in Table 8 from our running example (page 78) are resolved by application of the algorithm. The inferred entities *IE* with their associated concepts as determined by Algorithm 2 are shown in Table 9 in the column hypothetical concepts. The table lists all unidentified entities and shows for those that have not already been related to a concept by Algorithm 6 the hypothetical concepts generated by Algorithm 2. The different hypotheses are sorted according to how far the related entities are mentioned in the sentence. As can be seen, the correct interpretation for *Zeil* of being a *pedestrian street* is listed first. This motivates the use that ordering as heuristic in the search for a plausible interpretation. For *entrance* and *shopping* no sensible interpretation is possible with our approach.

5.4.1.2 *How to find the most plausible interpretation?*

For place description consisting of many nouns and relations the search space can become so large that a search strategy is required to make search feasible. In the following we discuss a heuristic that can be employed to gear search towards finding the most plausible interpretation in the sense that most nouns get geo-referenced.

We employ a simple heuristic based on estimating the likelihood that a constraint generated by one of the algorithms shown is correct. Such heuristic allows us to order all search branches and process them in order, starting with the most likely branch. Our heuristic is based on the observation that in place descriptions related words often appear in closer proximity than unrelated words. To this end, we determine the distance of related words within the sentence per constraint and sum up these distances in order to compute the average distance of words related.

Additionally, we consider a simple pruning method related to handling unidentified entities. In Algorithm 2, all combinations of relating an unidentified entity to an unnamed entity are considered (starting at line 10). This part of the algorithm is effective in recovering from failures of part-of-speech tagging to recognise a named entity, but it also generates several irrelevant search branches if the place description contains nouns that do not have a spatial meaning. In case of n unidentified entities with no spatial meaning and c concepts given by unnamed entities, we are facing $O(\frac{c!}{(n-c)!})$ irrelevant search branches. We therefore have chosen in our implementation to only explore the most likely search branches, pruning off all candidates but the one(s) with highest rank according to the distance heuristic.

Table 10: Top 5 search branches for the running example, sorted by heuristic (lower is better). Interpretations corresponding to ground truth are underlined.

rank	heuristic score	set of constraints with respect to inferred entities
1	10	$\underline{\text{is-a}((SC), C)} \wedge \text{is-a}(SP, PS) \wedge \text{is-a}(E, CC)$
2	11	$\underline{\text{is-a}(Z, PS)} \wedge \underline{\text{is-a}((SC), C)} \wedge \text{is-a}(E, CC)$
3	15	$\text{is-a}(Z, C) \wedge \text{is-a}(SP, PS) \wedge \text{is-a}(E, CC)$
4	15	$\text{is-a}(Z, C) \wedge \text{is-a}(SP, PS) \wedge \text{is-a}(E, CC)$
5	18	$\text{is-a}((SC), CC) \wedge \text{is-a}(SP, PS) \wedge \text{is-a}(E, C)$

5.4.1.3 *Example: heuristic ranking of hypotheses*

For our running example we obtain the set of search branches shown in Table 10 by application of the inference techniques explained above. The table lists all constraints that relate unidentified to unnamed entities, i.e., constraints relating to relations in the place description are omitted for readability. As can be seen in the table, the branch ranked second leads to the best interpretation possible in our approach which cannot provide georeferencing for concepts *shopping* and *entrance* [to Zeil].

Table 11: Comparison of how the extracted spatial constraints are resolved using reasoning for sentence "St. Catherine is the largest Protestant church, located in city center at the entrance to the Zeil, the central pedestrian street for shopping."

extracted spatial entities			final entity classification		
named	unnamed	unidentified	named (inferred)	unnamed	unidentified
(none)	city center	St. Catherine	St. Catherine Church	city center	shopping
	church	Entrance			entrance
	street	Zeil	Zeil street		
		shopping			

5.4.1.4 *Establishing reference to contextual factors*

In the contextualization stage, references to contextual factors are established. So far, we only consider geographical scale as influencing factor, which determines threshold distances for interpretation of spatial relations near/at. In our implementation the thresholds are hard-coded for simplicity. We determine geographical scale from the type of entities identified, using a simple classification of types occurring (country level vs. town level). It is subject to future research to determine means for evaluating spatial relations within an uncertain spatial context.

This module helps to resolve unrecognized objects and generate the output of object names and suitable spatial types to enrich the query process and provide initial ambiguity resolution for named entities. Furthermore, explicitly describing the entity type permits answering questions like where or what, which can aid in dealing with inquiries answered by geographic information systems.

The relational statements generated using approach presented in chapter 4 will be used in the next step 5.5 to generate a set of facts that can be used as query for locating the identified entities on the map.

5.5 INFERENCE

Once is-a and spatial relations have been instantiated, a deductive constraint propagation is performed to make implicit relations explicit. Whenever the deduction reveals an inconsistency – say a single entity is said to be a theatre and station at the same time – the current search branch gets rejected.

This inference can be seen in the example shown in Figure 13, in the inference column. It is also duty of the inference module to verify that spatial relations are satisfied when entities are matched to the geographic entities in the OSM database. If a relation cannot be verified by geo-references made, the respective search path is rejected.

In order to propagate facts, we employ path-consistency [94] for is-a relations, in case of spatial relations algebraic closure [101, Algorithm 1] which is commonly used in qualitative spatial reasoning [301]. We note that is-a relations are transitive and the path-consistency algorithm will enforce the desired transitive closure. Currently, our system only handles is-a relations and spatial containment as well as co-location (at, by).

The process is divided into two main steps:

- **First Step**, conjunctions are formed exploiting the transitive closure property. Details are provided in Figure 14. The first column explains how connections are formed (different colours represent the process), and the second column gives the output.

- **Second Step**, duplicate sets of hypotheses are pruned off, reducing the set of possible interpretations.

RELATIONAL STATEMENTS (p → q) AND (q → r)	APPLYING DEDUCTIVE REASONING (p → r)	PRUNING PHASE	FINAL OUTPUT
Is a (SC, C)			
in (SC, CC) V (C, CC)			
at (SC, E) V at (SC, Z) V at (SC, S) V at (SC, SP) V	at (SC, E) **V** at (SC, Z) **V** at (SC, S) **V** at (SC, SP) **V**	~~at (SC, E) V~~ ~~at (SC, Z) V~~ at (SC, S) **V** at (SC, SP) **V**	at (SC, S) **V** at (SC, SP) **V** at /in(SC, CC) **V** at (SC, Z) **V** at (SC, E) **V**
at (C, CC) V at (C, E) V at (C, Z) V at (C, S) V	at (SC, CC) V at (SC, E) V at (SC, Z) V at (SC, S) V	at (SC, CC) V ~~at (SC, E) V~~ at (SC, Z) V ~~at (SC, S) V~~	
at (CC, E) V at (CC, S) V at (CC, Z) V	at/in (SC, E) **V** at/in (SC, S) **V** at/in (SC, Z)	at/in (SC, E) **V** ~~at/in (SC, S) V~~ ~~at/in (SC, Z)~~	

Figure 14: Inference Steps Showing the Process for Generating Final Queries.

Figure 14 provides us with a step by step overview for generating the final set of declarative facts for example. The first columns present the method of how and where the transitive property is applied. The second column contains a combination of the deduced set of facts (in green color) and the remaining facts not used to form conjunctions. The third column, pruning, is carried out on the incoherent facts along with the duplicate ones. In this particular example, St.Catherine(SC) has already been identified as a church in the previous section. However, after applying deduction, its inferred as *at(SC, SP)* is incorrect; thus, the search path is removed. In the case of entrance, as the concept is unknown, the fact *in/at (SC, E)* is considered for further processing.

This allows us to exploit contextual knowledge to transform entities into named ones resulting in a set of data that can be converted to OSM query equivalents. Additionally, the methods from Section 5.5 provides us with the finalized sets of facts that will serve as queries with spatial relations.

5.6 DISCUSSION AND SUMMARY OF THE CHAPTER

Geo-referencing places from unconstrained natural language is a tough task comprising several unsolved sub-tasks. The primary focus of our research is to identify the OSM objects that correspond to spatial entities present in the place description. The overall aim of our approach is to address the handling of nouns that refer to spatial entities in the form of named, unnamed, and unidentified entities. The implications and limitations of the suggested approach are taken into consideration in this section.

As discussed in previous sections, most of the existing systems are unable to identify indirect references in the text like *park near Erba campus*. Our approach allows us to deal with named and unnamed entities,it is able to geo-reference unnamed entities and to exploit them for improving geo-referencing of named entities by connecting names to ontological types.Even in cases where we are unable to provide specific objects on the map, multiple entities are shown, all of which present reasonable interpretations. For example, *The Historical Museum of Bamberg* or located in the Alte Hofhaltung, both phrases are referring to a spatial location. The *Alte Hofhaltung* is an unknown object because we are unable to find any equivalent in our lexicon that can add to its type and geometry. On the other hand, the Historical Museum is a generic term that can refer to any number of museums. Even with multiple museums, as a result, one of them is the correct one,which is identified by the presence of Alte Hofhaltung right next to the museum indicating this is the expected outcome of the sentence.
The module of entity and relation extraction can extract the implicit key phrases present within the sentence, independent of the category of the spatial element. The module of contextualization help overcome the problem presented in [303]. It enables the handling of ambiguity in toponym resolution for named entity recognition to be resolved by assessing plausibility of an interpretation in context of the complete sentence (cp.[303]).The example is about a "furniture store located in Puebla state in Mexico City, and the other is located in La Paz in the neighborhood of Puebla." NER is unable to differentiate between the two Puebla's. Using the contextualization module, one Puebla is inferred as a state, and the other is deducted as a neighborhood. Besides, La Paz is labeled as a city. Moreover, with entity and relation extraction module, relational statements like Puebla in Mexico and the state of Puebla and the city of Lapaz helps identifying the correct place. Furthermore, SORS helps in dealing with the historical texts where the names are ambiguous or not known, and by applying the type information, we can reference them to a real place on the map.

Overall, this chapter presents an approach for automated interpretation of natural language place descriptions using an automated reasoning method that can supplement the information obtained utilizing natural language processing. This chapter's focus is to analyze the design and contribution of spatial-ontological reasoning for geo-referencing places from natural language input. We apply information extraction to obtain relational expressions from the input text, aiming for an over-generalization that avoids wrong parser commitments. A reasoning stage follows, which starts with an abductive contextualization step that generates hypothetical is-a relations and unification, using information available in the input context. A deductive inference stage then propagates relational expressions and rejects inconsistent interpretations.

Our system is implemented in a constraint-based programming approach that does not require a specific order of steps to be followed as graph-based approaches such as [371]. Chapter 7 will present a evaluation of our implemented system that reveals the effectiveness of reasoning, in particular the ability to resolve and interpret unnamed entities. Our approach does not require a correct parsing of the input sentence but can derive information from just light parsing. Despite the variety of open problems faced in understanding spatial language, the proposed method is already able to interpret composite phrases like *post office near the train station in Regensburg* correctly by the geo-referencing post office, train station, and the city of Bamberg; we regard this to constitute a potentially useful advancement.

In SORS, it is assumed that maximising the number of entities that can be geo-referenced leads to the most plausible interpretation. While this is certainly true for purely spatial descriptions, yet the assumption poses a clear limitation when handling text which is not of a pure spatial nature.

As it is an initial step in geo-parsing through reasoning,effectiveness of the proposed method depends on the availability of an mapping of unnamed entities to geographical types that can be used in the query of the OSM database. In other words, the quality of the underlying lexicon determines performance of the method. For example, for the sentence *I am at a local residence on the western side of KFC near a shoe shop in Bamberg*, the system is unable to handle *KFC* with the correct dependency (to recognise the named entity and find it on the map), but only tries to relate it to a *shoe shop* and *residence* – both do not lead to the correct result. Moreover, the sentences where there is no clear spatial concept present like *this place is best of coffee* require more information and external resources to interpret, in particular means to associate *coffee* with restaurant, café, etc. An approach to discover means to unfold more contextual information in order to enable a larger type of texts to be interpreted is presented in chapter 6.
Moreover, we like to mention that there is a significant gap between human concepts of place and geographic databases (see, for example, [394]) that needs to be bridged to arrive at a useful system. For example, there is no easy way to map *picnic place* to a category represented in OSM. Last but not least, the rich body of spatial relations has not been exploited, in our system only containment and proximity relations are implemented in a straightforward manner.

6 QUERY FORMULATION AND STRATEGY

This chapter is based on content from a published paper [404] in GIScience 2018 and [406] in GIR 2019. The chapter extends the query strategy discussions and how our system has been implemented using constrained specific queries. The query framework, which is briefly introduced in the paper, is formalized in this chapter. The chapter additionally proposes methods for improving querying by using contextual factors like type or generating type nouns. The corresponding (first) author's contributions include problem conceptualization, related work, a processing pipeline, experimentation, and paper writing. Furthermore, the last section of the chapter is based on the work which is still in progress and will soon be submitted fro review.
A Spatio-ontological reasoning system (SORS) retrieves a range of data about locations from place descriptions, such as the name of the relative location and the type of places indicated by place references and spatial relationships. A SORS should preferably support querying in order to utilize its captured knowledge in application scenarios. This chapter finds geographic entities in the OpenStreetMap database corresponding to every noun phrase in place descriptions that refer to real-world entities.
This chapter first provides a pipeline to compose the queries using the input presented. Then the chapter analyzes types of queries that can be answered by OSM and provides a query strategy that will help schedule the order of queries. Additionally, this chapter addresses the issue that arises while composing queries leading to incorrect or no results and how semantic similarity can be used to improve the effectiveness of the query process. Lastly, the chapter provides insight into dealing with non-spatial information and how semantic similarity can help us find similar terms in the form of unnamed nouns.

6.1 MOTIVATION AND GOALS

This research aims to develop an automated system capable of interpreting spatial language for dealing with place descriptions [406]. Effective automated interpretation is a combination of recognizing the geographic references and resolving them unambiguously to locations on the earth's surface [188] i.e., to relate all nouns in a sentence that represent geographic entities to corresponding entities in a geographic database([1] [406],[99]. In order to map the entities the GIS should be able to ask questions like *What is where?*and should allow the users to query and extract information [99].
In many situations, natural language is preferred from the description generation perspective, rather than other representations of spatial knowledge (such as maps) and is often non-intuitive with irrelevant details and mental interpretation costs [63]. As the volume of unstructured text documents increases [260], the growing need for place related information system also increases, which can provide locally valid and richer data on place, as users might benefit greatly by having such data in many situations [91].
The first component in such a system is a user interface, which mediates between the user and the system and supports user interaction, for exam-

1 which in our case is OSM

ple, helping users formulate queries, evaluate results and reformulate their search [188],[348]. It has been well-recognized in GIS literature ([377, 379, 122]) that natural language interfaces (NLIs) would effectively increase the use of more complex GIS characteristics, however, despite NLI's potential value for GIS, there has been somewhat limited work on this subject to date [379]. The design of NLIs for databases is generally considered a complicated problem because the interaction between human beings is often vague, ambiguous, or extremely contextual [377, 176].

After extracting the information using spatio-ontological reasoning, this thesis's last task focuses on composing the queries and scheduling them for mapping the entities.The chapter tends to focus on the final two components of the processing pipeline described in section 7 of chapter 2.The primary objective is to create queries and algorithms that might strengthen the geo-referencing processes. The queries will be in the form of constraints and based on the type and the name extracted in previous chapters.As a query engine we use OSCAR[2] [31],[30]. OSCAR is a new geospatial search engine based on freely available OSM-data. The framework can process general query expressions involving unions, intersections, set differences, substring searches, and geographic region constraints on the complete OSM planet data set [30]. The motivation here revolves around creating a natural language interface for OSCAR, which translates text descriptions into OSM queries by considering simple combinations that can help us with different scenarios present in a place description. Therefore, this chapter will first explain how queries are formulated from natural language to OSM querying format(from extracted set of entities). As it is challenging for non-expert users to formulate such queries, thus the system will automatically convert and be process by them.

The second task of this chapter is to provide with a query strategy to schedule the order of queries (e.g., A is north of B, which is north of C) and to decide whether a query should be submitted or the number of results expected is too large to be manageable. This chapter's third subtask is to identify appropriate key-value pairs for any type of spatial category in our input text. As language is rich and offers many words to communicate nuances of a single entity type and conceptualization of natural language entities in text and OSM often differ. Even the classification of a single entity type is subject to variations throughout the OSM database; thus, this task aims to identify similar OSM tags using semantic similarity for any word encountered [406].Last part of this section is mainly concerned with interpretation of *paraphrased* places, i.e., entities for which no name is given and which may only by described.Challenges arise from the open-endedness of language, its ambiguity, and context-sensitivity as well as from mismatches between human conceptualization of place and database ontologies. Our objective is to determine suitable entity types that allow querying the Open Street Map database for the respective place.

The remainder of this chapter is structured as follows. Section 6.2 will provide us an overview of the existing natural language interfaces and how the process of querying is carried out. The processing pipeline and steps for converting natural language place descriptions into OSM query format are presented in section 6.3. Section 6.3.4 addresses the second task by providing details about querying OSCAR. Section 6.3.5 provides an initial discussion of the limitations and gaps in the process of query conversion.

2 https://github.com/dbahrdt/oscar

Section 6.4 provides a solution to the existing limitation in the form of clustering algorithms using semantic similarity, which helps in improving the query process.Section 6.5, presents a clustering algorithm using semantic similarity to deal with non-spatial information. Section6.6 summarizes and concludes the chapter.

6.2 EXISTING QUERY METHODS

The spatial querying approaches used in Geographic Information Systems adopt a set of spatial query operators (for example, within, overlaps), which may be combined using Boolean operators (and, or, not) to build more complex queries. The use of these operators requires some training and expertise to understand their functions and formal meaning. Some GIS software packages provide spatial operators with natural language phrases (such as ArcGIS), but understanding the spatial operator's meaning is still needed. While these approaches are suitable for those who regularly use GIS and are accustomed to handling geographical information, it is difficult for non-experts to use them [348]. The objective of a natural language interface for GIS has been recognized as both significant and challenging for some time. Its advantages include ease of use for non-experts, the ability to more adequately describe requirements and speed of issue of command [122].

Several methods have been developed to support natural language spatial querying. The first is Natural Language Processing (NLP) methods to parse and interpret Geospatial queries. Examples include applying the theory of the possibility of parsing trees from geospatial queries to differentiate grammatically ambiguous sentences and 'de-fuzzify' uncertain spatial terms [378] and geospatial query parsing and part of speech tagging as a prerequisite of interpretation and inference [29].

Another approach is to apply constraints on the form of queries rather than trying to interpret a wide range of natural language sentences. It will become more user-friendly than the formal query languages and more comfortable for the machine to interpret and process. In [237], the system queries are based on templates for users' various dimensions (how, where, when, and what). [62] seems to provide us with a chat interface that supports specific GIS requests. In [157], they use Rabbit, a relatively limited natural language tool, to help construct OWL3 (Web Ontology Language) ontologies in a geospatial context. In our approach, we consider a combination of parsing and applying constraint on the form of queries based on the type of entity present in the description, and the details are provided in section 6.3.4.

Many researchers have used ontologies to aid with the interpretation of natural language geospatial user queries, either parsed or from a restricted set of expressions [416, 29, 129, 311]. Ontologies are commonly used to interpret domain concepts (so that the query can be expanded to include semantically equivalent terms) [311], but have been employed for geographical and query aspects, in some cases [129]. However, we have chosen not to employ formal ontology techniques and tools based on description logics for two reasons. First, the truth semantics of classical ontology languages is binary, i.e., entities belong to a particular class or do not. In the case of spatial entities and concepts, such crisp classification may be hard to achieve, and concepts may vary across individuals. Instead, relations like 'near' may be more adequately represented using a semantic capturing vagueness, e.g., using Fuzzy [406] or probabilistic models [99]. Second, existing ontology

languages do not support the spatial domain and manifold spatial relations to the extent required to empower spatial reasoning for computing likely interpretations of a locative phrase.
Likewise, we employ an ontology-like knowledge representation to capture constraints on potential interpretations of locative phrases. In order to compile all relevant background information in a single knowledge representation, we have developed a dedicated representation. Therefore, the key to this approach is building a knowledge base that captures possible semantics of natural language phrases (a lexicon).
Researchers consider a small set of spatial relations for querying, as discussed in [30],[129], which includes near, inside, and within a specified distance and along with specific cardinal relationships [129]. It has also been demonstrated how topological and metric concepts can evaluate the semantic similarities between various natural language queries using a set of 15 spatial relations. This previous work on spatial relationship querying focuses on interpreting specific spatial relationship terms, usually from English (although some are discussed in a cross-linguistic context.According to the hypothesis presented in [372] less specific spatial relations like *at* can be interpreted as *in*, *on* or *by* based on the granularity level and type of the reference feature are identified. Thus, based on this hypothesis, only containment and proximity relations are implemented straightforwardly in our system.
The available commercial search engines like Google and Bing Maps have been used by many applications for querying and geo-referencing, while current free/OSM-based search engines are still way behind. Our research uses OSCAR, which uses freely available OSM-data and, due to its substring searches, performs better in cases of the unknown spelling or incomplete names (as discussed in [30]) compared with Google and other search engines [30].The next sections will provide us with the processing pipeline and details about how our prototype works.

6.3 COMPOSING AND SCHEDULING OSM QUERIES

This section is divided into two parts. In the first part, the processing pipeline for composing and querying OSCAR is explained along with the basic concepts of OSM. The second part lists and introduces types of queries that can be scheduled to OSM .

6.3.1 *OSM basics and processing pipeline*

The following fundamental concepts of the OSM data model are essential to understand to deal with the difficulties that occur while processing the data to search for OSM queries effectively. There are three types of data in the OSM: 'nodes' (with latitude and longitude), polygonal paths or 'ways' (consisting of one or more nodes), and 'relations,' which are compositions thereof (i.e., referencing sets of nodes, ways, or other relations). All three forms may be expanded to include *tags* called key-value pairs, which allow to name and identify the data [30].

In general, entities are tagged with the type and name. Objects modeled through the OSM constructs node, way, and relation are administrative entities [30]. In queries, the desired target object is described using key-value information. Given a natural language description like "the campsite south of Lyon, near the river," an entity of type river could be found by specifying key "waterway" with value "river," *waterway=river* in short. The campsite can be found using *tourism=campsite* as key-value pair, whereas "Lyon" could be found using key-value pairs *place=city*, *name=Lyon*. Keys "tourism" and "waterway" are generic keys encompassing multiple values. Using matching key-value-pairs in the query is required to retrieve the desired objects. A query format consists of an expression E, which might contain intersection, union, set differences as operations and substring and prefix text searches as well as geographic region constraints as operands, hence covering all possible set operations. For example, the query (campsite Lyon @place:city) should return all items which contain the substrings' lyon' and 'campsite.'

Thus this section provides us with a processing pipeline that helps convert the entities extracted by spatio-ontological reasoning in the previous chapter 5 to the desired OSM query format. In our work we assume that an input text chunk has been processed by a part-of-speech tagging (POS tagging) and parsing to reveal potential structures of the sentence.

In the example shown in Figure 15, there is some ambiguity in the association of the phrase *on the river of Regnitz* to either *Upper Franconia* or *Bamberg*. All possible interpretations are processed simultaneously by combining them in a single disjunctive logic program. For every relation, a term is constructed combining any word (noun or named entity) appearing before the relation with any word occurring after the relation. The designator for each relation is retrieved from the lexicon. %, e.g., in(park, town) and of(town,Bamberg). For every noun an ontological "is-a" relation is generated in reference to any other noun or named entity, e.g., is-a(Bamberg, town) The basic idea of our approach is thus to generate all candidate interpretations of references and then apply reasoning to single out the most likely interpretation. We are using spatio-ontological reasoning for this phase where we are using an ontology-like representation to augment the semantic representation of words (the lexicon).The Figure 15 shows the first step of our information extraction phase. A lexicon is applied to associate spatial entity designators (e.g., *city*) with concept classes represented in the ontology-like knowledge representation. From their we extract possible spatial interpretations of all concepts in a disjunctive form. First one provides us information about the named entity by telling its type (which can be a water body or a land body)and geometry (which can be point line or region). While the second is based on spatial relations and the information we can get from them like the geometry of entity, axis which they represent(horizontal or vertical) and some information about the motion or position of the object(context).
We note that so far we have only realized a simple named entity recognizer that introduces the option to interpret every capitalized word or phrase as named entity – during final interpretation, a geographic database is then consulted for potential matches.

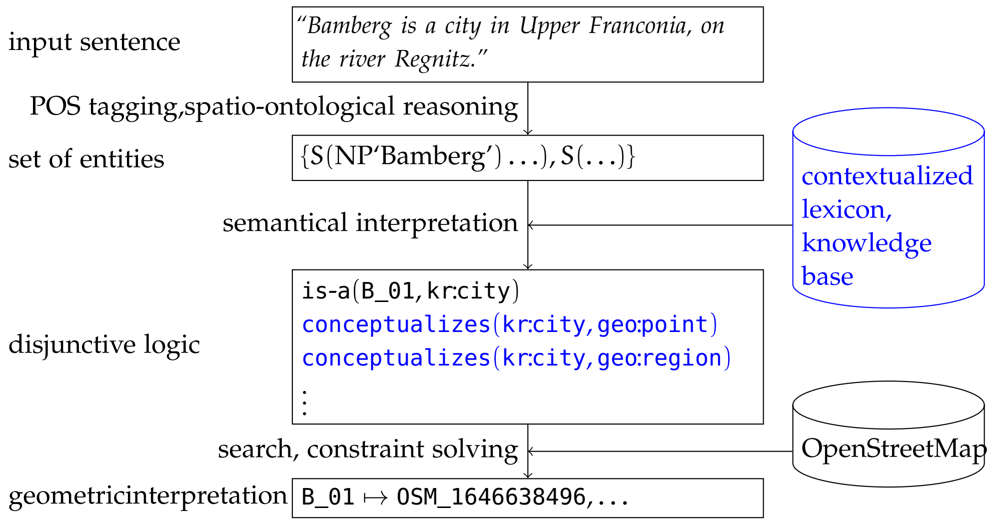

Figure 15: Overall processing pipeline for interpreting a place description

6.3.2 *Contextualized Lexicon and Knowledge Base*

Our contextualized lexicon aims to capture potential, alternative meanings of a natural language expression and its implication on interpreting other parts. In the figure, two potential conceptualizations of a city are indicated: entities of type city may be a host to specify containment (cp. [43]), yet they may also serve as point-like landmarks. With respect to interpretation of spatial relations, ambiguity also arises from resolving frames of reference . For example, the interpretation of a projective relations such as *the house left of the town hall* needs to related to an frame of reference that provides a reference direction which may not be given explicitly. It may be required to identify an appropriate frame of reference established in other parts of the input text (e.g., *suppose you are standing on the tower looking at city. . .*), established by objects in the environment (e.g., an entity that exhibits an intrinsic frame of reference), or derivable from (imaginary) movement in the environment (cp. [**tschander-SC:03**]). It is not the objective of our lexicon to determine an interpretation, but to list alternatives. Therefore, information provided by the lexicon is not a set of facts but logical templates.

6.3.3 *Geometric Interpretation*

In addition to above presented output our ontology like knowledge representations , provides information related to entities like the possible geometries for the spatial entities depicted by spatial ontology are *Bamberg :type : city,geo : Region* and *RiverRegnitz :type : river,geo : lines* and *Germany :type : country,geo : Region* and *RiverMain :type : river,geo : lines*.

Similarly,the possible geometries for the spatial entities depicted by spatial relations ontology are: *in : UpperFranconia :geo : Region* and *in : Germany :geo : Region* and *on : RiverRegnitz :geo : lines* and *on : RiverMain :geo : lines* and *on : UpperFranconia :geo : lines* and *on : Germany :geo : lines*,which will help us in our next step to query OSM and get desired results for our input which can be questions or simple declarative statements. The minimum information OSM requires for the query includes name, type and geometry of the entity and as we can see from above that we have all these three things with us and now all we need is to translate our natural language place description into OSM query syntax. The mapping of NLP to potential concepts of OSM results in the following syntax #Bamberg @ *place* : *city*, where geometries are shown as # for region and points and ! for lines.

Using contextualized lexicon and geometric interpretations the final entities inferred in our last Chapter 5 are converted into the desired OSM query formats. Table 12 provides with an detailed overview of how the entities are converted into key-value pairs based on their type for the running examples.

Table 12: 'Semantic interpretation to OSM query format for Running Example.'

final entity classification	OSM query format
named (inferred entities)	OSM Equivalent
St.Catherine	#St.Catherine@building : church
Zeil street	!Zeil@highway : street

6.3.4 *Query Strategy*

The second stage in our approach is concerned with scheduling queries to the OSM database. Aside from the technical side of composing queries from relational statements, it is the objective of the query strategy to schedule the order of queries (e.g., A is north of B, which is north of C) and to decide whether a query should be submitted or the number of results expected is too large to be manageable. Once the entities are converted into the required OSM format, the next step is to use the relational statements generated in the inference module in Section 5.5 to formulate the queries. The relational statements are updated according to the type of entity inferred by spatio-ontological reasoning in the previous Chapter 5. For example 2, the type business district is replaced by Victoria in *in(William street, business district)*, thus changing it to *in(William street, Victoria district)*. We neither query for entities if only their name, but not their ontological category or containing region is known, nor if only the ontological category is known. Since OSCAR enables us to query for key-value pairs like `amenity=theatre` restricted to local surroundings, we trigger a query whenever either the name or type of an entity and its spatial containment are known, i.e., given as relational expressions (e.g., *William street in/of Melbourne* or *st.Catherine at entrance*). After applying SORS, the entities are categorized either as named or as unnamed entities. The named entities are further categorized as country, city, or river, while the name of a shoe shop or a museum falls in the category of other named entities. The second category is for the type entities known as unnamed ones.

In the case of entities known to be countries, no further context is necessary. We refer to these kinds of queries as *constrained queries*. For formulating constrained queries, we have designed a further categorization based on the type of entity in the form of conjunctions.

Figure 16 below provides us with a comprehensive list of the various categories queries are labeled.

The first step in scheduling a query is identifying the relational statement category by determining its type. For example, if the relational statements fall under the category of unnamed * city category, as one entity is named and one unnamed *William street in/of Melbourne*, the query structure is pre-

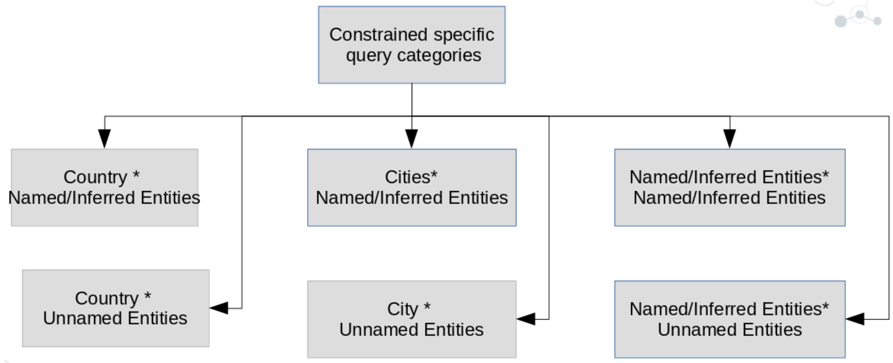

Figure 16: Constrained specific query categories for scheduling queries

sented in figure **??** where the city name follows the key-value pair for the entity we are supposed to find. Once all the osm equivalents for the street's value are retrieved, an extra check is imposed by specifying the entity's name, i.e., William, thus resulting in the geometric coordinates for the entity named William of type street.
However, if no country or city name is specified, the output contains all the possible results for that particular entity, making one of them a correct answer. The query structure is based on is-in relation, where a smaller entity is inside or contained in a larger entity. Thus while composing queries, entity types are determined. The larger entity name appears first in the query and then the type of smaller entity; thus, if we have an object named like the lunacy, it is specified in an additional check later.
In the case of a near relation, the query structure is different. Consider the sentence *post office near a train station in Bamberg*, and the first step is to identify the city and then look for train stations and post offices in that particular city. Once all the osm candidates are retrieved, the near function is used to retrieve all post offices situated close to Bamberg's train station. The near-specified concept identifies the distance between two entities, for example, 2 *meters or more*, and then identifies all possible post offices within the specified distance.

The output of these queries is in the form of geometric coordinates used to show the map's entities. The output for *Lunacy in/of St.pauli* provides us with a single result while in case of *post office near a train station in Bamberg* multiple options are retrieved.

Thus the query strategy provides us with an approach to simplify the automation process by defining constraint specific queries and running them one by one to retrieve the geometric coordinates and OSM id for the given entity by OSM. In our prototype, we use the Follium library to map the retrieved coordinates to OSM.

6.3.4.1 *Concluding results for Example*

In our running example, no city or country name is present; thus, we run our query by looking for inferred entities like St. Catherine church and Zeil street. Figure 17 shows the output obtained. As shown by the markers on the map, the system is able to identify the street, but as the church name is in German (*Katharinenkirche*), the query for St. Catherine does not match.[3]

3 Cross-lingual information retrieval techniques would be helpful here for a real system, but this is outside of the scope of this research

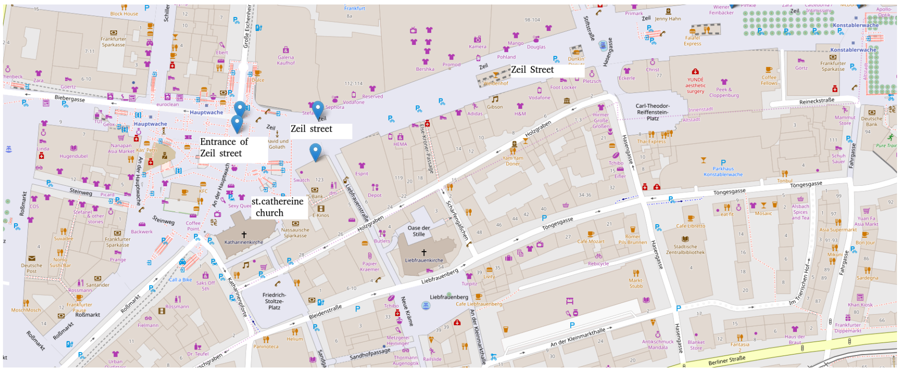

Figure 17: Indicating all geo-referenced spatial objects for running example.

6.3.5 *Limitations of presented approach*

This research presented how natural language words that designate spatial entity types (metropolis, city, creek, etc.) can automatically be translated into entities in the OpenStreetMap (OSM) database. We discovered two significant drawbacks of the provided technique, which are described in this section.

6.3.5.1 *Dealing with OSM key-value pairs which are correct but not accurate*

Querying the OpenStreetMap (OSM) requires describing the target object by assigning key-value tags to entities. The problem of identifying key-value pairs for querying OSM occurs in geographic information retrieval based on natural language text and is difficult for three reasons: First, the conceptualization of entities in natural language text and OSM often differ. Second, even classification of a single entity type is subject to variations throughout the OSM database. Third, language is rich and offers many words to communicate nuances of a single entity type.
OSM defines a set of standard keys and values that can be organized as a formal ontology like OSMonto [77]. Instructions are published for volunteers compiling OSM data on how spatial information should be tagged. For some real-world entities, it may be difficult for humans to identify the most appropriate key-value pairs, as some database variations occur. Moreover, the semantic model underlying OSM is an open tagging system with a recommended set of tags that provide users a catalog of core features.
It allows the contributors to add different key-value pairs for a particular entity. Words used as value fields may exhibit a very similar meaning, an example being "park" and "garden" which are both values of the key category "leisure". As a result, for a Botanical garden, the key-value pair could thus be expected to be *leisure=garden*, but we encountered the value "park" to be used too. In addition, sometimes different or undocumented terms are used, for example, instead of *amenity=parking*, *site=parking* is used, resulting in tags hard to anticipate.

The problem is to identify appropriate key-value pairs for any type of spatial category which occurs in our input text. We have to represent the spatial semantics of natural language as spatial concepts using the OSM taxonomy using key-value tags. The next section 6.4 focuses on how appropriate key-value pairs can be determined.

6.3.5.2 *Dealing with sentences having non-spatial information*

In addition, the existing approach works fine when the nature of place description is spatial. The categories of nouns present in a sentence can vary from names (a named entity like *London*) to type identifier nouns like a market, river, and so on. Other examples include abstract nouns that provide meta-level concepts like location, spot, place, position [7], and non-spatial place information such as place semantics, equipment, ethnicity, activities, or affordance.

The task of recognizing place names and identifying named places based on toponym resolution has received much attention and considerable progress has been made [229, 137, 252, 56, 230, 260]. The success of these methods hinges on the presence of named places for geoparsers to be effective. Text with few or no place mentions can hardly be handled. Given a natural language description like "Bao is my ideal lunch spot", it refers to a place where one can eat. Indeed, Bao is the name of (at least) one restaurant in the United Kingdom, yet it is also the name of a town in the Philippines. Using existing named entity recognition systems, Bao is resolved as the name of the town. These systems are unable to recognize Bao to be a restaurant as named entity recognition does not consider the ambiguity of locations and cannot handle the ambiguity of locations with the same label in a text [303]. The example shows that context information is required in place recognition, even to handle named entities.

In another example, "This Persian place on Elizabeth Street makes some of the best bread in the city" place represents an abstract concept that can either be a place where bread is made, a bakery, or some cafe. Given just the sentence, no definitive characterization of the place can be made. Several researchers have focused on identifying specific contextual factors that are postulated to have an impact on natural language interpretation [347].
One can argue that by using the abstract place concepts and details about typical food like what we can eat at a particular place, we can predict some contextual information that will ease place interpretation. There are numerous categories of contextual factors, out of which indications for place type can be derived. An explicit form is given in the form of the language phrase, and likewise ontological statement "is a" as in "Leeds is a city", which can be exploited to disambiguate interpretations and thus improve geo-referencing [405].
Thus we aim to advance the exploitation of context to paraphrased places in the coming section 6.5, by considering question:" How can information provided by a the sentence be used to determine an entity type of a place described?"

6.4 IDENTIFYING APPROPRIATE OSM KEY-VALUE PAIRS

This research is concerned with studying approaches that can map natural language words which designate a spatial entity type to appropriate key-value pairs for querying OSM. We investigate how well semantic similarity of words allows the correct key-value pairs to be identified. This study is solely based on data obtained from an input sentence, from a compilation of documented OSM tags, and WordNet for computing semantic similarity. We argue that even many intuitively phrased sentences are ambiguous in terms of appropriate key-value pairs and require a geo-referencing strategy in order to find matching map entities.

With our investigations, we are looking forward to these answers:

- To which extent does WordNet help to identify semantic similarity between words and OSM tags?
- What is the contribution of semantic similarity to improving a query strategy?

6.4.1 *Identifying Adequate OSM Tags using WordNet*

The main objective of this work is generating a set of tags adequate for querying for an entity described by a natural language word. Investigating Wikipedia sentences that describe geographic entities, we easily encounter situations in which the tagging required to retrieve the intended target object cannot be identified easily. We argue that a successful approach has to tackle three major challenges:

1. Acknowledging that conceptualisation of entities in text and OSM database may differ, since writers employ a commonsense understanding rather than a technical taxonomy.
2. Classification of a single entity type is subject to variations in OSM, since the semantic model underlying OSM is open-ended and variations across volunteers preparing the data are inevitable.
3. Language is rich and offers many words to communicate nuances of a single entity type.

In this work we propose a similarity-based ranking function to address these challenges. The use of a ranking function is motivated by the observation that the correct tag required to retrieve a target entity cannot be computed in a reliable manner due to inevitable variations. Instead, a ranking function allows us to start with a reasonable tag and to retry with a next-best option if a query is not successful.

6.4.1.1 *Semantic Similarity for Determining Tags*

In OSM, keys designate a certain category whereas values specify the respective element. Looking at OSM tags, we can identify that values belonging to the same category can be similar, for example "footway" and "path" as values for key "highway". However, similarity or even equality of values from distinct key categories can also occur.

For example, consider the text phrase from Wikipedia *"Seehof Schloss is located outside of Memmelsdorf (...)"*.[4]. The translation of German word *Schloss* is castle and, consulting OSM tags, we observe that both, *building=castle* and *historic=castle* are commonly used. Tagging instruction on the OSM website state that modern castles should be classified as buildings, whereas others are classified as historic. From the surrounding text it remains unclear which category is more appropriate. Thus, at time of query generation picking the most adequate tag remains ambiguous. In our example, the name of the castle, *name="Seehof"* would of course provide most valuable information, but in general we have to face situations in which entities remain unnamed. Even if entities are named – and even if their rough location is known – ambiguity still remains as several entities may share the same name.

4 https://en.wikipedia.org/wiki/Schloss_Seehof, last visit 2019/09/23

As a second example, consider the sentence *"The Holocaust Memorial is centrally located in Berlin Friedrichstadt district."*, abbreviated from English Wikipedia text.[5] In order to identify *Friedrichstadt district* within OSM, the key-value pair *place=district* seems most reasonable as there is a perfect match on the word level. However, *place=locality* is used to represent Friedrichstadt district in OSM. This exemplifies situations in which similarity information combined with category membership may help to resolve objects: if a key-value pair with a reasonable value does not match, other category members could be reasonable next-best options.

6.4.2 *Semantic Similarity using WordNet*

In order to identify an approach to determine key-value pairs automatically, we are motivated to study whether word similarity as provided by WordNet can help. WordNet [263, 116] is a special kind of digital English dictionary. Alongside with a definition of a word, WordNet presents distinct meanings associated with that word. These meanings are grouped and called *synsets* in WordNet. Each synset thus represents a distinct concept associated with the word. Words are assigned to (multiple) synsets and WordNet provides a semantic similarity measure among the distinct synsets. Similarity is measured as real-valued number in the range $[0, 1]$ with value 1.0 designating highest similarity.

As a first step, we gathered OSM tags from the documentation Wiki and recorded all key-value pairs in lists. The next step involves generating tag sets representing the key-value pairs that are associated with a given word representing an entity. To this end we constructed algorithms which address specific aspects of identifying reasonable key-value pairs. The algorithms are described in the following, experimental results and discussions are presented in the next section.

6.4.2.1 *Word-to-Value Similarity*

Given word E and a set of key-value pairs to consider, Algorithm 3 computes a list of tags sorted by the similarity of their tag value to the input word. Category information provided by the key is not considered. The algorithm is a straightforward application of WordNet similarity: for every word we look up its synsets to retrieve possible meanings, and compare them against each other. From the comparisons we choose the maximum similarity (line 6) of potential meanings.

6.4.2.2 *Word-to-Key-Value Similarity*

In a second algorithm we consider similarity to fields key and value. The algorithm is motivated by situations in which the same value is present in two different key categories. Applying Algorithm 3 would identify both key-value pairs as perfect matches (similarity yields 1.0 for identical words), but the order would be purely ambiguous. The idea underlying this algorithm is to consult the key field to resolve ambiguities. The algorithm is presented in Algorithm 4 and is a simple extension of Algorithm 3, computing similarity to key and value separately and summing up the results.

5 https://en.wikipedia.org/wiki/Memorial_to_the_Murdered_Jews_of_Europe, last visit 2019/09/23

Algorithm 3 WordNet similarity for entity E considering values within a set of tags (key-value pairs).

```
function TagsForEntity(E, Tags)
    S ← Synsets(E)                                   ▷ Collect all meanings of word
    T ← []                                           ▷ List of similar tags
    for (k, v) ∈ Tags do
        S_v ← Synsets(v)                             ▷ Meanings of value
        sim ← max_{s1∈S, s2∈S_v} WordNetSimilarity(s1, s2)
        Insert (k, v) into T with score sim
    end for
    return T sorted according to similarity score
end function
```

6.4.2.3 *Word-to-Key Similarity*

As a last algorithm we consider an algorithm that first sorts all tags according to their similarity to the key field and then sorts according to similarity to values within individual key groups. The algorithm is motivated by the observation that key-to-word similarity may be most informative in some situations. That principle is realised in Algorithm 5. For example, the interpretation "street" to *highway=secondary* as it occurs in our test data may more easily be derived from comparing "street" to "highway" which are similar words, but not from comparing dissimilar words "street" and "secondary". Algorithm 5 will in that situation rank all members of key category "highway", before including key-value pairs from other key categories.

Algorithm 4 WordNet similarity for entity E considering keys and values within a set of tags (key-value pairs).

```
function TagsForEntity2(E, Tags)
    S ← Synsets(E)                                   ▷ Collect all meanings of word
    T ← []                                           ▷ List of similar tags
    for (k, v) ∈ Tags do
        S_v ← Synsets(v)                             ▷ Meanings of value
        S_k ← Synsets(k)                             ▷ Meanings of key
        sim1 ← max_{s1∈S, s2∈S_v} WordNetSimilarity(s1, s2)
        sim2 ← max_{s1∈S, s2∈S_k} WordNetSimilarity(s1, s2)
        sim ← sim1 + sim2
        Insert (k, v) into T with score sim
    end for
    return T sorted according to similarity score
end function
```

In order to get better understanding of the motivation underlying the algorithms let us consider some examples.

6.4.2.3.1 EXAMPLE 1. Consider the sentence *"The smallest hotel in the world is in the town of Amberg known as «Little Wedding House»"*. Geographic entities in the sentence are hotel, town, named entity Amberg, and a compound name "Little Wedding House". As a result of successful text understanding, hotel should be identified to be a co-reference for Little Wedding House and town should be linked to Amberg. Mapping words to OSM tags, town can be mapped to *place=town* and hotel is ambiguously mapped to *tourism=hotel* or *building=hotel*.

Algorithm 5 WordNet similarity considering keys and values.

function $\textsc{TagsForEntityByCategory}(E, Tags)$
  $S \leftarrow \textsc{Synsets}(E)$ ▷ Collect all meanings of word
  $K \leftarrow []$ ▷ Set of similar keys
  $T \leftarrow []$ ▷ List of Tags
  for $(k, v) \in$ Tags **do**
    $S_k \leftarrow \textsc{Synsets}(k)$ ▷ Meanings of key
    $sim \leftarrow \max_{s_1 \in S, s_2 \in S_k} \textsc{WordNetSimilarity}(s_1, s_2)$
    Insert k with similarity sim into K
  end for
  Sort K according to similarity ▷ keys sorted by similarity
  for $k \in K$ **do** ▷ iterate over list
    $T_k \leftarrow \textsc{TagsForEntity}(E, \{(k', v') \in Tags | k' = k\})$
    Append T_k to end of T
  end forreturn T
end function

Amberg is a named entity and is represented with *name="Amberg"*.

As there are two categories available for Hotel, the possible solution here is to find similarity to the corresponding key fields as done in Algorithm 4. Doing so we obtain the ordered list [*tourism=hotel*, *building*, *hotel=*, ...] as hotel is semantically considered to be more similar to tourism than to building.

6.4.2.3.2EXAMPLE2. Let us consider another sentence from Wikipedia: *"The Berggarten is a botanical garden with the most varied collection of orchids in Europe"*. Starting with entity extraction, we have a garden, Europe and Berggarten. As "garden" is a member of category leisure, we have one tag attached to an entity. Despite being a reasonable OSM key-value pair, querying with *leisure=garden* does not provide us with the correct result. So we refer to Algorithm 3 and apply it for generating the tag set based on similarity measure to value fields within the same category of key "leisure", simply by supplying the key-value pairs for that key as input set of tags. The tag set looks like as shown in Table 13, the correct tag required to identify Berggarten is "park" which is ranked second. This example not only illustrates a potential utility of using semantic word similarity, but also shows that a specific strategy can be applied when employing the algorithms, for example by restricting the set of key-value pairs to consider. Computing the order of all key-value pairs and trying them out sequentially could possibly lead to a wrong interpretation of the input sentence.

6.4.3 *Effectiveness of identifying appropriate OSM Key-value Pairs*

To evaluate the results produced by our method, we consider the following evaluation criteria.

We aim to find out how effective the described methods are for determining the correct tag, i.e., to find the tag actually used in the OSM database as soon as possible. To this end we record the position in which the correct key-value pair occurs within the list of sorted tags produced by our algorithms. Lower position numbers indicate that the correct tag will be found earlier. For determining position numbers we follow a worst case assumption: if multiple tags receive the same similarity score and their order is thus

Table 13: Values for key "leisure", sorted by their WordNet similarity to "garden"

similarity	value
1.0	garden
0.82	park
0.63	pitch
0.53	playground
0.38	stadium
0.38	track
0.32	sauna
0.13	slipway

ambiguous, the correct tag is said to be the last of that set. Additionally, we record the similarity scores of the correct tags. Recall that similarity in WordNet is given in the range $[0, 1]$ with 1.0 indicating maximum similarity. Similarity values may be relevant to an overarching query strategy which has to decide when to give up on identifying a certain entity. Consistently high similarity values are thus desirable to ease such strategies.

6.4.3.1 *Baseline*

As a baseline, we have manually entered every word describing a spatial entity into *Tagfinder* and recorded whether it returns the correct key-value pair and, if so, at which position in the list. *Tagfinder* returns all possible key-value pairs which either have the input word in them or their definition. This includes words, which have the given word as substring in their spelling, e.g., "bar" is related to barrier, barn, barrack, bare rock, alongside the correct pair *amenity=bar*. The results obtained are shown in Table 14. In the table input word and correct key-value pairs are shown as well as the position of the correct tag in the result list. *Tagfinder* makes use of wildcards, e.g., *building=** to designate any value within the category represented by key "building". A matching wildcard is treated as correct result. Table 14 shows that *Tagfinder* is able in 5 out of 11 cases to list the correct key-value pair.

Out of 22 sentences in our corpus challenging key-value generation, 10 sentences fall in the category of Algorithm 4 where multiple categories for a single value are present in the tag set and by selecting the appropriate key we were able to obtain desired query results. Table 19 below lists the words and the categories associated with them. As we can see the value is present in two categories only, requiring only one additional query attempt at most.

Table 14: Determining tags with *Tagfinder*, position "—" indicates that the correct result is not listed.

entity in text	key-value in OSM	position in result list
park	*leisure=garden*	—
museum	*tourism=gallery*	—
castle	*building=**	5
district	*place=locality*	—
garden	*leisure=attraction*	—
street	*highway=secondary*	—
bar	*amenity=pub*	78
hub	*public_transport=station*	—
hotel	*tourism=hotel*	1
station	*railway=station*	2
office	*building=office*	5

Table 15: Words with multiple category associations

entity in text	categories (keys) in OSM tag sets
Castle	Building Historic
Hotel	Tourism Building
Station	Railway Publictransport
Office	Building amenity

For the remaining entities, as we were not facing multiple categories, the number of options for key-value pairs increases as we have to evaluate similarities to at least all values within a group. Table 18 presents results from Algorithm 3 applied to a single key category after the perfectly matching value failed. For example, consider the first line in Table 18. The word "park" is also a valid value in category "leisure", but the query for *leisure=park* fails. Instead, *leisure=garden* is the tag used in OSM which appears second in the list computed by Algorithm 3.

Out of 22 sentences in our corpus challenging key-value generation, 10 sentences fall in the category of Algorithm 4 where multiple categories for a single value are present in the tag set and by selecting the appropriate key we were able to obtain desired query results. Table 19 below lists the words and the categories associated with them. As we can see the value is present in two categories only, requiring only one additional query attempt at most.

Table 16: Results for finding alternative values within single class using Algorithm 3

word	category	correct value	similarity	position
park	leisure	garden	0.823	2
museum	tourism	gallery	0.133	9
district	place	locality	0.7142	9
gallery	tourism	attraction	0.133	4
street	highway	secondary	0.133	8
bar	amenity	pub	0.526	3
hub	public transport	station	0.5	2

Table 17: Words with multiple category associations

entity in text	categories (keys) in OSM tag sets
Castle	Building Historic
Hotel	Tourism Building
Station	Railway Publictransport
Office	Building amenity

For the remaining entities, as we were not facing multiple categories, the number of options for key-value pairs increases as we have to evaluate similarities to at least all values within a group. Table 18 presents results from Algorithm 3 applied to a single key category after the perfectly matching value failed. For example, consider the first line in Table 18. The word "park" is also a valid value in category "leisure", but the query for *leisure=park* fails. Instead, *leisure=garden* is the tag used in OSM which appears second in the list computed by Algorithm 3.

As we can see that for some words, the similarity measure is as low as 0.1333. The spread of similarity scores (0.133–0.823) and positions (2–9) indicates that WordNet similarity is effective but its performance hard to predict.

In order to make a more time-efficient query strategy, we conducted an experiment where we consider similarity to keys as presented in Algorithm 5. In Table 19 an association between selected examples of input word and the corresponding correct category (key) is shown. Based on these values, Algorithm 5 determines the order of categories in which values are sorted according to their word similarity. The results are presented in Table 20. As can be seen, this method improved the position measure when entity and correct values are neither synonyms nor share the same meaning. For example, since "street" is classified similar to key "highway", value "secondary"

Table 18: Results for finding alternative values within single class using Algorithm 3

word	category	correct value	similarity	position
park	leisure	garden	0.823	2
museum	tourism	gallery	0.133	9
district	place	locality	0.7142	9
gallery	tourism	attraction	0.133	4
street	highway	secondary	0.133	8
bar	amenity	pub	0.526	3
hub	public transport	station	0.5	2

Table 19: Table showing similarity measure between input and OSM keys

entity in text	correct category	similarity
park	leisure	0.133
museum	tourism	0.117
district	place	0.677
gallery	tourism	0.25
street	highway	0.823
bar	amenity	0.41
hub	public transport	0.45

gets considered for entity “street” earlier. Only evaluating similarity to the value field is not sufficient to arrive at this result.

In other cases like key category “tourism” with values “gallery” or “museum”, the range identified based on a similarity between key and word appears reasonable and helps us to construct a successful query in few attempts.

Table 20: Results after second evaluation

entity in text	correct key	correct value	similarity	position
park	leisure	garden	0.823	9
museum	tourism	gallery	0.133	3
district	place	locality	0.714	8
gallery	tourism	attraction	0.133	3
street	highway	secondary	0.133	8
public transport	hub	station	0.5	2

In Table 20 we can observe a decrease in the position for categories tourism/gallery and tourism/museum as compared to output of Algorithm 3 shown in Table 18. This indicates that a link between key and the entity as described in the input text exists and is helpful to exploit. Other values determined by Algorithm 5 remain the same as determined with Algorithm 3, except for “park” where the position number has increased. The reason is that we are no longer dealing with the restricted set of values from key category “leisure”, but Algorithm 5 considers all key-value pairs. Above all, the similarity value for leisure and garden is only 0.133, which also contributes to late consideration of the correct key-value pair.
From the experiments we can say that WordNet may play an important role in identifying reasonable key-value pairs for geo-referencing. We are able to resolve all queries given that it is acceptable to try up to 8 unsuccessful queries first. Although the overall scores do not speak for themselves, we regard the fact that correct key-value tags always appear in the top ten candidates to be a promising first step. In particular our second experiment shows that the approach using simple word-to-value similarities can be improved.This motivates further investigations to determine the most effective strategy, in particular to remedy problems in which the correct key/value keys are not directly similar to the entity.

6.4.4 *Summary*

We have implemented our approach as a research prototype in order to evaluate the contribution of WordNet for interpreting place descriptions, the experiments and results are discussed in the next chapter 7. The focus of our prototype is to improve querying in text-based geo-information retrieval by finding the correct tag at run time. Despite its promising first results, this approach has a number of limitations and open challenges. WordNet is a general-purpose semantic resource, and its coverage of geographic terms is limited. While the proposed mapping technique is effective with common terms (e.g., bay, city, university), it would not perform well with many technical terms in highly specialised vocabularies. To make things worse, words for tags can sometimes not be interpreted without considering OSM context. For example, key 'historic' serves as a special building category, i.e., it designates historic [buildings], whereas a general-purpose similarity function not sensitive to this context can easily fail to identify reasonable tag similarity. Another limitation of using WordNet arises from the fact that WordNet similarity is typically limited to within the same class of words, e.g., nouns to nouns. As such, no similarity between words "secondary" and "street" can be determined as would be necessary to pick *highway=secondary* as tag for street.[6]

Following conclusions can be drawn:

1. The WordNet similarity measure applied directly to geographic terms helps in posing successful queries and leads to an appropriate tag selection without further external knowledge being involved.
2. Using word-to-value similarity, the performance of WordNet is effective and accurate.
3. Performance of WordNet similarity is reasonable in case of common concepts such as amenity, building, tourism, and leisure. However, dealing with specialized vocabulary used for tags that must be interpreted in context of OSM, the performance is significantly worse.

6.5 DETERMINING ENTITY TYPES FOR QUERYING OSM

The last part of this chapter deals with paraphrased phrases or non-spatial information present in the place descriptions. The basic idea in this section is to generate semantically replaceable terms (which can be more than one) in the form of a type noun. For example, coffee will be linked to a cafe, coffee, and cafe linked to a restaurant, etc. By determining type nouns as replacements that correspond to types in a geographical database, we aim to improve the geo-referencing of spatial entities. Conceptual knowledge should ideally be contained in some semantic similarity measure, closely relating the terms coffee and cafe, and we can exploit this relation to make such inferences automatically. This research examines approaches that can map natural language words that designate a non-spatial noun to the appropriate place type using OSM type values. The idea is to generate a catalog of related types that will help while querying OSM. To this end, this paper investigates similarity measures and determines their effect on geo-information retrieval. We investigate how well semantic similarity of words allows related place types to be identified. Our study is solely based on data obtained from an input sentence crawled from the internet, a compilation of documented OSM tags, and different semantic similarity measures. We argue that even many intuitively phrased sentences are ambiguous in terms of entity types, and a context-sensitive geo-referencing strategy is thus needed. With our investigations, we are looking forward to these answers:

- To which extent can we exploit nouns present in a sentence to provide contextual information?
- To what extent does semantic similarity helps to identify semantically replaceable terms between words and OSM tags?
- What is the contribution of semantic similarity in improving the process of interpreting place descriptions?

6 In our algorithms, we use value 0.0 for undefined similarity measures.

6.5.1 *Problem Description and Approach*

The main objective of query formulation is to generate a set of OSM concepts adequate for querying for an entity described by a natural language word. Skimming over Wikipedia travel blogs that describe geographic entities, we easily encounter situations where it is challenging to identify the information required for mapping a target object to a precise type. We argue that a successful approach has to tackle these major challenges:

- Conceptualization of entities in text varies across sentences as writers employ an objective commonsense understanding rather than a scientific taxonomy.
- Classification of a single noun is subject to variations in OSM concepts since the semantic model underlying OSM is open-ended. Moreover, variations across volunteers preparing the data are inevitable.
- Language is rich and offers many words to communicate nuances of a single entity type. Conceptual boundaries of natural language words are not always aligned with those of OSM concepts.

Our approach is based on an explicit representation of non-spatial information extracted from the text. The method is divided into two primary modules; the first one is related to the extraction of nouns from the input, and the second is related to finding the semantic similarity between the extracted nouns and type nouns present in a lexicon of possible OSM concept.

6.5.1.1 *Noun Extraction*

The objective of the first module is to translate a single input sentence into a set of nouns that represent spatial and non-spatial information. In contrast to classic parsing, we do not aim to capture the full structure of the sentence, and we only retrieve noun phrases. Noun phrases can either be spatial entities like-named or unnamed entities or other nouns which provide further contextual information. For this extraction, the part-of-speech tagger from the Natural Language Toolkit (NLTK) [48] is used. If nouns follow after one another (e.g., "art" followed by "building"), then these are treated as one single noun[7]. Every token tagged as a noun phrase is then further analyzed, following a pragmatic approach to categorization, which is presented in 6. If the noun can be found in the geographic gazetteer GeoNames[8], it is categorized as a named entity. If the noun can be linked to a spatial category in OSM, it is categorized as an unnamed entity of the respective category. Furthermore, if none of those provide a match, it is added to a list of context words. The exploitation of context words is the subject of this work.

Consider for example the sentence "This place has great coffee". Using the above algorithm, we cannot identify a named or unnamed entity, but only the contextual words *place* and *coffee*. As humans we are likely able to infer that this place might be a cafe, based on the contextual information that it offers coffee. But as far as we know, such non-spatial context has not been exploited in a meaningful way in geo-referencing yet. We use the extracted context words as input for our second component, which automatically derives the most likely OSM concepts for the given input and can thus lead to identifying the place mentioned.

---

7 Currently, we ignore attribute adjectives, determiners, etc., contained in the noun phrase.
8 `https://www.geonames.org`

Algorithm 6 Noun extraction

function PROCESSINPUT(S) ▷ input sentence
  Apply part-of-speech tagging to S
  I ← ∅ ▷ information extracted
  for W is noun in S **do**
    if W is listed in gazetteer as spatial category C **then**
      Add {(W has-name W), (W is-a C)} to I ▷ named entities
    else if W is listed as spatial category C in OSM lexicon **then**
      Add {(W is-a C)} to I, using a normalized type name ▷ unnamed entities
    else if W is an abstract concept or non-spatial noun CN **then**
      Add {(W occurs-in-context CN)} to I ▷ context nouns
    end if
  end for
  return I
end function

6.5.2 *Inferring Spatial Knowledge from Non-Spatial Context Words*

Given a number of previously identified context words, this module is concerned with finding the semantically closest related OSM concepts in the form of key-value tags. Previous work often relies on WordNet [263] to establish semantic relationships [33, 406]. We thus relate our approach to an approach using WordNet in order to identify its potential and possible weaknesses (Section 6.5.2.1). In a proof-of-concept, we implemented two approaches based on Word2Vec [262] word embeddings: Noun-to-value similarity (Section 6.5.2.2) and value-to-value similarity (Section 6.5.2.3).

6.5.2.1 *Using WordNet as a Baseline*

Following the research of [406], we used WordNet as a starting point to calculate the similarity measure between context nouns and OSM tag values. WordNet [263] is a special digital English dictionary. Alongside the definition of a word, WordNet presents distinct meanings associated with that word. These meanings are grouped and called synsets, where each synset represents a distinct concept. Words accordingly can be assigned to one or more synsets.

WordNet further builds an ontological structure for nouns and verbs using a hierarchy of *is-a* relations [286]. Using this structure, similarity can easily be computed as a distance measure in the conceptual structure. Additionally, each word sense (synset) also comes with a short definition or gloss. Using this information, further similarity measures can be offered. Based on those two methods, initially six similarity measures for WordNet have been proposed [286]. Three of them are based on the path lengths between different concepts (*lch*, *wup*, and *path*) and the other three use the information contained in the concept definitions (*res*, *lin*, and *jcn*). Further, there exist three measures of relatedness that aim at expressing related concepts (*hso*, *lesk*, and *vector*).

For an initial test, we extracted a small set of OSM tag values and compared all similarity and relatedness measures summarized in [286][9] to a number of non-spatial context words. We quickly noticed that the similarity

9 Only the *vector* relatedness measure is not considered, as it is not implemented in the used toolset WordNet Similarity for Java (WS4J) (`https://code.google.com/archive/p/ws4j`).

measures returned inconsistent results regarding our use case, or did not render any meaningful results overall. To demonstrate this, we use 10 arbitrary OSM tag values from different domains and compare the similarity and relatedness for two different context words, namely *lunch* and *bread*. As possible OSM tag values we choose *hotel*, *bakery*, *cafe*, *garden*, *bridge*, *library*, *picnic-table*, *fast-food*, *church*, and *university*. In case of multiple word senses, we automatically select the one with the highest similarity value. The results are displayed in Table 21 and Table 22, respectively. For all measures, higher values represent a higher similarity. Values for *lch*, *res*, *jcn*, and *lesk* are in the interval of $[0, \infty]$, *wup*, *path*, *lin* in the interval of $[0, 1]$, and *hso* in $[0, 16]$.

Table 21: Comparison of WordNet similarities for lunch

OSM tag value	lch	wup	path	res	lin	jcn	hso	lesk
hotel#n#1	1.124	0.333	0.077	0.614	0.072	0.063	0	23
bakery#n#1	1.124	0.333	0.077	0.614	0	0	0	14
cafe#n#1	1.105	0.316	0.071	0.614	0.065	0.057	0	24
garden#n#1	1.050	0.316	0.071	0.614	0.067	0.059	0	19
bridge#n#1	1.204	0.353	0.083	0.614	0.065	0.056	0	**38**
library#n#5	1.124	0.333	0.077	0.614	0	0	0	13
picnic-table	-	-	-	-	-	-	-	-
fast_food#n#1	**2.303**	**0.824**	**0.250**	**6.782**	0	0	**4**	8
church#n#2	1.050	0.316	0.071	0.614	**0.073**	**0.064**	0	21
university#n#2	1.124	0.333	0.077	0.614	0.069	0.060	0	20

Table 22: Comparison of WordNet similarities for bread

OSM tag value	lch	wup	path	res	lin	jcn	hso	lesk
hotel#n#1	1.204	0.353	0.083	0.614	0.075	0.066	0	72
bakery#n#1	1.204	0.353	0.083	0.614	0	0	0	63
cafe#n#1	1.124	0.333	0.077	0.614	0.067	0.059	0	49
garden#n#1	1.124	0.333	0.077	0.614	0.070	0.061	0	114
bridge#n#1	1.291	0.375	0.091	0.614	0.067	0.056	0	**149**
library#n#5	1.204	0.353	0.083	0.614	0	0	0	38
picnic-table	-	-	-	-	-	-	-	-
fast_food#n#1	**1.897**	**0.706**	**0.167**	**5.463**	0	0	**2**	20
church#n#2	1.124	0.333	0.077	0.614	**0.076**	**0.067**	0	140
university#n#2	1.204	0.353	0.083	0.614	0.071	0.063	0	94

It can be noted that some OSM tag values like *picnic-table* do not yield an entry in WordNet and thus no similarity can be computed. For other vastly different concepts (e.g. *bakery* and *library*, or *church* and *cafe*), all or most values are identical, making the similarity measure almost meaningless. Only fast-food renders surprisingly high results, making it the potentially best match for both terms. We argue that an adequate measure should at least render higher values for the combination of *bread* and *bakery*. Lastly, the relatedness measure *lesk* always returns highest values for bridge, which is objectively not the most related term to any of the context words. Similar observations were made with other context words and also when including further OSM tag values.

6.5.2.2 *Noun-to-Value Similarity*

As an alternative to WordNet, we first designed a simple algorithm based on the popular Word2Vec[10][262] embeddings. Word embeddings like Word2Vec map words to real-valued vectors in a way that synonyms are mapped to vectors that are close with respect to Euclidean distance. Large text corpora are used to train the mapping function. We rely on the pretrained model contained in Spacy, particularly the medium model for English (en_core_web _md)[11]. Using this model, we then retrieve the vector representation of all input tokens. They are then individually compared to all available OSM values from the tags, and the respective vector similarities are computed. Before further processing, all results for a single word are normalized such that the most similar OSM tag value set to 1.0 (maximum similarity). This step is motivated by the assumption that each place must correspond to some entity in the OSM database.

For example, consider again the sentence "This place has great coffee". We first extract the nouns of the sentence, giving us the context words *place* and *coffee*. For each, the Word2Vec similarity to all OSM tag values is calculated. For example, *coffee* is most similar to OSM tag values *cafe* (0.65), *drinking-water* (0.53), and *bar* (0.51), and least similar to *track* (0.11), *bunker* (0.11), and *planetarium* (0.06). *place* is more ambiguous and thus yields more random results with highest values for *picnic-site* (0.56), *camp-site* (0.55), and *food-court* (0.53), and lowest values for *icecream* (0.20), *breweries* (0.15), and *planetarium* (0.12). These values are then normalized, such that the highest value represents 100% (i.e. cafe and picnic-site both are 1.00). The two values for each tag are then added and the result is the final "score" of the tag. Using this method, the best five tags are *food-court* (1.66), *cafe* (1.65), *fast-food* (1.64), *picnic-table* (1.63), and *drinking-water* (1.58). When querying for the described entity, the OSM tag value with highest similarity would be used. If no match can be obtained, the next best alternative is tried and so forth.

6.5.2.3 *Value-to-Value Similarity*

This algorithm is motivated by the need to handle words that are similar to several OSM values. In situations where a similarity measure fails to single out likely candidates, one might fail to identify a place because trying to many options is not feasible. The idea underlying this algorithm is to reduce the number of candidates by adding a pruning phase to resolve ambiguities, which yields a more compact set. We calculate the similarity at two different levels. The first level is that of noun-to-value similarity, but by introducing a pruning phase we reduce the number of candidates. Algorithm 7 presents the pruning phase in detail.

---

10 https://code.google.com/archive/p/word2vec

11 https://spacy.io/usage/vectors-similarity

Algorithm 7 Noun-to-value Similarity using Word2Vec Model

1. Calculate the similarity measure between all ContextNouns and OSM values. More precisely, given a noun CN_1 and a list of potential OSM values v =(v_1, v_2, v_3, v_4,.... v_N) we compute the similarities of (CN_1, v_1), (CN_1, v_2), (CN_1, v_3), (CN_1, v_4) till (CN_1, v_N).

2. We will store the semantic similarity measure(sim_M) as a list

```
["CN_1": v_1, "sim_M", "CN_1": v_2, "sim_M" ... v_N, "sim_M"]
["CN_2": v_1, "sim_M", "CN_2": v_2, "sim_M" ... v_N, "sim_M"]
```

   a) Normalize the data on basis of similarity value (sim_M) for all CN

   b) Compute Mean for CN_1 based on similarity values (sim_M) and repeat the process for all ContextNouns

$$A = \frac{1}{N} \sum_{v=1}^{N} v_i = \frac{sim_{M1} + sim_{M2} + \cdots + sim_{MN}}{N}$$

   c) IF (sim_M) is greater than mean for v_1 then insert v_1 in a list. Repeat the process for all values till v_N for CN_1 till CN_N and save them in list respectively.

```
["CN_1": values, "CN_2": values, .... "CN_N": values]
```

   d) Check if the value is present for CN_1 AND CN_2 AND CN_N Add it into the list of Inferred-types If the value is only present for CN_1 AND CN_2 and not for CN_N Ignore the value Repeat the process for all context nouns.

3. Output List of Inferred-types

The second level is to calculate the semantic similarity measure between the set inferred-types generated in initial pruning phase. Even after pruning being applied in the above step, we still observed an unmanageable large candidate set of 25 to 39 in some circumstances. Thus the idea of this step is motivated by the observation that all these type values are also related to one another. We thus explore whether a generation of clusters or classes containing types closely related to one another will reveal information that allows us to improve predicting types. The second part of Algorithm 8 explains the steps in details for clustering process.

Algorithm 8 Value-to-value Similarity Between Inferred-types

1. Calculate the similarity measure between all values in Inferred-types. More precisely, given a value v_1 and a list of potential OSM values v =(v_1, v_2, v_3, v_4,.... v_N) we compute the similarities of (v_1, v_1), (v_1, v_2), (v_1, v_3), (v_1, v_4) until (v_1, v_N).
   Repeat the process for all values (v_i, v_N).

2. We will store the semantic similarity measure (sim_v) as a list.

```
["v_1": v_1, "sim_v", "v_1": v_2, "sim_v" ... v_N, "sim_v"]
["v_2": v_1, "sim_v", "v_2": v_2, "sim_v" ... v_2, "sim_v"]
```

   a) Compute Mean for v_1 based on similarity values (sim_v) and repeat the process for all values

$$A = \frac{1}{N} \sum_{v=1}^{N} v_i = \frac{sim_{v1} + sim_{v2} + \cdots + sim_{vN}}{N}$$

   b) IF (sim_M) is greater than mean for v_1 then insert v_1 in a list. Repeat the process for all values from v_1 until v_N and save them in list respectively.

```
["v_1": values, "v_2": values,....  "v_N": values]
```

   c) Count the number of times v_1 is related to all values until v_N

```
["v_1": 10, "v_2": 4, "v_3":10,....  "v_N": 12]
```

   Repeat the process and count for all values
   Save the values and their number count in a list.

   d) Sort the values in ascending order on basis of their number count
   Values with the same number form a cluster

```
["v_N": 12, ["v_1": 10, "v_3":10,]....  "v_N": 12]
[position v_N:1, position v_1,v_3:2, .... position v_2:10]
```

   The higher the number, the closer the relation between the values

3. Output clusters of closely related type values

In order to get a better understanding of the motivation underlying the algorithm, let us consider an example. Given the sentence "I came across an Italian place whose bread is to die for in Berlin.", the extracted nouns consist of *bread* and *place*, along with *Berlin*. Using noun-to-value similarity, we get these best five values: *food-court*, *fast-food*, *bakehouse*, *bakery*, and *picnic-table*. These values are correct concerning food items but are not appropriate as for this sentence the required output is a *cafe* which, in noun-to-value similarity, appears at position 25 in the sorted list of all candidate types.

Using value-to-value similarity, the type nouns are first reduced to a set of 12 OSM type values which include *bar*, *cafe*, *farm*, *bird-hide*, *drinking-water*, *house*, *houseboat*, *restaurant*, *food-court*, *fast-food*, *garden*, and *picnic-table*. Then again, we find how closely these 12 types are related to one another by means of Algorithm 8. The output indicates that we get clusters of terms that are related to one another. A higher number will indicate that it might be an appropriate OSM tag value for our input sentence. Out of these 12 types, *picnic-table* is related closely to 6 and therefore to the most other type values, *garden*, *food-court*, and *fast-food* are each related to 5 other types, *house*, *restaurant*, *houseboat*, *drinking-water*, and *bird-hide* are related to 4, and finally *cafe*, *farm*, and *bar* are related to 3 other types.

Putting these clusters in a list, we obtain [[picnic-table], [garden, food-court, fast-food], [house, houseboat, drinking-water, restaurant], [cafe, farm, bar]]. In the list of clusters, the correct output *cafe* can be found at position four. Thus, the second algorithm provides us results based on the noun-to-value plus the similarity between those values.

We have implemented our approach as a research prototype to evaluate the contribution of using semantic similarity to exploit context in the interpretation of place descriptions. The focus of our prototype is to improve querying and resolve ambiguity in text-based geo-information retrieval by finding the correct OSM tag.

6.5.3 *Experimental Setup and Data*

For a first evaluation, we collected a corpus of 100 sentences that contain place descriptions from English Wikipedia by scanning the summary part from articles about geographical entities along with different place descriptions from travel blogs. For now, we focused mostly on texts about places related to food.

Further, we are specifically interested in sentences containing contextual non-spatial information. Taking those context words, we then evaluate how well we can infer the correct OSM tag value. From the corpus of 100 sentences, 70 include a named or unnamed entity which we are able to directly match with an entry in GeoNames or OSM (cf. Section 6.5.1.1). We thus focus specifically on the 30 sentences where spatial inference is not feasible with existing methods.

The evaluation of both proposed methods is conducted analogous. We first execute the algorithm on the given input sentence. As a result we retrieve a sorted list of possible OSM tag values. We then identify the position at which the correct OSM tag is listed. Lower position numbers indicate that the correct type will be found earlier. In case multiple tags receive the same score and the order is thus ambiguous, we follow a worst-case assumption as such that the correct tag is said to be the last of this set.

As a baseline, we have manually annotated the input sentences with the correct OSM tag. This was done by querying OSM using the original named entity, taken from the source of the sentence. We then extracted the nouns according to the method described earlier. The annotated sentences are displayed in Table 23. For each sentence, the corresponding context words are highlighted and the correct OSM tag is given.

The highlighted nouns indicate that there are three different kinds of combinations of context words. In the case of bread and place, the first one is the name of the food, and the second one is representing a more general concept about places. The second combination contains only words that are representing a specific context, like bread and pasta, where both are a type of food and thus belong to the same category. The third combination is a result of more abstract words not directly related to place descriptions like spot or location, but rather to a general food category like breakfast and

Table 23: Annotated sentences with correct OSM tag

Sentence with highlighted context words	**Correct OSM tag**
This cake shop is one of the most photographed places in Belgravia, and a fun one to have tea and cupcakes with friends.	cafe
This Persian place on Elizabeth Street makes some of the best bread in the city.	bakery
There is an excellent lunch place in Maxvorstadt which features a smart and super casual menu that's vegan	restaurant
From good bread to great pasta, Belgaravia has a lot going for it when it comes to food.	restaurant
Sironi is an Italian place with bread to die for.	cafe
A very nice breakfast and lunch place with a typical Berlin vibe is located along Reichenberger Strasse.	restaurant
Bao is my ideal lunch spot	cafe
From the Aussie breakfasts to the winning cocktails, Fitzrovia has something delicious for every hour of the day.	cafe
If in the mood for pizza one can go to this Canadian place in Berlin with great pizza you can buy by slices.	fast-food
I did order a breakfast meal. My avocado sandwich had a small twist and was delicious. Poached egg was just like I like it.	cafe

lunch. We are interested to see if those differences in the input will affect the prediction of the correct tag.

6.5.4 *Discussion of Results*

Both algorithms were executed on the same set of sentences to allow for a direct comparison of the different methods. The results are displayed in Table 24, where each line represents a sentence with its extracted context words, the correct OSM tag, and the position in the sorted result list for both the noun-to-value (NV) similarity and value-to-value (VV) similarity method.

Before discussing the differences between the algorithms, we analyze the results for both individually. The noun-to-value similarity approach was able to identify the correct OSM tag at the first position for a combination with one abstract noun (*lunch* and *vegan*) or both abstract nouns belonging to the same category (*breakfast* and *lunch*). A similar constellation with *breakfast* and *cocktail* however only finds the correct tag at position 6, which is probably due to the arguably more unrelated context word combination. For the specific words related to food, the best position is 4 in case of *tea* and *cupcakes*, but drops to 10 in case of *bread* and *pasta*. In other cases, the correct value is identified at a position somewhere as high as 4 all the way down to 25.

For the value-to-value similarity method, we first notice that a value for *place* and *bread* was not found. This is due to the pruning phase in combina-

Table 24: Results

context words	OSM tag	NV sim.	VV sim.
tea, cupcakes	cafe	4	**3**
place, bread	bakery	**4**	-
lunch, vegan	restaurant	**1**	6
bread, pasta	cafe	10	**2**
bread, place	cafe	25	**4**
breakfast, lunch	restaurant	**1**	8
lunch, spot	cafe	**10**	11
breakfast, cocktails	cafe	6	**2**
pizza, slices	fast-food	**7**	**7**
breakfast, meal, avocado sandwich, eggs	cafe	12	**5**

tion with one highly ambiguous and one more specific word. Both words in such a scenario can return vastly different similarity values for the same OSM tag values, and as result the correct value can be pruned when both are combined. This however is not always the case, as the combination of the same context words (*bread* and *place*) for a different target tag can achieve position 4.

Comparing both algorithms, we can see that the numbers are improved in cases where the words are belonging to specific food categories, either being *tea* and *cupcakes* or *bread* and *pasta*. Even in the case where one noun is abstract and the other is specific but both belong to the same category of food like *breakfast* and *cocktails*, the numbers are improving. As the types extracted for both nouns belong to the same category, clustering will further favor values within this category and contribute to a higher position for the correct answer. However, as mentioned above, this clustering of values does not render helpful in cases where the nouns belong to different categories like *place* and *bread*. Here the output type after pruning at the initial level does not include bakery, resulting in no correct answer. On the other hand, *place* helps to identify more general places such as *cafe*, which provides a huge improvement over the noun-to-value similarity. Similarly, *lunch* and *spot* also belong to different categories, hence no improvement can be achieved.

From those first experiments, we can conclude that word embeddings like Word2Vec can play an essential role in identifying OSM tag values, which can benefit the task of georeferencing. In most situations, we were able to detect the correct tag within the first few candidates. Using pruning and clustering techniques, we were able to improve reults compared to the simple noun-to-value similarity method.

In general, we can conclude that the simple noun-to-value similarity approach is more effective in cases where the categories are more general either for food or for the place. If words are more specific, the value-to-value similarity showed better performance. These results motivate further investigations to determine the most effective strategy. For example, maybe a combination of both methods can produce better results consistently in all three identified categories of nouns.

6.5.5 *Summary*

We can conclude that word embeddings like Word2Vec can play an important role in identifying OSM tag values, which can benefit geo-referencing. In most situations, we were able to detect the correct tag within the first few candidates. Using pruning and clustering techniques, we improved results compared to the simple noun-to-value similarity method.

In general, we can conclude that the simple noun-to-value similarity approach is more effective in cases where the categories are more general, either for food or for the place. If words are more specific, the value-to-value similarity showed better performance. These results motivate further investigations to determine the most effective strategy. For example, maybe a combination of both methods can produce better results consistently in all three identified categories of nouns.

6.6 DISCUSSION AND SUMMARY OF THE CHAPTER

This chapter presents a deductive method based on logic programming techniques that, given a locative phrase, a spatial lexicon, and a spatial database, can interpret the locative phrase, i.e., to geo-reference all spatial entities mentioned in the phrase. A specific feature of our approach is that it employs an explicit model of context using dedicated *context variables* [398]. During geo-referencing, spatial reasoning aims to simultaneously match spatial entities to ground objects to match context variables to specific contexts. We obtain a context-sensitive interpretation by explicating conceptual decisions and enforcing jointly consistent decisions within one interpretation task.
On the one hand side, our approach relies on a formal model of spatial language, while on the other hand, the interpretation process simulates a cognitive language interpretation process on a functional level. Therefore we expect algorithmic techniques described in this chapter to be useful to empirical researchers who test specific spatial models or steps in the interpretation process. We implemented the automated reasoning method using these sources in order to arrive at a spatial interpretation of the constraints. To make the approach efficient, a query strategy has been suggested based on the entity type present in the set of queries that will help avoid costly queries by serializing queries and focusing on reasonable candidate locations. Besides, the presented approach can be applied to both questions and declarative statements to get the desired output.

This chapter investigates the problem of identifying key-value tags for querying OSM from natural language input. We consider the semantic similarity of words to keys and values as a method for identifying tags. Based on the intuition that similar terms tend to be defined using similar tags, the proposed approach to compute the semantic similarity of key-value pairs is made using WordNet.We present a strategy to identify key-value pairs for a word from natural language text using WordNet and analyze its effectiveness. With its limitations and some first promising results, we are motivated to further study automatic tag generation.

Lastly, we investigate how contextual information can be exploited to help geo-referencing places, which are only paraphrased in natural language place descriptions (“a place to eat near river Regnitz”). For identifying those places in the OSM database, one requires OSM type names to query for the object (an entity of type restaurant near river Regnitz, for example). We consider the semantic similarity of words to OSM type names to infer and predict the most likely OSM type of the place in the OSM database. Natural language place descriptions can be ambiguous and imperfect, so a perfect interpretation is not possible in general. Approaches can only determine a set of (ranked) candidates, aiming to assign a high likelihood to the correct type name. Word2Vec-based similarity provides us with reasonable results without further external knowledge being involved. Using noun-to-value similarity as introduced in this paper, we obtain approximately 65 to 70 percent correct answers, and they are also present in the top ten sets of candidates determined. By adding a pruning phase and using value-to-value similarity, we can improve correctness to almost 80 percent.

Open research areas involve studying other means to compute semantic similarity tailored to geographic concepts. We also plan to investigate algorithms for computing tags based on word similarity to better address specifics of OSM tag names and their organization. The aim is to explore further ways of exploiting context and integrate this method within an automated system for the interpretation of place descriptions in order to determine the overall effect of contextual information. We are nevertheless motivated to continue our work as it can contribute to advanced capabilities of geo-referencing.

7 EXPERIMENTS AND EVALUATION

The experiments presented in this chapter are based on the content from the previously published paper by Yousaf and Wolter [407] in an interdisciplinary journal- Spatial cognition and computation in 2021. Furthermore, the experiments from a publication by Yousaf and Wolter in 2018 [405] are also reorganized and presented here. In addition, it includes some experiments from work under review as well. The chapter extends the experiments and defines the evaluation criteria and success measures in detail. The contributions of the corresponding first author include designing the experiments and identifying the measures which should be used to identify the impact of the presented system (SORS) through discussion with my supervisor.

A Spatio-ontological reasoning system (SORS) comprises three modules. Two are related to the geographic information retrieval phase, where declarative statements and the spatial entities are extracted and then used in the third and final phase of the system, i.e., querying an OSM database. This chapter focuses on identifying the impact of reasoning techniques presented in this research study. We first provide an overview of how to evaluate geographic information systems and what insights can be expected. The chapter then explains the various experiments conducted on SORS and provides details about the experimental setup and evaluation data. Finally, the last section of the chapter discusses the results of experiments and highlights how the system performs and contributes to the existing research gaps.

7.1 MOTIVATION AND GOALS

GIS focuses on the importance of *where things are: what they do there, how they got there, who put them there, what they connect to, and what their initial or residual values are* [423]and a key remaining question is establishing how well the GIS answer these questions. Numerous studies on the advantages of GIS have been conducted, including [141],[362]. In a study by [273] and [272], this is referred to as the impact of GIS. Thus, following the development of a GIS framework's required components, the next step is to assess and discover the effectiveness of GIS (our system) and determine how well the presented method performs its objectives [188]. An evaluation may be used to determine the impact. According to [273],[272],[138], there is still skepticism about whether GIS and related technologies are delivering on their commitments to society. In order to determine whether systems meet their intended objectives, it is crucial to develop or identify a mechanism for measuring the success derived from implementing such systems [219]. Nevertheless, to build reliable, effective, efficient, and functional systems that allow aspects of the search process and its outcomes to be quantified and measured, evaluation activities are required. It allows to benchmark, monitor results and investigate the potential benefits of new features, and make improvements [188].
The word evaluate has been defined as "to judge or calculate the quality, im-

portance, amount, or value of something [1]." IT evaluation is characterized by Farbey et al. [113] as "a process, or group of parallel processes, that occurs at different points in time or continuously, for searching and making explicit, quantitatively or qualitatively, all the impacts of an IT project and the program and strategy of which it is a part." Evaluation is conducted by matching a set of parameters to a set of requirements or benchmarks. Evaluation provides data for informing a wide range of stakeholders about a project's success or failure [123]. The importance of projects can be calculated using the data obtained during the assessment process. According to Remenyi et al.[300], evaluation is an essential activity in ensuring information systems success. According to [306], the following questions should be considered while planning any evaluation,

- What is the evaluation's goal/purpose
- What should be the performance metrics and measures
- How should the evaluation be performed.

The evaluation aims to determine the success of the implemented methods as well as the lack of success.

This chapter identifies the metrics and approaches used to assess the impact and efficacy of our framework. For decades, evaluation has been an active investigation area in information retrieval (IR), especially experimental evaluation [158]. However, the majority of studies have failed to include specific methods for GIS evaluation [219]. This research aims to bridge the gap between theoretical information retrieval concepts and how they can assess GIS. Thus, We consider evaluation in the context of information retrieval in general when evaluating geographic information systems [188].

The evaluation has been conducted on two different levels. The first evaluation is conducted to assess how spatio-ontological reasoning affects the process of triplet extraction and compares the amount of ambiguities introduced by our over-generalizing method for information extraction to wrongly identified references by parsers. The second set of experiments aims to assess the contribution of spatio-ontological reasoning and the effectiveness of our general approach, and it does not aim to measure performance. We only compare performance measures between different experimental conditions or between different kinds of entities. This evaluation is conducted on two levels: in the first, we investigate the overall ability of our approach to geo-reference places correctly by considering both named and unnamed entities—the second analysis singles out performance concerning the gist of place information known as primary entities. Besides we also measure the average number of queries that get enabled due to the interpretation.

The last experiment aims to compare the performance of our system against that of available approaches. Since, to the best of our knowledge, there is no ready-to-run system that can handle unnamed entities, we restrict this experiment to named entities by comparing our system with existing named entity recognition systems.

The remaining chapter is arranged in the following manner. Section 7.2 provides us with an overview of evaluation in information retrieval, includ-

1 Cambridge Dictionary: http://dictionary.cambridge.org/dictionary/english/evaluate

ing the success criteria and measures and how these methods can be used to evaluate GIS. Section 7.3 provides the evaluation criteria for SORS and explains different kinds of experiments conducted to evaluate the system's performance.Section 7.3.4 discusses the results of the evaluation and summarizes the results.

7.2 EVALUATION METHODS

As per our knowledge, no specific criteria or methods has been specified on how to evaluate a GIS. Therefore,this section explains how the evaluation is conducted in information retrieval systems and then explains how these concepts and evaluation criteria can be borrowed to analyze and assess the performance of geographic information systems.

7.2.1 *Evaluation in Information Retrieval*

People use IR systems to solve problems, complete assignments, and are part of a broader context that should be considered while evaluating [100]. According to [158], the notion of IR evaluation is defined as a process or methodology (e.g., a controlled experiment, such as a lab-based experiment, or controlled online experimentation, the use of test collection); the need to establish criteria against which success can be measured, which may include the use of measures or metrics to quantify the criteria; and the need to define goals and objectives: typically system or program goals but could also include the goals of the evaluation, such as research hypotheses [366],[365]. Multiple methods are needed to evaluate IR systems holistically, both during the implementation of the prototype or its sub-systems (formative evaluation) and at the ending of development (summative evaluation) [188]. There may be on-going monitoring and assessment practices in developing operating processes, usually through developing relevant so-called Key Performance Indicators [188].

According to [390], as IR systems move beyond only providing a search box to providing a rich array of functionality (e.g., categories and facets, suggestions, visualizations) to help the user's search and discovery, evaluation becomes more critical. In these instances, merely treating evaluation as a method for assessing query–results is not enough. However, in IR assessment, evaluating the search result's quality remains a dominant practice [188]. There has been continual debate about performance measures, including how search result quality affects user satisfaction or predicts task success [366],[12]. In IR, the following levels are frequently distinguished: (1) evaluation within the IR framework context; (2) evaluation within the job context; and (3) evaluation within the socio-organizational and cultural context ;(4) evaluation within the information-seeking context [317],[181]. Considering these levels helps determine to what extent users and the broader context are included in the system and are taken into account, partially exploring the 'what' aspect of the evaluation project. At different levels, the notions of relevance change are significant since one of an IR system's primary objectives is to return information relevant to the user's information need [188].

The predominant approach to IR assessment has been system-oriented for decades. System-oriented approaches concentrate on retrieval algorithms and their outputs, such as the ability to differentiate relevant from irrelevant documents and rank the results effectively (considering the quality, diversity, and position of relevant documents in the ranked list). This evaluation method employs test collections, which allow systematic and repeatable evaluations to be conducted in a controlled environment [72],[316],[155],[74]. On the other hand, user-oriented evaluations strive to assess how well the system as a whole helps a user's search and retrieval of information [52],[195],[386]. Criteria for evaluating retrieval systems usually focus on how well users accomplish their goals or assignments and their performance and satisfaction with search results [195],[166]. Comparing findings from different systems side by side is an alternative form of user study [356],[325], [168].

Once the what aspect has been identified, the next step is to determine how the answers or evaluation results can be obtained. The following section provides us with an overview of evaluation methodologies, success measures and criteria that help us establish how well our system performs.

7.2.2 *Existing Evaluation Methodologies,Criteria and Success Measures*

This section provides an insight into how various methods contribute to collecting data and using it for evaluation. White in [386] describes nine common approaches for evaluation processes: field experiments, instrumented committees, log analyses, interviews and focus groups, laboratory studies, crowdsourcing, surveys, online methods like A/B research, and offline methods Cranfield Model and simulations. In practice, the methods used can differ based on whether the evaluation is user or system-based, the type of data available, and if the evaluation is being performed on an operational system or not. Offline testing uses three main methodologies the Cranfield Model and test collections [155],[316], user testing [52],[195],[386], and online testing [168]. Whether evaluating a system, a user, or an operating system, an experimental setup would be used, containing performance measures (dependent variables) and various systems/components under test(independent variables) [188]. The test collection may provide a controlled experimental environment where the dependent variable (i.e., precision) is calculated based on search performance in systems-oriented evaluation. Different systems, algorithms, and parameters are examples of independent variables. In the case of a user-based design, dependent variables could include measures of the user's search process and technique and measures of the search output such as precision and user satisfaction. Different systems or components may be independent variables [188].
These dependent and independent variables provide us with the criteria for evaluation. Cleverdon et al. [73] suggested six criteria for evaluating an IR system: (1) coverage, (2) recall, (3) precision, (4) time lag, (5) user effort, and(6) performance presentation. Precision and recall (see [325]), which capture an IR system's ability to distinguish relevant from non-relevant records, are the most well-known measures. The proportion of relevant documents retrieved is calculated by precision, while the recall measures the proportion of relevant documents retrieved. Su [349] defines four key evaluation criteria for web search engines: (i) relevance (ii) efficiency (iii) utility and (iv) user satisfaction.
White [386] describes two main categories of performance measures in his book on interaction in search systems: (i) those based on the search process and (ii) those based on the search outcomes after the search process is complete. In [195], Kelly presents four categories for measures: (i) perfor-

mance measures: these are used to evaluate the retrieval process and typically reflect usage activities (e.g., Precision, Recall, P10, GMAP, BPref, DCG, response time, informativeness, cost and utility measures); (ii) interaction measures: these are used to assess the search process and typically reflect usage activities (e.g., number of queries, number of clicked items); (iii) usability measures: the degree to which users may use a device to complete a task in a particular context of use (usually with parameters of efficacy, performance, and satisfaction); and (iv) contextual measures: these are typically aspects of the user (e.g., human characteristics such as intelligence, imagination, attitude, memory, and cognitive style) or the context of use.

The preceding material set out some basics of evaluation in information retrieval. It is important to emphasize that most efforts in evaluating GIR and GIS have followed similar approaches [188]. The following section provides us with various methodologies to evaluate geographic information retrieval systems and overall geographic information systems.

7.2.3 *Evaluation in Geographic Information Systems*

In [254], Mandl describes four levels of evaluation for GIR systems: (i) component level (i.e., evaluating individual parts of an IR system); (ii) system level (i.e., evaluating the outputs of a complete IR system); and (iii) user-system-interaction level (i.e., evaluating an entire GIR system including interfaces and visualizations in a controlled laboratory) and (iv) at the user performance level (i.e., assessing an operational system in use and its impact on daily work tasks). GIR evaluation has typically emphasized system-oriented approaches and measuring search performance, such as the development of standardized standards for systems and individual components [280],[57],[376],[149]. However, it is critical to analyze GIR systems and their constituent components for geo-referencing, presentation/visualization, and ranking using various methods, taking into account both system functionality and user's encounters and expectations for specific interfaces [254],[57]. Assessing spatial parameters, such as relevance, and assessing visualizations typically used in GIR interfaces, such as maps, are two factors that must be addressed in evaluations [188]. According to Robertson [308], relevance should be viewed as a continuous variable. Researchers have experimented with different levels or scales for assessing relevance; for example, the SPIRIT framework's text significance was assessed in terms of spatial and thematic relevance [75],[292].

Gorgiadou and Stoter [138] propose four evaluation orientations based on evaluation approaches [330], i.e., pure and mixed types. Control and exploratory evaluation are part of the pure form, whereas sense-making and learning evaluation are part of the mixed type [138]. Nedovic-Budic [272] applies the Delone and Mclean [93] information systems success model to assess GIS technology results, incorporating societal impact as a category or metric of success. They suggest that understanding the effects of GIS usage requires understanding what should be measured, whether the measurements should be direct or indirect and whether contextual or technical factors should precede.

Concluding, the above-mentioned work indicates that the measure used for information retrieval system evaluation can be used to evaluate the effectiveness and performance of geographic information systems, whether it is related to geo-referencing or determining the system's overall performance. In conclusion, for evaluating SORS we will be using the following evaluation methods:

- Using Recall, Precision and F1 measure for evaluating the overall performance of SORS,
- For finding the effectiveness of reasoning stages parameters of standard deviation and average are used,
- Concept of spatial relevance is used to evaluate the triplets generated by SORS.

7.3 EVALUATION CONDUCTED ON SORS

This section demonstrates the experimental setups and evaluation approaches used to assess our SORS system's performance. System-oriented evaluation approaches have ng been used to assess our reasoning framework's performance at the component level. For the information extraction phase using an exhaustive algorithm, the relevance of output has been determined based on the number of spatial triplets extracted by our approach, which provide a piece of implicit or extra information compared to the existing parsing systems.

The second set of experiments is designed to measure the system's overall performance by identifying the number of entities correctly geo-referenced on the OSM map after the contextualization stage using precision and recall measures. The third experiment assesses how well our system comprehends the context or gist of provided inputs by determining if the prototype can identify the target object. The fourth experiment evaluates the average number of correct OSM interpretations generated after applying the reasoning steps.

The fourth set of experiments uses comparison and compares our prototype with other existing named entity recognition systems and provides an overview of how the reasoning layer affects and improves identifying named entities and unnamed entities.

The fifth experiment is conducted to find out the overall performance of the information extraction phase using our exhaustive algorithm and to give a clear idea about that how well does a piece of implicit or extra information is handled and is extracted by our system but not by the existing parsing systems[2]. The spatial relevance of output is based on the entities of spatial triplets extracted by our approach.

For assessing the role of reasoning in interpreting place descriptions, we implemented our approach as a research prototype.

7.3.1 *Experimental setup and data*

To the best of our knowledge, there exists no dataset for evaluating georeferencing of spatial text which addresses complex place descriptions involving unnamed entities. Therefore, we collected a corpus of sentences that contain place descriptions from travel blogs and English Wikipedia by scanning the summary part from articles. Additionally, we have added sentences discussed in the literature on related approaches. Our evaluation dataset comprises 105 sentences and multi-sentence text fragments. We manually determined ground truth for all spatial entities in the dataset (about 430

2 NLTK and spacy

entities) by identifying OSM objects that correspond to the spatial entities occurring in the place descriptions. Our evaluation aims to assess the contribution of spatio-temporal reasoning and the effectiveness of our general approach, it does not aim to *measure* performance. Any performance measure for complex tasks critically depends on how the evaluation dataset is composed: complexity of sentences contained, distribution between named and unnamed entities, potential bias in geographic scope, the type of text it can adequately represent, etc. We therefore only compare performance measures between different experimental conditions, or between different kinds of entities.

Our evaluation is based on running our implementation on all sentences of the dataset. For each word corresponding to a spatial entity we check whether our system has identified the correct OSM entity. Correctly referenced entities are counted as true positives, missed entities as false negatives. For all nouns that do no have a spatial interpretation we check that that our system did not match them to any OSM entity. Those unmatched nouns are counted as true negatives, others false positives. We then compute precision, recall and the combined F_1 score.

7.3.1.1 *Overall performance*

In a first analysis we investigate the overall ability of our approach to geo-reference places correctly. These numbers are broken down per type of noun. We expect significant differences between simple queries for an entity by its name and complex queries for an unnamed entity, e.g., a river as in *river flowing through Prague*. We compare the results obtained by running the reasoning-based pipeline against a variant in which the reasoning steps are skipped except for querying, i.e., querying directly follows information extraction (cp. Figure 7). This condition is interesting as it implements an approach to toponym resolution in which one would greedily search for OSM entities, thus mainly relying on named entities. Such approach is an fully automated version of the approach using place graphs [371, 64], which starts by toponym resolution for named entities and then continues by searching for entities related to entities already geo-referenced in a strictly forward-chaining manner. Data for this analysis is presented in Table 25.

Table 25: Results obtained by our approach. Column Δ_{F_1} gives the improvement of performing reasoning steps.

	without reasoning			with reasoning			
condition	precision	recall	F_1	precision	recall	F_1	Δ_{F_1}
named entities	0.98	0.97	0.97	0.98	0.97	0.97	0.00
unnamed entities	0.4	0.37	0.38	0.70	0.52	0.6	0.22
inferred named entities	–	–	–	0.90	0.73	0.81	0.81

7.3.1.2 *Primary entities*

The second analysis singles out performance with respect to the gist of a place information. *Primary entities* are said represent the target place mentioned in a place description. For example, the primary place of the sentence *There is a famous bridge in Prague* is Charles Bridge. In order to interpret the primary place information correctly, *bridge* needs to be geo-referenced to

Charles Bridge's OSM entity. Identifying primary entities can be highly ambiguous particularly for unnamed entities (as in case of Prague, there exist multiple bridges) and may require a high degree of text understanding, but this measure addresses the ultimate goal of understanding place descriptions. For every sentence in our data set we have marked those words that refer to the primary entity, the data-set has been released and is present online on GitHub. Evaluation is performed as in the first analysis, but only primary entities are considered. Results obtained are presented in Table 26.

Table 26: Results for geo-referencing primary entities.

primary entity	Correctly Identified			
type	total number	precision	recall	F_1
Primary named entities				
named entities	18	1	1	1
newly named entities	49	1	0.94	0.96
Primary unnamed entities				
unnamed entities	27	0.51	0.84	0.64
Primary unknown entities				
unknown entities	6	0	0	0

7.3.1.3 *Effectiveness of reasoning stages*

Our third analysis addresses a differentiation of noun types along the processing pipeline. We count the number of nouns, differentiating between named entities, unnamed entities, and unidentified entities. Measures are taken for the contextualisation and inference stage separately to study the effect of the two methods. Unnamed entities may be interpreted as ontological types or co-reference, unidentified nouns may be interpreted as named entities. A decrease of unnamed and unidentified entities indicates an effective interpretation (correctness of the results as addressed by the first evaluation). We also measure the average number of queries that get enabled due to the interpretation. Data obtained for this evaluation is presented in Table 27.

7.3.1.4 *Evaluation based on Number of OSM Queries Enabled*

Our third analysis addresses a differentiation of noun types along the processing pipeline. We count the number of nouns, differentiating between named entities, unnamed entities, and unidentified entities. Measures are taken for the contextualisation and inference stage separately to study the effect of the two methods. Unnamed entities may be interpreted as ontological types or co-reference, unidentified nouns may be interpreted as named entities. A decrease of unnamed and unidentified entities indicates an effective interpretation (correctness of the results as addressed by the first evaluation). We also measure the average number of queries that get enabled due to the interpretation. Data obtained for this evaluation is presented in Table 27.

Table 27: Amount of entity types per sentence across the steps of the processing pipeline.

(a) from information extraction

entity type	avg. and std. dev.	
nouns	6.16	±2.33
named entities	1.26	±1.27
unnamed	1.92	±1.06
unidentified	2.97	±1.87

(b) exploitable by reasoning

entity type	avg. and std. dev.	
unnamed	59.3%	±35.9
unidentified	47.0%	±40.0

(c) after contextualisation

entity type	avg. and std. dev.	
unnamed	1.08	±0.96
unidentified	2.13	±1.69
contextual queries	1.66	±1.61

(d) after inference

entity type	avg. and std. dev.	
unnamed	0.89	±0.91
unidentified	1.98	±1.62
contextual queries	0.39	±0.69

Table 29: Results of entities successfully and correctly converted to OSM tags before and after reasoning.

	without reasoning		with reasoning	
condition	input	matched to OSM	input	matched to OSM
tags for named				
named Entities	96	96	96	96
tags for unnamed				
unnamed Entities	161	74	56	56
tags for unidentified				
unidentified Entities	176	0	63	3
tags for inferred named				
inferred named entities	–	–	218	210
total nouns converted				
total nouns	433	152	433	365

7.3.2 *Comparison to geo-taggers*

This experimental setting aims to compare the performance of our system against that of available approaches. Since to the best of our knowledge there is no ready-to-run system that can handle unnamed entities, we restrict this experiment to named entities. Moreover, available systems only perform geo-tagging, i.e., they identify toponyms in text, but they do not perform toponym resolution or just in special cases. The comparison is nevertheless interesting as it sheds light from a new angle on the effectiveness of reasoning in understanding place descriptions: As our approach also uses a part-of-speech tagger that performs named entity recognition (Stanford Core NLP), we can determine to which degree the proposed reasoning pipeline can improve performance as compared to a pure named entity recognition system. We consider the following systems:

- Standford Core NLP `https://stanfordnlp.github.io/CoreNLP/`
- spaCy `https://spacy.io/`

- DBpedia Spotlight https://www.dbpedia-spotlight.org/
- OpenCalais https://permid.org/onecalaisViewer
- our system SORS

In our evaluation dataset we marked all named entities to obtain a ground truth. For the evaluation we say that a named entity has been correctly identified by a tagger, if it is labeled as name of an unspecified or spatial entity. For example, if a tagger labels *Europe* as the name of a person, we regard this as a wrong interpretation. For SORS, we run our system as usual and then compare the set of named entities with the ground truth for the respective sentence. The set of named entities, i.e., the set of words for which a has-name constraint has been generated $\{W| \text{ has-name}(W, N) \in \Gamma\}$ comprises nouns tagged as named entities and unidentified nouns has been inferred. We do not consider whether named entities are geo-referenced to the correct OSM entity or not. For all systems we count the number of named entities in the dataset that have been correctly identified, as well as the number of words which are wrongly labeled as named entities. The results we obtain are presented in Table 30.

Table 30: Comparison of our system SORS with named entity recognition systems

system	entities recognised	precision	recall	F_1
SORS	**293**	**1.00**	**0.93**	**0.96**
DBpedia Spotlight	197	0.80	0.62	0.70
OpenCalais	156	0.86	0.50	0.63
spaCy	153	0.80	0.49	0.61
Standford Core NLP	134	0.77	0.42	0.51
ground truth	314			

7.3.3 *Spatial Relevance of Extracted Spatial Triplets*

This last experiment evaluates the output of the approach presented in chapter 4 for triplet extraction. Our approach relies on the concept of using a set of logical statements as an intermediate representation that over-generalizes knowledge conveyed in a sentence. The experiment is conducted to assess how much this over-generalization adds to the number of spatial triplets extracted using existing dependency parsers[3], along with measuring the spatial relevance of the extracted facts.

The spatial relevance concept allows determining whether a spatial triplet is relevant or non-relevant for a given information need or query. Thus, we are using the concept of binary relevance scale. The binary relevance scheme is based on the type of entities present in the triplet and consists of two scores:

- Score 1: The triplet is relevant because it includes either named/inferred entities or unnamed entities.
- Score 2: The triplet is not relevant if it does not contain named, inferred, or unnamed entities

3 Standford NLTK or Spacy

A comparison is carried out between the triplets generated by dependency parsers and our system, and a set of common triplets is identified (i.e., identified by all SORS, NLTK, and Spacy). The remaining triplets are assigned a relevance score based on binary scale definitions, and if the spatial triplets fall into the score one category, then we count the number of relevant facts for each sentence; otherwise, we ignore the fact. This process helps us in pruning the extra triplets and judges how practical the presented approach is in dealing with implicit knowledge. The results can be seen in the table 31 below.

Table 31: Results showing the additional spatially relevant triplets extracted by SORS with comparison to NLTK and Spacy.

Number of Additional facts	1	2	3	4	5	6	7	8	9	10	11
Number of Sentences	11	19	22	12	8	11	4	2	2	1	1
Percentage	10.4%	18%	21%	11.4%	7.6%	10.4%	3.8%	1.9%	1.9%	0.95%	0.95%

7.3.4 *Discussion of Results*

Geo-referencing places from unconstrained natural language is a tough task comprising several unsolved sub-tasks, so it comes at no surprise to see overall performance as presented in Table 25 to be unsatisfactory for practical applications. Only for geo-referencing named entities, precision and recall scores reach mostly 0.9, i.e., most of the objects get identified correctly. This is remarkable since toponym resolution is challenged by ambiguous interpretations using the complete OSM database (the so-called planet file), in particular names of German cities got re-used by immigrants in northern America. Moreover, we found a mismatch caused by English translations of originally German named entities which limits recognition success. As can be seen, reasoning does not affect performance when querying for named entities (first row in Table 25), which is to be expected. Only in one case a query for a city within a country, reasoning made required geographic context explicit, which was not already revealed at the information extraction stage. Analysis of the failures reveals that further contextual information may be required in some cases, which is not provided by the input sentence. In the second analysis we re-evaluated the results by focusing on the primary entity to see how well SORS can interpret the gist of a sentence. Comparing Table 25 to Table 26 we see that performance is comparable for primary named entities as for the set of all named entities. Also for sentences in which the primary entity is an unnamed entity, the performance is nearly identical (0.60 for all unnamed entities versus 0.64 for primary unnamed entities). This shows that SORS can interpret both kinds of spatial concepts in a sentence equally well. We note that Table 27 (b) indicates a substantial standard deviation of about $\pm$ 35–40%. This is simply due to the fact that sentences in the data set contain a varying amount of non-spatial words (unidentified) and unnamed spatial words for which no interpretation could be found. The bottom row of Table 26 lists primary entities which are not catered by our approach as they are concepts like *car* in *My car is parked near the Erba park close to the university in Bamberg*, for which no geo-referencing by means of identifying OSM entities can be performed.

Concerning the ability to resolve unnamed or unidentified entities, reasoning helps the results to be improved, see lower part of Table 25. From the first analysis alone it can be concluded that contextualization and spatio-ontological reasoning are effective means in place recognition. We particularly point to the fact that for queries like *post office near a train station in* X, where X designates a city, the approach is often able to identify the desired post office and train station within the city correctly. This is a clear advancement over many existing approaches. Effectiveness of the two reasoning procedures can be gauged from Table 27. Both of the main inference stages contribute to resolving uninterpreted parts of the input as the number of unnamed and unidentified entities decreases, and constrained queries are generated, as shown in Table 27 (c) and (d) in contrast to Table 27(a). We note that the number of words given in Table 27 (a) does not sum up to the total number of nouns in the dataset as sometimes distinct words are combined to a single entity, for example, in *townhall building*. <In the case of unnamed entities, i.e., nouns representing geographic concepts represented in our ontology, the average number of 1.92 entities per sentence obtained by information extraction decreases to 1.08 by contextualization and further to 0.89 by inference. Indicating that about 48% of the unnamed entities occurring in the input text get interpreted as either co-reference or ontological specification of a named entity. In total, 59.3% of all unnamed entities are exploited by reasoning in the sense that the ontological information provided by the word was used to form a query. Roughly 10% of the unnamed entities extracted initially refer to an entity not further specified in the text that our approach can identify, e.g., *post office near train station*. Since the application of reasoning resulted in strictly better performance concerning precision and recall, as shown in Table 25, one can regard the values in Table 27 to show actual improvements made.

We observe that the number of unidentified entities is considerably large (2.97 out of 6.16 words per sentence on average). Most of the unidentified entities in our corpus have indeed spatial entities, which means that named entity recognition as well as recognition of unnamed entities can further be improved.

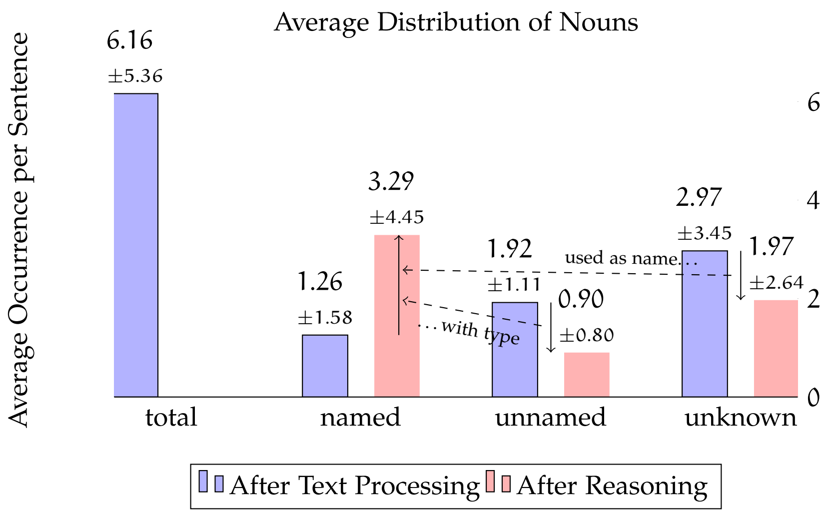

Figure 18: Nouns in dataset that represent spatial entities, classified into named entities recognised, unnamed spatial concepts, and nouns unknown to the lexicon. Output from text processing is compared to output from the proposed reasoning procedure, arrows indicate effect of the proposed reasoning steps to resolve nouns.

In several cases where part-of-speech-tagging is unable to identify a named entity, sufficient contextual information is generated to allow a successful query using the unidentified entity as a name (see Table 25, row inferred entities) In Figure 18 we present how the amount of unidentified and unnamed entities decreases as they get hypothesised to be a named entity which is then confirmed by querying the OSM database.

In the fourth analysis we compare SORS against existing systems for named entity recognition (NER). Existing NER systems perform well, yet there still exist issues with ambiguous names. For example, *Europe* may be classified as name of a women or as toponym. All pure NER systems compared rely on machine learning to grasp the context of a sentence and to choose the correct label for each word. As Table 30 shows, there is still much headroom for all NER systems. The performance of SORS is significantly better, scoring 1.00 in precision. While the numbers came at a surprise to us at first sight, there is a simple explanation. Precision score 1.0 means that only named entities which refer to a spatial entity are classified as such. After information extraction, SORS is confronted with names identified by Standford Core NLP (precision 0.77). During contextualisation, SORS aggressively tries to interpret all unidentified entities as names which have to be confirmed by a match in the OSM database. Therefore, precision score 1.0 comes naturally for our approach. Recall that in this analysis correctness of toponym resolution is not considered as it is not performed by the other systems. The recall score of 0.73 for SORS means that the aggressive approach of trying to interpret unidentified entities as names also produces a considerable amount of false positives. Scores for SORS are better than for NER systems which do not perform any reasoning that improves plausibility (as shown in previous analysis) and are still far from perfect. They easily miss the structure of the sentence and, for example, cannot find *Northumberland Street* but can only identify *Northumberland* as a county in North East England, tagging *Street* as a separate noun. Consider the example in Table 32, where the location for spatial entities like *Sherlock Holmes' apartment* and *The Sherlock Holmes pub* is described. As shown in the table, existing NER systems are unable to identify both entities as a location; instead, they identify it as a detective's name, which is correct but is does not fit the sentence. This evaluation should however been taken with a grain of salt: our evaluation dataset consists of natural language place descriptions, which naturally involve few non-spatial words. If SORS would be applied to arbitrary sentences that include some mentioning of place, say from a novel, SORS would still aggressively try to interpret nouns as named spatial entities, likely producing many false positives.

Table 32: Table comparing which entities are recognised by named entity recognition systems

Another version of Sherlock Holmes' apartment is at The Sherlock Holmes pub in a street called Northumberland near Charing Cross station.

Spatial Entities	SORS (our system)	Dbpedia Spotlight	OpenCalais	spaCy	Standford
Sherlock Holmes' apartment	✓	✗	✗	✗	✗
Sherlock Holmes pub	✓	✗	✓	✗	✗
Northumberland street	✓	✗	✗	✗	✗
Charing Cross station	✓	✓	✓	✓	✓

In the fifth experiment, we analyze the output of explicit spatio-ontological reasoning and how well it deals with the implicit information present in the given input sentence by comparing SORS results with existing dependency parser *Standford NLTK and Spacy*. Out of 105, we came across ten sentences with no common spatial triplets, while two sentences did not have any additional facts. The details for the remaining 93 sentences can be seen in table 31, where we provide the number of sentences by the number of additional facts identified in them. It can be seen from the table that most sentences have three or two other spatially relevant facts generated, increasing the likelihood of geo-referencing for a specific entity. For example, in *"Hannover on the River Leine, is the capital and largest city of the German state"*, initially, even if the system is unable to identify Hannover on River Leine and identifies just Hannover, the fact *of(River leine, German State)* generated by presented reasoning approach makes sure to provide a match for the river as well, along with Hannover. Overall, 88% of the sentences from our data set has spatial triplets besides the ones generated by parsers, which itself is an indication implicit knowledge is manipulated and extracted during the information extraction process.

As spatio-ontological reasoning provides us with a more compact set of facts, we also came across situations where the spatial triplets are spatially relevant but are incorrect. Consider the following sentence, *"Birmingham Repertory is a Britain longest-established producing theatre, Other producing theatres include the Blue Orange and the Birmingham Rep."*, Birmingham Repertory is identified as a theater, which is correct. However, it is also associated with Blue Orange and Birmingham Rep, and both are names of the different theaters. The extracted triplets are incorrect and make no sense, so they are rejected during the query phase.
Exhaustive search does not address the core issue of language understanding; instead, it is based on the premise that the large set of statements will make the process of reaching the intended interpretation flexible and manageable by finding a match to a regional database.

7.4 SUMMARY AND CONCLUSION

This chapter identifies and understands the idea of evaluation in geographic information systems and how the evaluation methodologies and success measures from IR and GIR fields can assess the geographic information systems and analyze our research prototype's performance.

The experiments are designed to assess the impact of SORS on the component as well as system level. The system's overall performance is measured by using precision-recall and F1 scores, criteria being the entities are successfully mapped on the open street map (OSM). The analysis is performed at the component level to determine how accurately the primary entities are mapped and how many desirable queries are produced. A comparison with existing systems has also been carried out to analyze the impact of reasoning methods and see if our system provides solutions to existing research gaps or performs in the same manner as existing approaches.
All these experimental evaluations of SORS allow us to understand that using the presented reasoning methods as an additional layer in the information retrieval phase of a geographic information system is quite effective. The reasoning method generates an extensive set of declarative facts, which resolves indirect references, eases querying, and incorporates contextual in-

formation in the form of type or unnamed entities present in human descriptions, which helps solve the named entity disambiguation process at an initial level. All this contributes to the geo-referencing process, which is improved as the query comprises the name and the type of spatial entity, making the search process more convenient.
However, all these evaluations are from a system-oriented perspective, and yet we need to evaluate our system from the user's perspective. A research prototype for this purpose is under consideration where the output can be evaluated and analyzed according to the user's point of view.

8 DISCUSSION AND CONCLUSION

As a human spatial concept widely used in communicative and cognitive tasks, **place** is not directly representable by the geographic information systems (GISs) and spatial databases developed so far. Whereas the traditional approach of spatial representation for GIS and spatial databases is usually focused on precise, crisp, and metric geometries and coordinates, places are often spatially uncertain, influenced by human experience, and context-dependent. The fundamental disparity prevents the possibility of information systems to capture and depict the way people think and reason about space that can be used to smooth out and simplify human-computer interactions. Despite the recognized significance of place-based research, however, little progress has been made in geographic information science (GIScience), and no commonly accepted place-based computational data model has been developed. This thesis aims to contribute in linking human and technical concepts related to place.

Place descriptions convey how people are mentally perceiving and communicating about places and information related to them.Natural language place descriptions implies that complementary place-based approaches to human information representation are possible. Since the web contains a myriad of place descriptions in different types, it presents a unique opportunity for VGI to collect rich place description datasets. Towards developing a place-based GIS, this thesis aims to contribute to several research aspects. *Regarding place descriptions as a new type of unstructured data source that existing GISs and spatial databases do not explicitly assist, this research discussed incorporating context, spatial information extraction and modeling, reasoning, place geo-referencing, and querying with information derived from place descriptions using reasoning techniques.*

This research makes a two-fold contribution. As a scientific contribution, the first fold includes approaches to solve these five tasks: a) a computational model of place capable of incorporating context from human-generated place descriptions, b) extracting both implicit and explicit relational information from text, c) reasoning with the extracted spatial entities as human spatial knowledge of places, d) locating places from descriptions through the contextualization of spatial type entities and disambiguation, and e) generating constraint specific queries, identifying appropriate OSM tag and type entity and querying OSM. As contribution to basic research, this thesis analyzed limitations of existing approaches to geo-referencing and it demonstrated that spatio-ontological reasoning plays a central role in understanding place descriptions.

The second fold is a prototype that integrates and incorporates functions that relate to the five activities. With such a system using an essential layer of reasoning to provide all the spatial entities present in the place descriptions, it would be possible to explore further how this system could improve current services involving human-computer interaction. The ultimate aim of such a system is the intelligent machine [396] that understands and produces human spatial language, enabling smooth communication between man and machine.

In what follows, we will review the main results, contrast them to the research questions raised in Chapter 1 and respond to the hypotheses sketch in Chapter 2. The shortcomings of this study are described in Section 8.3, and future research directions are discussed.

8.1 DISCUSSION OF MAJOR RESULTS

This section will summarize the approaches developed for the primary research tasks and discuss the results obtained.

8.1.1 *Incorporating Context in Automated Interpretation of Text: A Computational Model*

Natural language uses place names and descriptions in written and oral forms as references for essential locations. The first step in creating a spatially-aware system is to determine what needs to be presented concerning the place and which parts of the human context are related to people's understanding of the place, including the name and type of the place, semantic description, and even the semantic relationship between physical objects and places. Before interpreting the description of the place, we must decide the definition of the place, leading to the following questions: What is a place and how to define it to grasp human intuition, leading to a fully defined computational data model.
Chapter 3 discusses how contextual information affects the meaning of place descriptions and describes a computational model that represents the vague spatial knowledge and context found in the description of places. It uses three elements to define a place: a place is identified by a name (which can be identified by a unique name or a circumscription), its spatial extent, and through its conceptualization, that is, through the process of connecting specific concepts to spatial areas. Conceptualization uses a bottom-up approach to define *context variables*, which affect the most reasonable interpretation of the description of the place and are derived from the available information. Context variables are divided into three categories: the *environmental factors*, the *human factors* who generates a place description, and the *linguist factors of place description* . This classification adapts a characterization of context when performing map-based tasks by [126] to a text-based task. Our implicit linguistic context in the description of the place corresponds to its map context, but we separate objective environmental factors from all cognitive factors.

The last section of Chapter 3 goes through the model's specifics and focuses on how the presented model considers the spatial dimensions established in the literature—the conceptualization of place fully embodies the three aspects of Agnews [98]. Considering the relationship with other places usually requires identifying clear objects and understanding their types and is implied by the environment and human descriptions of places. Conceptualization using three defined context variables helps to conclude various functions, activities, events, and language. These conclusions help us understand place as a cognitive concept. In addition, the model records spatial, semantic, and context-related information on places that can be linked to the five central concepts of spatial information proposed by Kuhn (2012): location, field, object, network, and event, as discussed in Section 3.2.3

The following Chapters 4 and 5 explain how contextual variables (especially the type of location) focus on improving the collection and analysis of geographic information by considering place as a context.

8.1.2 *Spatial Information Extraction*

The main aim of our research is to extract spatial entities and to reason about them and represent them. The automatic interpretation of text remains a challenge, mainly due to language analysis and ambiguity between names and human conceptualization. Traditional learning or rule-based methods are insufficient to cope with the volume and dimensions of unstructured data. In addition, no open access parser correctly resolves in-text references. A misidentified reference can easily hinder the correct interpretation of a sentence.Additionally, there is no definite solution for methods that cannot successfully parse the text.
Therefore, Chapter 4 explores the extent to which reasoning about spatial and ontological attributes can help extract spatial information and overcome problems in natural language parsing. Our purpose is to identify informative triples by considering context dependencies—a reasoning method rather than the machine learning method selected for extracting the triplets, instead of depending on parsers.

Spatial-ontological reasoning integrates the concepts of geographic scope resolution with the best possible level of spatial entity granularity in order to resolve place descriptions by adjusting the semantics of relationships and queries to focus on results that correspond to scope identified by type and location of the geographic entities that appear in the same text. Our method uses logic programming techniques to declaratively express dependencies as constraint satisfaction problems (CSP) corresponding to spatial objects. Our approach uses a series of logical statements as intermediate representations, over-generalizing the information expressed in the sentences. Therefore, it is an exhaustive search composed of *Generation* and *pruning* stages(used to eliminate unreasonable explanations), relying on the exact source of information.It searches for the conjunctive terms which can be best satisfied by searching a suitable instantiation from a geographical database for un-referenced nouns that agrees with the relations. The approach also tests relations between named entities and/or referenced nouns.

The exhaustive search (for the concept of mental model theory) contains all correct interpretations by construction, and some statements that do not follow from the input text, and some are incorrect. In comparison with parsers that provide us with limited information regarding relations from complex language constructs, exhaustive search provides us with explicit and implicit information as a densely connected set of facts.This overgeneralization helps resolve reference problems or geographical ambiguities by relying on the input sentence, while others rely on external knowledge and heuristics based on additional knowledge of world locations such as population and land area.

By making spatial ontological reasoning explicit in interpretation, we can take context-related dependencies into account. The presented approach conforms to the assumption that the most extensive set of statements matched to a geographic database corresponds to the intended interpretation. Compared with parsers using Wikipedia English sentences that describe geographic features, we see that reasoning eliminates most misunderstandings, while natural language parsing leads to misleading commitments that cannot be detected in a later post-processing step.

8.1.3 *Spatio-Ontological Reasoning using Unnamed Entities and Geo-referencing*

Place descriptions are regarded as descriptive or subjective references for describing geographic locations. To this end, the input text is processed using spatial and ontology (type)relationships to form a series of declarative relation statements. The third goal of this work is geo-referencing. The central hypothesis of our research is that spatial and ontological reasoning allows us to overcome the limitations of previous geo-referencing methods, which include the limited reliability of name recognition using NER systems and the limited use of context for disambiguation and reliance on natural language parsing.

Chapter 5 interprets the information extracted from language employing reasoning and investigates how implicit information present in the description can be exploited. This chapter provides an automated approach to geo-reference entities in natural language place descriptions, capable of interpreting unnamed entities and a method for pruning off implausible interpretations using context-sensitive reasoning by improving the performance of text understanding components, including named entity recognition. The presented approach uses the context information in the form of type nouns to enrich the interpretation, improve the geo-referencing process semantically, and use abductive and deductive reasoning methods to fill in the missing information to narrow down the interpretations.

SORS consists of four main modules, in which the reasoning is expanded into three different modules. The entity and relation extraction module extracts information in the form of named, unnamed, and unidentified entities and their relationships. Spatial and ontological relations are used to resolve ambiguous named and unnamed entities. The contextualization module involves associating statements with contextual information and performing abductive reasoning by linking (named) entities with unnamed entities based on the same ontological type. In the inference phase, the derived information is propagated deductively. The reasoning is employed in an exhaustive search manner, aiming to identify an interpretation of the input sentence that maximizes the number of noun phrases matched to the OSM database. In addition, a simple likelihood-based heuristic is proposed that can be used to guide the search to the most reasonable explanation, in case most nouns are geo-referenced. The heuristic is based on the observation that related words often appear in closer proximity than unrelated words in descriptions.

The proposed geo-referencing approach enables the mapping of locations and establishes links to coordinates in GIS and geo-databases. The contribution of this phase is that it allows us to deal with named and unnamed entities and can geo-reference unnamed entities and exploit them for improving geo-referencing of named entities by connecting names to ontological types. Moreover, through contextualization, the ambiguity in the resolution of toponyms for recognizing named entities can be resolved by evaluating the plausibility of an interpretation in the context of the entire sentence. The effectiveness of the suggested method depends on the availability of a mapping of unnamed entities to geographic types that can be used when querying the OSM database. While this is true for purely spatial descriptions, it is a clear limitation when it comes to text that is not purely spatial. Overall, this thesis shows that the reasoning provides a solution to the existing limitations of research.

8.1.4 *Querying*

After extracting information with spatial ontological reasoning, this final task focuses on composing queries and scheduling them to map the entities.As a query engine we use OSCAR[1] , which is a new geographic data search engine based on free OSM data.
Chapter 6 first explains how to formulate natural language queries in the OSM query format (from the set of extracted entities). Using a contextualized lexicon that associates spatial entities with concept classes and geometric interpretations, the final entities inferred in Chapter 6 is converted to the required OSM Query format. In addition to the technical aspects of composing queries, this chapter also provides a query strategy for ordering queries and determining whether to send the query or whether the number of expected results is too large to manage. These types of queries are called *constrained queries*. For formulating them, additional classifications based on entity types are provided in the form of conjunctions.
In addition to query formulation and strategy, this research investigates how well semantic similarity of words allows the correct key-value pairs to be identified when tagging cannot retrieve the intended target object without precisely hand-crafted background information to apply the technique to any word encountered. This subtask aims to generate a set of adequate tags to query entities described in natural language words. A similarity-based ranking function is proposed, allowing us to start with a reasonable tag and retry with a next-best option if a query is not successful.
The last part of Chapter 6 mainly deals with interpreting entities that are not given names and can only be described. A context-sensitive method for determining place types is proposed. This method is based on the semantic similarity of words and generates semantically replaceable terms in the form of noun types. Identifying type nouns to replace non-spatial nouns in the description corresponds to types in geo-databases improves geo-referencing spatial entities. The result is a catalog of related types which help in answering OSM queries.

Overall, the query strategy based on entity type help avoids costly queries by serializing queries and focusing on reasonable candidate locations. Besides, the approach presented can also be applied to questions and declarative statements to achieve the desired result. We hope, therefore, that the algorithmic techniques described in this chapter will help empirical researchers who wish to test specific spatial models or steps in the interpretation process. WordNet's similarity measure applied directly to geographic

1 https://github.com/dbahrdt/oscar

terms aids in posing successful queries and results in appropriate tag selection without using any outside knowledge. We used these resources to implement automatic reasoning techniques to interpret constraints spatially.

8.2 SUMMARY OF CONTRIBUTIONS AND EVALUATION AGAINST HYPOTHESES

This thesis has addressed the five identified significant challenges in defining place using context, modeling, and reasoning with type nouns as context variables, geo-referencing with unnamed entities, and querying human place knowledge extracted from NL place descriptions. The major contributions and outcomes are summarized as follows:

- A computational model has been presented which can capture vague and context-sensitive information about places. This model allows tackling the interpretations as a constrained optimizing problem using discrete and continuous variables.
- A reasoning framework for spatial information extraction is presented to progress towards the long-standing goal of solving spatial language parsing. The framework uses exhaustive search for generating all plausible parse options and then keep the valid ones by pruning. The framework allows the consideration of contextual dependencies.
- A new method to automatically interpret natural language place descriptions by using spatio-ontological reasoning techniques, enabling us to explicate implied information in the form of Unnamed entities embedded in place descriptions. Reasoning allows the generation of new information across three components: contextualization (incorporating type information), ontological reasoning (is-a relation), and inference to generate map-able data.
- A query strategy is presented to schedule queries to the underlying geographic database and supplement the set of relational statements accordingly. A strategy becomes necessary to avoid queries that lead to more matches in the database that can be handled.
- We investigated means to identify the entity type used for a place in OSM by ranking and identifying semantic similarity tags.

The remainder of this section compares the contributions and assesses them against the hypotheses described in Chapter 2.

1. **H1: A more precise understanding of place in geographic information systems can be achieved by defining it in terms of context.**
   For the first hypothesis, a computational model has been presented in Chapter 3, which defines place concerning context and how it shapes the meaning of place. The place has been defined to have a name, spatial extent, and conceptualization, which depends on how humans describe the place in natural language. The conceptualization module identifies contextual variables using a bottom-up approach, which indicates that context can be inferred in three different categories environment, human, and place description. The advantages of the model has been demonstrated in Chapters 4, 5 and in 6. Specifically, the categorization presented in this model can be used to extract more context-sensitive information than parsers [63],which will improve the

process of automation of place-based systems. Thus defining place with respect to context makes it easy to identify the factors of the human context ascribed to people's understanding of places like its name, type, and semantics of entities and relations, proving the first hypothesis that context-sensitive information can lead to an efficient geographic information system.

2. **H2: By incorporating context, reasoning approaches aid in automating the process of understanding place descriptions (using the presented computational model). It could improve the information extraction from unstructured text and can overcome the limitations of the existing parser.**

   For the second hypothesis, a detailed state of the art is presented in Chapter 2, where despite extensive work in understanding spatial knowledge, reasoning, and representation, a more in-depth consideration of context is required. Existing techniques, including machine learning and rule-based methods or QSR techniques, have significant challenges impeding the automation and analysis of natural language place descriptions. Thus, this thesis investigates the contribution of reasoning over state-of-the-art geo-referncing techniques. Chapters 3, 4, 5 and 2 proves the superiority of reasoning and indicates that considering an additional layer of reasoning based on context helps improve spatial information retrieval, modeling and reasoning, geo-referencing, and querying, thus proving the second hypothesis.

3. **H3: Using the finest level of granularity possible for relation extraction aids in resolving indirect or implied information in place descriptions and can improve the parser's co-reference resolution problem.**

   For the third hypothesis, a comprehensive reasoning approach is presented in Chapter 4, in which the description of the places is solved with the best possible granularity. Logical statements are used as an intermediate representation for expressing the over-generalized information present in the sentence. The initial evaluation and comparison with existing parsers provided in Chapter 4 and Chapter 7 indicate that the algorithm can exploit the implicit information along with the explicit one. It provides us with the densely connected set of facts which, by exploiting implicit information, solve a parser's co-reference resolution task and resolve indirect references (as shown in the discussion section 4.3 of Chapter 4), making spatio-ontological reasoning effective.

   Another advantage of the presented approach is that it helps solve geo-ambiguity issues (example in Section 4.3) by generating the most extensive set of statements that can be matched to the geographic databases. Besides, this method uses contextual dependency given the context variable (place description) introduced in the chapter 3. These variables allow generating relational statements based on spatial relationships rather than verbs, resulting in better performance than the parsers irrespective of the sentence structure.

   Making spatio-ontological reasoning an explicit step generates the maximum number of queries based on the entity level, which helps find the intended interpretation in the corresponding geographic database, proving the third hypothesis.

4. **H4: Explicit use of the type information in the description helps to interpret and geo-reference spatial entities in different languages and unnamed entities. It could also improve the existing NER systems.**

   Regarding the fourth hypothesis, Chapter 5 provides a reasoning approach using contextual information presented in the form of type nouns to enrich the interpretation and semantically improve the geo-referencing process. The process uses abductive and deductive reasoning to fill in missing information and uses a likelihood heuristic to narrow the scope of the interpretations. Therefore, this method overcomes limitations of the traditional named entity recognition systems proposed in the literature, which only considers the named places. The advantages of the approach are that it allows named entity disambiguation at an information retrieval phase and allows named and unnamed entities to be geo-referenced, and it only yields limited numbers of false positives.
   The evaluation presented in the discussion section 5.6 of chapter 5 and in chapter 7 provides an insight into how the contextualization step that generates hypothetical is-a relations and unification improves the number of geo-referenced entities with a reasonable level of precision and recall for unnamed entities as well.Thus, proving the fourth and second hypotheses.

5. **H5: Constrained specific queries and the ranking algorithm based on semantic similarity could improve the quality of geo-referencing.**

   For the last hypothesis, place queries have been identified and categorized, and a contextualized querying approach is introduced in Chapter 6. Besides, Chapter 6 provides us with two different algorithms based on semantic similarity to identify likely entity types used. The first algorithm helps to generate a set of OSM key-value pairs using a ranking-based algorithm where the designated type of entity is not appropriate. It improves querying by providing the next best options for queries instead of halting the process and generating no solution upon unsuccessful queries.

6. **H6: Using semantic similarity measure and clustering technique can help in finding what to query for by generating similar terms in the form of unnamed entities.**
   The algorithm 7 generates the semantically replaceable terms for the non-spatial nouns in the forms of type nouns using clustering and pruning techniques. Once the type nouns are generated, the reasoning approach presented in this thesis can be applied to the sentences, and entities can be geo-referenced for both named and unnamed objects. Both these algorithms improve the quality of querying and geo-referencing, thus proving the last hypothesis.

8.3 LIMITATIONS AND FUTURE WORK

This section is divided into two parts. First we highlight the limitations of this study and discuss how future studies may overcome these limitations. Secondly we discuss several promising future research topics based on the approaches and results of this thesis.

8.3.1 *Spatial Relation Modeling*

Interpreting and modeling spatial relationships from NL expressions is one of the core challenges in spatial language understanding [63]. Even though spatial relation modeling is beyond the scope of this thesis, it is critical for an intelligent spatially aware system to have a defined set of spatial relations semantics. In this research work, only containment and proximity relations are implemented in a straightforward manner, which works well for topological relations like on, at, in, by, and so on, as suggested by Varsadani [367],that with specified granularity level, these relations can be replaced with each other. However, dealing with directional relations requires an even deeper consideration of context along with relation semantics.

Thus, future work in this area requires work on the models that can be employed better to model the semantics of spatial relationships from NL expressions. Existing reasoning models like fuzzy or defeasible logic-based models can also be used or based on the contextual factors and their definitions; the relations can be depicted after identifying and mapping spatial entities.

8.3.2 *Evaluation Data Set*

For place geo-referencing, the major limitation in determining the utility of the proposed approach is that there is no data-set for evaluating georeferencing of spatial text that addresses complex place descriptions involving unnamed entities. Existing evaluation data sets usually focus on question-answering systems, and sub tasks such as NER. We collected a corpus of sentences that contain place descriptions from travel blogs and English Wikipedia by scanning the summary part from articles. Additionally, we have added sentences discussed in the literature on related approaches.

A proper community wide data-set can help improve the evaluation and thus contributes in finding further improvements in geo-referencing. Therefore, by collecting and investigating means to annotate the test data from Wikipedia, the next step involves publishing and promoting a data-set along with the ground truth that will facilitate to evaluate our system and a comparison between existing approaches and our system.

8.3.3 *Querying Spatial Databases*

The query system researched in this thesis is using OSCAR as a query database for spatial mapping entities. However, representing the rich repertoire of spatial relations that co-exists and interrelates any entities, GeoSPARQL has been standardized as query language, which features data types to represent geographical entities and represent spatial relations.

The GeoSPARQL is attractive as a general approach to query geographical data since it allows semantic information to be used for filtering, like using key-value tags in OSM. Thus in the future, shifting from OSCAR/OSM to GeoSPARQL might be a beneficial step for improving the quality of query answering for place references. Considering current research work, the OSM query format can be converted to GeoSPAQRL format by linking object entities to a formal ontology.

8.4 APPLICATIONS OF A SPATIO-ONTOLOGICAL REASONING SYSTEM

Regarding a SORS the next step is to explore how it can be used for facilitating place-related research and services. This subsection discusses different research directions.

1. This work contributes to a better understanding of NL place descriptions as input for GIS in general. Because spatial information is ubiquitous, many applications will benefit from these capabilities because they can learn about the location information from people, web searches, automated emergency response, or robots instructed using spatial language. Exploitation of NL is currently not available or restricted in these applications. In those applications problems with unstructured input, vernacular references, and context are typical.

2. There exists rich sources of historical data. While this thesis has considered OSM as a contemporary map, it could be interesting to adapt the approach to historical data, possibly helping to discover new findings

3. Current methods for emergency content mapping are usually based on text analysis on social media and usually rely on geo-tagged user information and keyword analysis. When the geotag information is not available, our system can interpret the spatial NL expression, which is helpful for this task. Therefore, this work provides methods and ideas for future research in the field of emergency content mapping.

4. Location-based search engines treat place-based NL queries as crisp points or polygons, ignoring the type of spatial entities and relations. The queries and geo-referencing techniques included in this study provide information on how to understand these queries properly.

With the promising results achieved on how reasoning can be applied to NL understanding and how it benefits spatial text understanding, we are looking forward to new applications that may emerge.

BIBLIOGRAPHY

[1] "A Comparison of String Similarity Measures for Toponym Matching." In: *Proceedings of The First ACM SIGSPATIAL International Workshop on Computational Models of Place*. COMP '13. Orlando FL, USA: Association for Computing Machinery, 2013, 54–61. ISBN: 9781450325356. DOI: 10.1145/2534848.2534850. URL: https://doi.org/10.1145/2534848.2534850.

[2] Asma Ben Abacha and Pierre Zweigenbaum. "Medical entity recognition: a comparaison of semantic and statistical methods." In: *Proceedings of BioNLP 2011 workshop*. 2011, pp. 56–64.

[3] Zahraa S. Abdallah, Mark Carman, and Gholamreza Haffari. "Multi-domain evaluation framework for named entity recognition tools." In: *Computer Speech & Language* 43 (2017), pp. 34–55. ISSN: 0885-2308. DOI: https://doi.org/10.1016/j.csl.2016.10.003. URL: https://www.sciencedirect.com/science/article/pii/S0885230815300504.

[4] Benjamin Adams and Krzysztof Janowicz. "On the Geo-Indicativeness of Non-Georeferenced Text." In: *Proceedings of the Sixth International Conference on Weblogs and Social Media, Dublin, Ireland, June 4-7, 2012*. Ed. by John G. Breslin, Nicole B. Ellison, James G. Shanahan, and Zeynep Tufekci. The AAAI Press, 2012. URL: http://www.aaai.org/ocs/index.php/ICWSM/ICWSM12/paper/view/4629.

[5] Benjamin Adams and Krzysztof Janowicz. "Thematic signatures for cleansing and enriching place-related linked data." In: *Int. J. Geogr. Inf. Sci.* 29.4 (2015), pp. 556–579. DOI: 10.1080/13658816.2014.989855. URL: https://doi.org/10.1080/13658816.2014.989855.

[6] Benjamin Adams and Grant McKenzie. "Inferring thematic places from spatially referenced natural language descriptions." In: *Crowdsourcing Geographic Knowledge, Volunteered Geographic Information (VGI) in Theory and Practice*. Ed. by Daniel Sui, Sarah Elwood, and Michael Goodchild. Springer, 2013, pp. 201–221.

[7] Marco D. Adelfio and Hanan Samet. "Structured toponym resolution using combined hierarchical place categories." In: *Proceedings of the 7th Workshop on Geographic Information Retrieval, GIR 2013, 5th November, 2013, Orlando, Florida, USA*. Ed. by Christopher B. Jones and Ross Purves. ACM, 2013, pp. 49–56. DOI: 10.1145/2533888.2533931. URL: https://doi.org/10.1145/2533888.2533931.

[8] Kiran Adnan and Rehan Akbar. "Limitations of information extraction methods and techniques for heterogeneous unstructured big data." In: *International Journal of Engineering Business Management* 11 (2019), p. 1847979019890771. DOI: 10.1177/1847979019890771. eprint: https://doi.org/10.1177/1847979019890771. URL: https://doi.org/10.1177/1847979019890771.

[9] Pragya Agarwal. "Operationalising 'Sense of Place' as a Cognitive Operator for Semantics in Place-Based Ontologies." In: *Spatial Information Theory*. Ed. by Anthony G. Cohn and David M. Mark. Berlin, Heidelberg: Springer Berlin Heidelberg, 2005, pp. 96–114. ISBN: 978-3-540-32020-3.

[10] Eugene Agichtein and Luis Gravano. "*Snowball*: extracting relations from large plain-text collections." In: *Proceedings of the Fifth ACM Conference on Digital Libraries, June 2-7, 2000, San Antonio, TX, USA*. ACM, 2000, pp. 85–94. DOI: 10.1145/336597.336644. URL: https://doi.org/10.1145/336597.336644.

[11] John Agnew. "Space and place." In: *The Sage Handbook of Geographical Knowledge* (Jan. 2011), pp. 316–330. DOI: 10.4135/9781446201091.n24.

[12] Azzah Al-Maskari, Mark Sanderson, and Paul Clough. "The Relationship between IR Effectiveness Measures and User Satisfaction." In: *Proceedings of the 30th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. SIGIR '07. Amsterdam, The Netherlands: Association for Computing Machinery, 2007, 773–774. ISBN: 9781595935977. DOI: 10.1145/1277741.1277902. URL: https://doi.org/10.1145/1277741.1277902.

[13] Ahmed N. Alazzawi, Alia I. Abdelmoty, and Christopher B. Jones. "What can I do there? Towards the automatic discovery of place-related services and activities." In: *International Journal of Geographical Information Science* 26.2 (2012), pp. 345–364. DOI: 10.1080/13658816.2011.595954.

[14] Christopher Alexander. *The nature of order: an essay on the art of building and the nature of the universe. Book 2, Book 2,* English. Berkeley, Calif.: Center for Environmental Structure, 2002. ISBN: 0972652922 9780972652926 0972652906 9780972652902.

[15] James Allan, Bruce Croft, Alistair Moffat, and Mark Sanderson. "Frontiers, Challenges, and Opportunities for Information Retrieval: Report from SWIRL 2012 the Second Strategic Workshop on Information Retrieval in Lorne." In: *SIGIR Forum* 46.1 (May 2012), 2–32. ISSN: 0163-5840. DOI: 10.1145/2215676.2215678. URL: https://doi.org/10.1145/2215676.2215678.

[16] Christopher Allen, Thomas Hervey, Sara Lafia, Daniel Phillips, Behzad Vahedi, and Werner Kuhn. "Exploring the Notion of Spatial Lenses." In: vol. 9927. Sept. 2016, pp. 259–274. ISBN: 978-3-319-45737-6. DOI: 10.1007/978-3-319-45738-3_17.

[17] Saad Aloteibi and Mark Sanderson. "Analyzing geographic query reformulation: An exploratory study." In: *Journal of the Association for Information Science and Technology* 65 (2014).

[18] Einat Amitay, Nadav Har'El, Ron Sivan, and Aya Soffer. "Web-a-where: geotagging web content." In: *SIGIR 2004: Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Sheffield, UK, July 25-29, 2004*. Ed. by Mark Sanderson, Kalervo Järvelin, James Allan, and Peter Bruza. ACM, 2004, pp. 273–280. DOI: 10.1145/1008992.1009040. URL: https://doi.org/10.1145/1008992.1009040.

[19] Ivo Anastácio, Bruno Martins, and Pável Calado. "Classifying Documents According to Locational Relevance." In: *Progress in Artificial Intelligence, 14th Portuguese Conference on Artificial Intelligence, EPIA 2009, Aveiro, Portugal, October 12-15, 2009. Proceedings*. Ed. by Luís Seabra Lopes, Nuno Lau, Pedro Mariano, and Luis Mateus Rocha. Vol. 5816. Lecture Notes in Computer Science. Springer, 2009, pp. 598–609. DOI: 10.1007/978-3-642-04686-5_49. URL: https://doi.org/10.1007/978-3-642-04686-5_49.

[20] Geoffrey Andogah, Gosse Bouma, and John Nerbonne. "Every document has a geographical scope." In: *Data Knowl. Eng.* 81-82 (2012), pp. 1–20. DOI: 10.1016/j.datak.2012.07.002. URL: https://doi.org/10.1016/j.datak.2012.07.002.

[21] Geoffrey Andogah, Gosse Bouma, John Nerbonne, and Elwin Koster. "Placename Ambiguity Resolution." In: *Proceedings of Language Resources and Evaluation Conference (LREC)*. 2008, pp. 4–10.

[22] Ballatore Andrea, Hegarty Mary, Kuhn Werner, and Parsons Ed. "Spatial Search, Final Report." In: *UC Santa Barbara: Center for Spatial Studies*. 2015. URL: https://escholarship.org/uc/item/33t8h2nw.

[23] Maria Varsadani. Arbaz khan and Stephan Winter. "Extracting Spatial Information From Place Descriptions." In: *Proceedings of The First ACM SIGSPATIAL International Workshop on Computational Models of Place*. COMP '13. New York, NY, USA: Association for Computing Machinery, 2013, 62–69. ISBN: 9781450325356. DOI: 10.1145/2534848.2534857.

[24] Linda H. Armitage and Peter G.B. Enser. "Analysis of user need in image archives." In: *Journal of Information Science* 23.4 (1997), pp. 287–299. DOI: 10.1177/016555159702300403. eprint: https://doi.org/10.1177/016555159702300403. URL: https://doi.org/10.1177/016555159702300403.

[25] John Atkinson and Veronica Bull. "A multi-strategy approach to biological named entity recognition." In: *Expert Systems with Applications* 39.17 (2012), pp. 12968–12974.

[26] Isabelle Augenstein, Diana Maynard, and Fabio Ciravegna. "Distantly supervised Web relation extraction for knowledge base population." In: *Semantic Web* 7.4 (2016), pp. 335–349. DOI: 10.3233/SW-150180. URL: https://doi.org/10.3233/SW-150180.

[27] Amittai Axelrod. "On building a high performance gazetteer database." In: *Proceedings of the HLT-NAACL 2003 Workshop on Analysis of Geographic References*. 2003, pp. 63–68. URL: https://aclanthology.org/W03-0110.

[28] Nguyen Bach and Sameer Badaskar. "A Review of Relation Extraction." In: 2007.

[29] Miriam Baglioni, Emiliano Giovannetti, Maria V. Masserotti, Chiara Renso, and Laura Spinsanti. "Ontology-supported Querying of Geographical Databases." In: *Transactions in GIS* 12 (2008).

[30] Daniel Bahrdt and Stefan Funke. "OSCAR: OpenStreetMap Planet at Your Fingertips via OSm Cell ARrangements." In: *Proceedings, Part I, of the 16th International Conference on Web Information Systems Engineering — WISE 2015 - Volume 9418*. Berlin, Heidelberg: Springer-Verlag, 2015, 153–168. ISBN: 9783319261898. DOI: 10.1007/978-3-319-26190-4_11. URL: https://doi.org/10.1007/978-3-319-26190-4_11.

[31] Daniel Bahrdt, Rick Gelhausen, Stefan Funke, and Sabine Storandt. "Searching OSM Planet with Context-Aware Spatial Relations." In: *Proceedings of SIGSPATIAL'17, November 7–10, 2017, Los Angeles Area, CA, USA*. 2017.

[32] Kanwalpreet Singh Bajwa and Amardeep Kaur. "Hybrid approach for named entity recognition." In: *International Journal of Computer Applications* 118.1 (2015).

[33] Andrea Ballatore, Michela Bertolotto, and David C. Wilson. "Geographic Knowledge Extraction and Semantic Similarity in OpenStreetMap." In: *Knowl. Inf. Syst.* 37.1 (Oct. 2013), pp. 61–81. ISSN: 0219-1377. DOI: 10.1007/s10115-012-0571-0.

[34] Andrea Ballatore, Werner Kuhn, Mary Hegarty, and Ed Parsons. "Special issue introduction: Spatial approaches to information search." In: *Spatial Cogn. Comput.* 16.4 (2016), pp. 245–254. DOI: 10.1080/13875868.2016.1243693. URL: https://doi.org/10.1080/13875868.2016.1243693.

[35] Michele Banko, Michael J. Cafarella, Stephen Soderland, Matthew Broadhead, and Oren Etzioni. "Open Information Extraction from the Web." In: *IJCAI 2007, Proceedings of the 20th International Joint Conference on Artificial Intelligence, Hyderabad, India, January 6-12, 2007*. Ed. by Manuela M. Veloso. 2007, pp. 2670–2676. URL: http://ijcai.org/Proceedings/07/Papers/429.pdf.

[36] John A. Bateman. "Language and Space: a two-level semantic approach based on principles of ontological engineering." In: *Int. J. Speech Technol.* 13.1 (2010), pp. 29–48. DOI: 10.1007/s10772-010-9069-x. URL: https://doi.org/10.1007/s10772-010-9069-x.

[37] John A. Bateman, Joana Hois, Robert J. Ross, and Thora Tenbrink. "A linguistic ontology of space for natural language processing." In: *Artif. Intell.* 174.14 (2010), pp. 1027–1071. DOI: 10.1016/j.artint.2010.05.008. URL: https://doi.org/10.1016/j.artint.2010.05.008.

[38] John A. Bateman, Joana Hois, Robert Ross, and Thora Tenbrink. "A linguistic ontology of space for natural language processing." In: *Artificial Intelligence* 174 (2010), pp. 1027–1071.

[39] John Bateman, Thora Tenbrink, and Scott Farrar. "The Role of Conceptual and Linguistic Ontologies in Interpreting Spatial Discourse." In: *Discourse Processes* 44.3 (2007), pp. 175–212. eprint: https://doi.org/10.1080/01638530701600912. URL: https://doi.org/10.1080/01638530701600912.

[40] Mary Bazire and Patrick Brézillon. "Understanding context before using it." In: *In the proceedinds of Context 2005*. 2005, pp. 29–40.

[41] Lamia Belouaer, David Brosset, and Christophe Claramunt. "From verbal route descriptions to sketch maps in natural environments." In: Oct. 2016, pp. 1–10. DOI: 10.1145/2996913.2997003.

[42] Brandon Bennett. "Spatial Vagueness." In: vol. 256. Jan. 2011. ISBN: 978-3-642-14754-8. DOI: 10.1007/978-3-642-14755-5_2.

[43] Brandon Bennett and Pragya Agarwal. "Semantic Categories Underlying the Meaning of 'Place'." In: *Spatial Information Theory*. Ed. by Stephan Winter, Matt Duckham, Lars Kulik, and Ben Kuipers. Berlin, Heidelberg: Springer Berlin Heidelberg, 2007, pp. 78–95. ISBN: 978-3-540-74788-8.

[44] Pavel Berkhin. "A Survey of Clustering Data Mining Techniques." In: *Grouping Multidimensional Data - Recent Advances in Clustering*. Ed. by Jacob Kogan, Charles K. Nicholas, and Marc Teboulle. Springer, 2006, pp. 25–71. DOI: 10.1007/3-540-28349-8_2. URL: https://doi.org/10.1007/3-540-28349-8_2.

[45] Jules Berman. "Providing Structure to Unstructured Data." In: Dec. 2013, pp. 1–14. ISBN: 9780124045767. DOI: 10.1016/B978-0-12-404576-7.00001-0.

[46] Tim Berners-Lee, James Hendler, and Ora Lassila. "The Semantic Web." In: *Scientific American* 284.5 (May 2001), pp. 34–43. URL: http://www.sciam.com/article.cfm?articleID=00048144-10D2-1C70-84A9809EC588EF21.

[47] Claudio Bettini, Oliver Brdiczka, Karen Henricksen, Jadwiga Indulska, Daniela Nicklas, Anand Ranganathan, and Daniele Riboni. "A survey of context modelling and reasoning techniques." In: *Pervasive and Mobile Computing* 6.2 (2010), pp. 161–180.

[48] Steven Bird, Ewan Klein, and Edward Loper. *Natural Language Processing with Python: Analyzing Text with the Natural Language Toolkit*. Beijing: O'Reilly, 2009. ISBN: 978-0-596-51649-9. DOI: http://my.safaribooksonline.com/9780596516499. URL: http://www.nltk.org/book.

[49] Thomas Bittner and John G. Stell. "Rough Sets in Approximate Spatial Reasoning." In: *Rough Sets and Current Trends in Computing, Second International Conference, RSCTC 2000 Banff, Canada, October 16-19, 2000, Revised Papers*. Ed. by Wojciech Ziarko and Y. Y. Yao. Vol. 2005. Lecture Notes in Computer Science. Springer, 2000, pp. 445–453. DOI: 10.1007/3-540-45554-X_55. URL: https://doi.org/10.1007/3-540-45554-X_55.

[50] Thomas Blaschke and Helena Merschdorf. "Geographic information science as a multidisciplinary and multiparadigmatic field." In: *Cartography and Geographic Information Science* 41.3 (2014), pp. 196–213. DOI: 10.1080/15230406.2014.905755. eprint: https://doi.org/10.1080/15230406.2014.905755. URL: https://doi.org/10.1080/15230406.2014.905755.

[51] Thomas Blaschke, Helena Merschdorf, Pablo Cabrera-Barona, Song Gao, Emmanuel Papadakis, and Anna Kovacs-Györi. "Place versus Space: From Points, Lines and Polygons in GIS to Place-Based Representations Reflecting Language and Culture." In: *ISPRS International Journal of Geo-Information* 7.11 (2018). ISSN: 2220-9964. DOI: 10.3390/ijgi7110452. URL: https://www.mdpi.com/2220-9964/7/11/452.

[52] Pia Borlund. "User-Centred Evaluation of Information Retrieval Systems." In: Oct. 2009, pp. 21 –37. ISBN: 9780470033647. DOI: 10.1002/9780470033647.ch2.

[53] Bert Bredeweg and Peter Struss. "Current Topics in Qualitative Reasoning." In: 24.4 (2004), pp. 13–16.

[54] Sergey Brin. "Extracting Patterns and Relations from the World Wide Web." In: *The World Wide Web and Databases, International Workshop WebDB'98, Valencia, Spain, March 27-28, 1998, Selected Papers*. Ed. by Paolo Atzeni, Alberto O. Mendelzon, and Giansalvatore Mecca. Vol. 1590. Lecture Notes in Computer Science. Springer, 1998, pp. 172–183. DOI: 10.1007/10704656_11. URL: https://doi.org/10.1007/10704656_11.

[55] Boyan Brodaric. "Geo-Pragmatics for the Geospatial Semantic Web." In: *Transactions in GIS* 11.3 (2007), pp. 453–477. DOI: https://doi.org/10.1111/j.1467-9671.2007.01055.x. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1467-9671.2007.01055.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-9671.2007.01055.x.

[56] Shawn Brunsting, Hans De Sterck, Remco Dolman, and Teun van Sprundel. "GeoTextTagger: High-Precision Location Tagging of Textual Documents using a Natural Language Processing Approach." In: *CoRR* abs/1601.05893 (2016). arXiv: 1601.05893.

[57] Bénédicte Bucher, Paul Clough, Hideo Joho, Ross Purves, and Awase Syed. "Geographic IR systems: requirements and evaluation." In: (Jan. 2005).

[58] Razvan C. Bunescu and Raymond J. Mooney. "Subsequence Kernels for Relation Extraction." In: *Advances in Neural Information Processing Systems 18 [Neural Information Processing Systems, NIPS 2005, December 5-8, 2005, Vancouver, British Columbia, Canada]*. 2005, pp. 171–178. URL: https://proceedings.neurips.cc/paper/2005/hash/2e0bff759d057e28460eaa5b2cb118e5-Abstract.html.

[59] Davide Buscaldi. "Approaches to disambiguating toponyms." In: *ACM SIGSPATIAL Special* 3.2 (2011), pp. 16–19. DOI: 10.1145/2047296.2047300. URL: https://doi.org/10.1145/2047296.2047300.

[60] Davide Buscaldi and Paulo Rosso. "A Conceptual Density-Based Approach for the Disambiguation of Toponyms." In: *Int. J. Geogr. Inf. Sci.* 22.3 (2008), 301–313. ISSN: 1365-8816. DOI: 10.1080/13658810701626251. URL: https://doi.org/10.1080/13658810701626251.

[61] Guoray Cai. "Contextualization of Geospatial Database Semantics for Human—GIS Interaction." In: *Geoinformatica* 11.2 (2007), 217–237. ISSN: 1384-6175. DOI: 10.1007/s10707-006-0001-0. URL: https://doi.org/10.1007/s10707-006-0001-0.

[62] Guoray Cai, Hongmei Wang, Alan M. MacEachren, and Sven Fuhrmann. "Natural conversational interfaces to geospatial databases." English (US). In: *Transactions in GIS* 9.2 (Mar. 2005), pp. 199–221. ISSN: 1361-1682. DOI: 10.1111/j.1467-9671.2005.00213.x.

[63] Hao Chen. "A place graph database as a qualitative human place knowledge base." PhD thesis. University of Melbourne, Parkville, Victoria, Australia, 2019.

[64] Hao Chen, Maria Vasardani, Stephan Winter, and Martin Tomko. "A Graph Database Model for Knowledge Extracted from Place Descriptions." In: *ISPRS International Journal of Geo-Information* 7.6 (2018). ISSN: 2220-9964. DOI: 10.3390/ijgi7060221.

[65] Juan Chen, Anthony G. Cohn, Dayou Liu, Shengsheng Wang, Jihong Ouyang, and Qiangyuan Yu. "A survey of qualitative spatial representations." In: 30.1 (2013), pp. 106–136.

[66] Zhiyuan Cheng, James Caverlee, and Kyumin Lee. "You are where you tweet: a content-based approach to geo-locating twitter users." In: *Proceedings of the 19th ACM Conference on Information and Knowledge Management, CIKM 2010, Toronto, Ontario, Canada, October 26-30, 2010*. Ed. by Jimmy Huang, Nick Koudas, Gareth J. F. Jones, Xindong Wu, Kevyn Collins-Thompson, and Aijun An. ACM, 2010, pp. 759–

768. DOI: 10.1145/1871437.1871535. URL: https://doi.org/10.1145/1871437.1871535.

[67] Zhiyuan Cheng, James Caverlee, and Kyumin Lee. "A content-driven framework for geolocating microblog users." In: *ACM Trans. Intell. Syst. Technol.* 4.1 (2013), 2:1–2:27. DOI: 10.1145/2414425.2414427. URL: https://doi.org/10.1145/2414425.2414427.

[68] Chien-Lung Chou and Chia-Hui Chang. "Named Entity Extraction via Automatic Labeling and Tri-training: Comparison of Selection Methods." In: *Information Retrieval Technology*. Ed. by Azizah Jaafar, Nazlena Mohamad Ali, Shahrul Azman Mohd Noah, Alan F. Smeaton, Peter Bruza, Zainab Abu Bakar, Nursuriati Jamil, and Tengku Mohd Tengku Sembok. Cham: Springer International Publishing, 2014, pp. 244–255. ISBN: 978-3-319-12844-3.

[69] Gobinda Chowdhury. "Natural language processing." In: *ARIST* 37 (Jan. 2005), pp. 51–89. DOI: 10.1002/aris.1440370103.

[70] Eliseo Clementini and Paolino Di Felice. "An Algebraic Model for Spatial Objects with Indeterminate Boundaries." In: Jan. 1996, pp. 155–169. ISBN: 9781003062660. DOI: 10.1201/9781003062660-15.

[71] Eliseo Clementini, Paolino Di Felice, and Daniel Hernández. "Qualitative representation of positional information." In: *Artificial Intelligence* 95.2 (1997), pp. 317–356. ISSN: 0004-3702. DOI: https://doi.org/10.1016/S0004-3702(97)00046-5. URL: https://www.sciencedirect.com/science/article/pii/S0004370297000465.

[72] Cyril W. Cleverdon. "The Significance of the Cranfield Tests on Index Languages." In: *Proceedings of the 14th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. SIGIR '91. Chicago, Illinois, USA: Association for Computing Machinery, 1991, 3–12. ISBN: 0897914481. DOI: 10.1145/122860.122861. URL: https://doi.org/10.1145/122860.122861.

[73] Cyril W Cleverdon, Jack Mills, and E Michael Keen. "Factors determining the performance of indexing systems,(Volume 1: Design)." In: *Cranfield: College of Aeronautics* 28 (1966).

[74] P. Clough and M. Sanderson. "Evaluating the performance of information retrieval systems using test collections." In: *Information Research* 18.2 (2013). © Year The Author(s). This is an Open Access article distributed under the terms of the Creative Commons Attribution Licence (http://creativecommons.org/licenses/by/3.0), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. URL: https://eprints.whiterose.ac.uk/78886/.

[75] Paul D. Clough, Hideo Joho, and Ross Purves. "Judging the Spatial Relevance of Documents for GIR." In: *Advances in Information Retrieval, 28th European Conference on IR Research, ECIR 2006, London, UK, April 10-12, 2006, Proceedings*. Ed. by Mounia Lalmas, Andy MacFarlane, Stefan M. Rüger, Anastasios Tombros, Theodora Tsikrika, and Alexei Yavlinsky. Vol. 3936. Lecture Notes in Computer Science. Springer, 2006, pp. 548–552. DOI: 10.1007/11735106_62. URL: https://doi.org/10.1007/11735106_62.

[76] Richard Coates. "Properhood." In: *Language* 82 (June 2006), pp. 356–382. DOI: 10.1353/lan.2006.0084.

[77] Mihai Codescu, Gregor Horsinka, Oliver Kutz, Till Mossakowski, and Rafaela Rau. "OSMonto – An Ontology of OpenStreetMap Tags." In: *State of the map Europe (SOTM-EU) 2011*. 2011.

[78] Mihai Codescu, Daniel Couto Vale, Oliver Kutz, and Till Mossakowski. "Ontology-based Route Planning for OpenStreetMap." In: *Proceedings of the Terra Cognita Workshop on Foundations, Technologies and Applications of the Geospatial Web, Boston, USA, November 12, 2012*. Ed. by Dave Kolas, Matthew Perry, Rolf Grütter, and Manolis Koubarakis. Vol. 901. CEUR Workshop Proceedings. CEUR-WS.org, 2012, pp. 62–73.

[79] Anthony G. Cohn and Nicholas Mark Gotts. "The egg-yolk representation of regions with in determinate boundaries." In: *GIS specialist meeting on geographical objects with undetermined boundaries*. 1996, pp. 171–187.

[80] Anthony G. Cohn and Shyamanta M. Hazarika. "Qualitative Spatial Representation and Reasoning: An Overview." In: *Fundam. Informaticae* 46.1-2 (2001), pp. 1–29. URL: http://content.iospress.com/articles/fundamenta-informaticae/fi46-1-2-02.

[81] Anthony G. Cohn, Sanjiang Li, Weiming Liu, and Jochen Renz. "Reasoning about Topological and Cardinal Direction Relations Between 2-Dimensional Spatial Objects." In: 51 (2014), pp. 493–532.

[82] Helen Couclelis. "Location, place, region, and space." In: (1992).

[83] Jim Cowie and Wendy Lehnert. "Information Extraction." In: *Commun. ACM* 39.1 (Jan. 1996), 80–91. ISSN: 0001-0782. DOI: 10.1145/234173.234209. URL: https://doi.org/10.1145/234173.234209.

[84] Tim Cresswell. *Place: A Short Introduction*. Jan. 2004. ISBN: 1405106727.

[85] Tim Cresswell. "Place." English. In: *International Encyclopedia of Human Geography*. Ed. by Nigel Thrift and Rob Kitchen. Vol. 8. Elsevier, 2009, pp. 169–177.

[86] Jennifer Cross. "What is Sense of Place?" In: Jan. 2001.

[87] Aron Culotta and Jeffrey Sorensen. "Dependency Tree Kernels for Relation Extraction." In: *Proceedings of the 42nd Annual Meeting of the Association for Computational Linguistics (ACL-04)*. Barcelona, Spain, July 2004, pp. 423–429. DOI: 10.3115/1218955.1219009. URL: https://aclanthology.org/P04-1054.

[88] Michael R. Curry. *Digital Places: Living with Geographic Information Technologies*. London: Taylor & Francis, 1998.

[89] Jennifer D'Souza and Vincent Ng. "UTD: Ensemble-Based Spatial Relation Extraction." In: Jan. 2015, pp. 862–869. DOI: 10.18653/v1/S15-2146.

[90] Soham Dan, Parisa KordJamshidi, Julia Bonn, Archna Bhatia, Zheng Cai, Martha Palmer, and Dan Roth. "From Spatial Relations to Spatial Configurations." In: *Proceedings of The 12th Language Resources and Evaluation Conference, LREC 2020, Marseille, France, May 11-16, 2020*. Ed. by Nicoletta Calzolari et al. European Language Resources Association, 2020, pp. 5855–5864. ISBN: 979-10-95546-34-4. URL: https://www.aclweb.org/anthology/2020.lrec-1.717/.

[91] Clare Davies, Ian Holt, Jenny Green, Jenny Harding, and Lucy Diamond. "User Needs and Implications for Modelling Vague Named Places." In: *Spatial Cognition & Computation* 9 (Aug. 2009), pp. 174–194. DOI: 10.1080/13875860903121830.

[92] Ernest Davis. *Representations of Commonsense Knowledge*. Morgan Kaufmann, 1990.

[93] William H DeLone and Ephraim R McLean. "Information systems success: The quest for the dependent variable." In: *Information systems research* 3.1 (1992), pp. 60–95.

[94] Rina Dechter. *Constraint processing*. Morgan Kaufmann, 2003.

[95] Curdin Derungs, Damien Palacio, and Ross Purves. "Resolving fine granularity toponyms: Evaluation of a disambiguation approach." In: Sept. 2012. DOI: 10.5167/uzh-67546.

[96] Curdin Derungs and Ross S. Purves. "From text to landscape: locating, identifying and mapping the use of landscape features in a Swiss Alpine corpus." In: *Int. J. Geogr. Inf. Sci.* 28.6 (2014), pp. 1272–1293. DOI: 10.1080/13658816.2013.772184. URL: https://doi.org/10.1080/13658816.2013.772184.

[97] Mariana Dias, João Boné, João C. Ferreira, Ricardo Ribeiro, and Rui Maia. "Named Entity Recognition for Sensitive Data Discovery in Portuguese." In: *Applied Sciences* 10.7 (2020). ISSN: 2076-3417. DOI: 10.3390/app10072303. URL: https://www.mdpi.com/2076-3417/10/7/2303.

[98] Katie Digan. "An Introduction to Space and Place." In: *Places of Memory: The Case of the House of the Wannsee Conference*. London: Palgrave Macmillan UK, 2015, pp. 12–18. ISBN: 978-1-137-45642-7. DOI: 10.1057/9781137456427_2. URL: https://doi.org/10.1057/9781137456427_2.

[99] B Dixon, V Uddameri, et al. "GIS analysis: fundamentals of spatial query." In: *GIS and geocomputation for water resource science and engineering* (2016), pp. 111–128.

[100] Mark Dunlop. "Reflections on Mira: Interactive evaluation in information retrieval." In: *JASIS* 51 (Dec. 2000), pp. 1269–1274. DOI: 10.1002/1097-4571(2000)9999:9999<::AID-ASI1042>3.0.CO;2-7.

[101] Frank Dylla, Till Mossakowski, Thomas Schneider, and Diedrich Wolter. "Algebraic Properties of Qualitative Spatio-temporal Calculi." In: *Spatial Information Theory, 11th International Conference, COSIT 2013 Scarborough, UK, September 2–6, 2013 Proceedings*. Vol. 8116. LNCS. Springer, 2013, pp. 516–536.

[102] Miguel P. Eckstein. "Visual search: A retrospective." In: *Journal of Vision* 11.5 (Dec. 2011), pp. 14–14. ISSN: 1534-7362. DOI: 10.1167/11.5.14. eprint: https://arvojournals.org/arvo/content_public/journal/jov/933487/jov-11-5-14.pdf. URL: https://doi.org/10.1167/11.5.14.

[103] Max J. Egenhofer. "Toward the Semantic Geospatial Web." In: *Proceedings of the 10th ACM International Symposium on Advances in Geographic Information Systems*. GIS '02. McLean, Virginia, USA: Association for Computing Machinery, 2002, 1–4. ISBN: 1581135912. DOI: 10.1145/585147.585148. URL: https://doi.org/10.1145/585147.585148.

[104] Max J. Egenhofer and David Mark. "Naive geography." In: *Spatial information theory: a theoretical basis for GIS*. Ed. by Andrew U. Frank and Werner Kuhn. 1995, pp. 1–15.

[105] Jacob Eisenstein, Brendan O'Connor, Noah A. Smith, and Eric P. Xing. "A Latent Variable Model for Geographic Lexical Variation." In: *Proceedings of the 2010 Conference on Empirical Methods in Natural Language Processing*. Cambridge, MA: Association for Computational Linguistics, Oct. 2010, pp. 1277–1287. URL: https://aclanthology.org/D10-1124.

[106] Safaa Eltyeb and Naomie Salim. "Chemical named entities recognition: a review on approaches and applications." In: *Journal of cheminformatics* 6.1 (2014), pp. 1–12.

[107] Sarah Elwood, Michael Goodchild, and Daniel Sui. "Prospects for VGI research and the emerging fourth paradigm." English (US). In: *Crowdsourcing Geographic Knowledge*. Vol. 9789400745872. Netherlands: Springer Netherlands, Mar. 2013, pp. 361–375. ISBN: 9400745869. DOI: 10.1007/978-94-007-4587-2_20.

[108] C. Ergler. "Place: an introduction by T. Cresswell, Wiley-Blackwell, Chichester, UK, 2nd edition, 2014. 232 pp. ISBN 978-0-470–65562-7." In: *New Zealand Geographer* 71 (2015), pp. 219–220.

[109] M. Erwig and M. Schneider. "Vague Regions." In: *Proceedings of the 5th International Symposium on Advances in Spatial Databases*. 1997, pp. 298–320.

[110] Oren Etzioni, Michael J. Cafarella, Doug Downey, Ana-Maria Popescu, Tal Shaked, Stephen Soderland, Daniel S. Weld, and Alexander Yates. "Unsupervised named-entity extraction from the Web: An experimental study." In: *Artif. Intell.* 165.1 (2005), pp. 91–134. DOI: 10.1016/j.artint.2005.03.001. URL: https://doi.org/10.1016/j.artint.2005.03.001.

[111] Oren Etzioni, Anthony Fader, Janara Christensen, Stephen Soderland, and Mausam. "Open Information Extraction: The Second Generation." In: *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*. Ed. by Toby Walsh. IJCAI/AAAI, 2011, pp. 3–10. DOI: 10.5591/978-1-57735-516-8/IJCAI11-012. URL: https://doi.org/10.5591/978-1-57735-516-8/IJCAI11-012.

[112] Anthony Fader, Stephen Soderland, and Oren Etzioni. "Identifying Relations for Open Information Extraction." In: *Proceedings of the 2011 Conference on Empirical Methods in Natural Language Processing, EMNLP 2011, 27-31 July 2011, John McIntyre Conference Centre, Edinburgh, UK, A meeting of SIGDAT, a Special Interest Group of the ACL*. ACL, 2011, pp. 1535–1545. URL: https://aclanthology.org/D11-1142/.

[113] B. Farbey, F. Land, and D. Targett. "Moving IS evaluation forward: learning themes and research issues." In: *The Journal of Strategic Information Systems* 8.2 (1999), pp. 189–207. URL: https://eprints.soton.ac.uk/36373/.

[114] Keith Feldman, Louis Faust, Xian Wu, Chao Huang, and Nitesh V. Chawla. "Beyond Volume: The Impact of Complex Healthcare Data on the Machine Learning Pipeline." In: *CoRR* abs/1706.01513 (2017). arXiv: 1706.01513. URL: http://arxiv.org/abs/1706.01513.

[115] Susan Feldman. *NLP Meets the Jabberwocky: Natural Language Processing in Information Retrieval*. Online magazine *Online* http://www.onlinemag.net. 1999. URL: http://www.onlinemag.net/OL1999/feldman5.html.

[116] Christiane Fellbaum. *WordNet: An Electronic Lexical Database*. Cambridge (MA), USA: MIT Press, 1998.

[117] Dieter Fensel. *Ontologies - a silver bullet for knowledge management and electronic commerce*. Springer, 2001. ISBN: 978-3-540-41602-9.

[118] Paolo Fogliaroni and Heidelinde Hobel. "Implementing naive geography via qualitative spatial relation queries." In: *Proceedings of the 18th AGILE international conference on geographic information science*. 2015.

[119] Frederico T. Fonseca. "Geospatial Semantic Web." In: *Encyclopedia of GIS*. Ed. by Shashi Shekhar and Hui Xiong. Springer, 2008, pp. 388–391. DOI: 10.1007/978-0-387-35973-1_513. URL: https://doi.org/10.1007/978-0-387-35973-1_513.

[120] Cidália Costa Fonte and Weldon A. Lodwick. "Areas of fuzzy geographical entities." In: *Int. J. Geogr. Inf. Sci.* 18.2 (2004), pp. 127–150. DOI: 10.1080/13658810310001620933. URL: https://doi.org/10.1080/13658810310001620933.

[121] Andrew U. Frank. "Qualitative spatial reasoning about distances and directions in geographic space." In: 3.4 (1992), pp. 343–371.

[122] Andrew U Frank and David M Mark. "LANGUAGE ISSUES FOR GEOGRAPHICAL INFORMATION SYSTEMS1." In: (1990).

[123] Joy A. Frechtling, Henry T. Frierson, Stafford Hood, Gerunda B. Hughes, and Conrad G. Katzenmeyer. "The 2002 User Friendly Handbook for Project Evaluation." In: 2002.

[124] Christian Freksa. "Qualitative spatial reasoning." In: *Cognitive and linguistic aspects of geographic space*. Springer, 1991, pp. 361–372.

[125] Christian Freksa. "Using Orientation Information for Qualitative Spatial Reasoning." In: *Theories and Methods of Spatio-Temporal Reasoning in Geographic Space, International Conference GIS - From Space to Territory: Theories and Methods of Spatio-Temporal Reasoning, Pisa, Italy, September 21-23, 1992, Proceedings*. Ed. by Andrew U. Frank, Irene Campari, and Ubaldo Formentini. Vol. 639. Lecture Notes in Computer Science. Springer, 1992, pp. 162–178. DOI: 10.1007/3-540-55966-3_10. URL: https://doi.org/10.1007/3-540-55966-3_10.

[126] Christian Freksa, Alexander Klippel, and Stephan Winter. "A Cognitive Perspective on Spatial Context." In: *Spatial Cognition: Specialization and Integration*. Ed. by Anthony G. Cohn, Christian Freksa, and Bernhard Nebel. Dagstuhl Seminar Proceedings 05491. Dagstuhl, Germany: Internationales Begegnungs- und Forschungszentrum für Informatik (IBFI), Schloss Dagstuhl, Germany, 2007.

[127] Christian Freksa, Jasper van de Ven, and Diedrich Wolter. "Formal representation of qualitative direction." In: *International Journal of Geographical Information Science* 32.12 (2018), pp. 2514–2534. DOI: 10.1080/13658816.2017.1420794. eprint: https://doi.org/10.1080/13658816.2017.1420794. URL: https://doi.org/10.1080/13658816.2017.1420794.

[128] Gaihua Fu, Christopher B. Jones, and Alia I. Abdelmoty. "Building a Geographical Ontology for Intelligent Spatial Search on the Web." In: *IASTED International Conference on Databases and Applications, part of the 23rd Multi-Conference on Applied Informatics, Innsbruck, Austria, February 14-16, 2005*. Ed. by M. H. Hamza. IASTED/ACTA Press, 2005, pp. 167–172.

[129] Gaihua Fu, Christopher B. Jones, and Alia I. Abdelmoty. "Ontology-Based Spatial Query Expansion in Information Retrieval." In: *On the Move to Meaningful Internet Systems 2005: CoopIS, DOA, and ODBASE, OTM Confederated International Conferences, CoopIS, DOA, and ODBASE 2005, Agia Napa, Cyprus, October 31 - November 4, 2005, Proceedings, Part II*. Ed. by Robert Meersman, Zahir Tari, Mohand-Said Hacid, John Mylopoulos, Barbara Pernici, Özalp Babaoglu, Hans-Arno Jacobsen, Joseph P. Loyall, Michael Kifer, and Stefano Spaccapietra. Vol. 3761. Lecture Notes in Computer Science. Springer, 2005, pp. 1466–1482. DOI: 10.1007/11575801_33. URL: https://doi.org/10.1007/11575801_33.

[130] Qingqing Gan, Josh Attenberg, Alexander Markowetz, and Torsten Suel. "Analysis of Geographic Queries in a Search Engine Log." In: *Proceedings of the First International Workshop on Location and the Web*. LOCWEB '08. Beijing, China: Association for Computing Machinery, 2008, 49–56. ISBN: 9781605581606. DOI: 10.1145/1367798.1367806. URL: https://doi.org/10.1145/1367798.1367806.

[131] Song Gao, Krzysztof Janowicz, and Helen Couclelis. "Extracting urban functional regions from points of interest and human activities on location-based social networks." In: *Trans. GIS* 21.3 (2017), pp. 446–467. DOI: 10.1111/tgis.12289. URL: https://doi.org/10.1111/tgis.12289.

[132] Song Gao, Krzysztof Janowicz, Daniel R. Montello, Yingjie Hu, Jiue-An Yang, Grant McKenzie, Yiting Ju, Li Gong, Benjamin Adams, and Bo Yan. "A data-synthesis-driven method for detecting and extracting vague cognitive regions." In: *Int. J. Geogr. Inf. Sci.* 31.6 (2017), pp. 1245–1271. DOI: 10.1080/13658816.2016.1273357. URL: https://doi.org/10.1080/13658816.2016.1273357.

[133] Eric Garbin and Inderjeet Mani. "Disambiguating Toponyms in News." In: *HLT/EMNLP 2005, Human Language Technology Conference and Conference on Empirical Methods in Natural Language Processing, Proceedings of the Conference, 6-8 October 2005, Vancouver, British Columbia, Canada*. The Association for Computational Linguistics, 2005, pp. 363–370. URL: https://aclanthology.org/H05-1046/.

[134] Peter Gardenfors. *The geometry of meaning: Semantics based on conceptual spaces*. MIT press, 2014.

[135] Simon Garrod and Kenny Coventry. "Saying, Seeing and Acting: The Psychological Semantics of Spatial Prepositions." In: *Saying Seeing and Acting: The Psychological Semantics of Spatial Prepositions* (July 2004). DOI: 10.4324/9780203641521.

[136] Johanna Geiß, Andreas Spitz, Jannik Strötgen, and Michael Gertz. "The Wikipedia location network: overcoming borders and oceans." In: *Proceedings of the 9th Workshop on Geographic Information Retrieval, GIR 2015, Paris, France, November 26-27, 2015*. Ed. by Ross S. Purves

and Christopher B. Jones. ACM, 2015, 2:1–2:3. DOI: 10.1145/2837689.2837694. URL: https://doi.org/10.1145/2837689.2837694.

[137] Judith Gelernter, Gautam Ganesh, Hamsini Krishnakumar, and Wei Zhang. "Automatic Gazetteer Enrichment with User-geocoded Data." In: *Proceedings of the Second ACM SIGSPATIAL International Workshop on Crowdsourced and Volunteered Geographic Information*. GEOCROWD '13. Orlando, Florida: ACM, 2013, pp. 87–94. ISBN: 978-1-4503-2528-8. DOI: 10.1145/2534732.2534736.

[138] Y. Georgiadou and J. Stoter. "Towards a social agenda for SDI evaluation (Unpublished), in CENUS." In: (2008).

[139] Alfonso Gerevini and Jochen Renz. "Combining Topological and Qualitative Size Constraints for Spatial Reasoning." In: *Proceedings of the 4th International Conference on Principles and Practice of Constraint Programming*. CP '98. Berlin, Heidelberg: Springer-Verlag, 1998, 220–234. ISBN: 3540652248.

[140] Alfonso Gerevini and Jochen Renz. "Combining topological and size information for spatial reasoning." In: 137.1-2 (2002), pp. 1–42.

[141] Stephen R Gillespie. "Measuring the benefits of GIS use: Two transportation case studies." In: *URISA Journal* 6.2 (1994), pp. 62–67.

[142] Yoav Goldberg and Michael Elhadad. "An Efficient Algorithm for Easy-First Non-Directional Dependency Parsing." In: *Human Language Technologies: Conference of the North American Chapter of the Association of Computational Linguistics, Proceedings, June 2-4, 2010, Los Angeles, California, USA*. The Association for Computational Linguistics, 2010, pp. 742–750. URL: https://aclanthology.org/N10-1115/.

[143] Reginald G. Golledge and Aron N. Spector. "Comprehending the Urban Environment: Theory and Practice." In: *Geographical Analysis* 10.4 (1978), pp. 403–426. DOI: 10.1111/j.1538-4632.1978.tb00667.x. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1538-4632.1978.tb00667.x.

[144] M. F. Goodchild and L. L. Hill. "Introduction to digital gazetteer research." In: *International Journal of Geographical Information Science* 22.10 (2008), pp. 1039–1044. DOI: 10.1080/13658810701850497. eprint: https://doi.org/10.1080/13658810701850497. URL: https://doi.org/10.1080/13658810701850497.

[145] Michael F. Goodchild. "Geographical information science." In: *International Journal of Geographical Information Systems* 6.1 (1992), pp. 31–45. DOI: 10.1080/02693799208901893. eprint: https://doi.org/10.1080/02693799208901893. URL: https://doi.org/10.1080/02693799208901893.

[146] Michael F. Goodchild. "Citizens as Voluntary Sensors: Spatial Data Infrastructure in the World of Web 2.0." In: *Int. J. Spatial Data Infrastructures Res.* 2 (2007), pp. 24–32. URL: https://ijsdir.sadl.kuleuven.be/index.php/ijsdir/article/view/28.

[147] Michael F. Goodchild. "Formalizing place in geographic information systems." In: *Communities, Neighborhoods, and Health: Expanding the Boundaries of Place*. Springer, 2011, pp. 21–34.

[148] Michael Goodchild and Linna Li. "Formalizing space and place." In: *CIST2011 - Fonder les sciences du territoire*. Fonder les sciences du territoire. Collège international des sciences du territoire (CIST). Paris, France, Nov. 2011, pp. 177–183. URL: https://hal.archives-ouvertes.fr/hal-01353206.

[149] Milan Gritta, Mohammad Taher Pilevar, Nut Limsopatham, and Nigel Collier. "What's missing in geographical parsing?" In: *Language Resources and Evaluation* 52 (June 2018). DOI: 10.1007/s10579-017-9385-8.

[150] Christian Grothe and Jochen Schaab. "Automated Footprint Generation from Geotags with Kernel Density Estimation and Support Vector Machines." In: *Spatial Cogn. Comput.* 9.3 (2009), pp. 195–211. DOI: 10.1080/13875860903118307. URL: https://doi.org/10.1080/13875860903118307.

[151] Qinghua Guo, Yu Liu, and John Wieczorek. "Georeferencing locality descriptions and computing associated uncertainty using a probabilistic approach." In: *Int. J. Geogr. Inf. Sci.* 22.10 (2008), pp. 1067–1090. DOI: 10.1080/13658810701851420. URL: https://doi.org/10.1080/13658810701851420.

[152] Ehsan Hamzei, Hao Chen, Maria Vasardani, Martin Tomko, and Stephan Winter. "Deriving place graphs from spatial databases." In: *Research@Locate*. Vol. 2087. CEUR Workshop Proceedings. CEUR-WS.org, 2018, pp. 25–32.

[153] Bo Han, Paul Cook, and Timothy Baldwin. "Text-Based Twitter User Geolocation Prediction." In: *J. Artif. Intell. Res.* 49 (2014), pp. 451–500. DOI: 10.1613/jair.4200. URL: https://doi.org/10.1613/jair.4200.

[154] Erum Haris, Keng Hoon Gan, and Tien-Ping Tan. "Spatial information extraction from travel narratives: Analysing the notion of co-occurrence indicating closeness of tourists." In: *Journal of Information Science* (2019), p. 0165551519837188. DOI: 10.1177/0165551519837188.

[155] Donna Harman. "Information retrieval evaluation." In: *Synthesis Lectures on Information Concepts, Retrieval, and Services* 3.2 (2011), pp. 1–119.

[156] Steve R. Harrison and Paul Dourish. "Re-place-ing Space: The Roles of Place and Space in Collaborative Systems." In: *CSCW '96, Proceedings of the ACM 1996 Conference on Computer Supported Cooperative Work, Boston, MA, USA, November 16-20, 1996*. Ed. by Mark S. Ackerman, Gary M. Olson, and Judith S. Olson. ACM, 1996, pp. 67–76. DOI: 10.1145/240080.240193. URL: https://doi.org/10.1145/240080.240193.

[157] Glen Hart, Martina Johnson, and Catherine Dolbear. "Rabbit: Developing a Control Natural Language for Authoring Ontologies." In: *The Semantic Web: Research and Applications, 5th European Semantic Web Conference, ESWC 2008, Tenerife, Canary Islands, Spain, June 1-5, 2008, Proceedings*. Ed. by Sean Bechhofer, Manfred Hauswirth, Jörg Hoffmann, and Manolis Koubarakis. Vol. 5021. Lecture Notes in Computer Science. Springer, 2008, pp. 348–360. DOI: 10.1007/978-3-540-68234-9_27. URL: https://doi.org/10.1007/978-3-540-68234-9_27.

[158] Stephen Harter and Carol Hert. "Evaluation of Information Retrieval Systems: Approaches, Issues, and Methods." In: *Annual Review of Information Science and Technology (ARIST)* 32 (Jan. 1997).

[159] Jordan T. Hastings and Linda L. Hill. "Georeferencing." In: *Encyclopedia of Database Systems*. Springer US, 2009, pp. 1246–1249.

[160] Daniel Hernández, Eliseo Clementini, and Paolino Di Felice. "Qualitative distances." In: *Spatial Information Theory A Theoretical Basis for GIS*. Ed. by Andrew U. Frank and Werner Kuhn. Berlin, Heidelberg: Springer Berlin Heidelberg, 1995, pp. 45–57. ISBN: 978-3-540-45519-6.

[161] Annette Herskovits. *Language and spatial cognition*. Cambridge university press Cambridge, 1986.

[162] Linda L. Hill. "Core Elements of Digital Gazetteers: Placenames, Categories, and Footprints." In: *Research and Advanced Technology for Digital Libraries*. Ed. by José Borbinha and Thomas Baker. Berlin, Heidelberg: Springer Berlin Heidelberg, 2000, pp. 280–290. ISBN: 978-3-540-45268-3.

[163] Linda L. Hill. *Georeferencing: The Geographic Associations of Information*. The MIT Press, Aug. 2006. ISBN: 9780262275385. DOI: 10.7551/mitpress/3260.001.0001. URL: https://doi.org/10.7551/mitpress/3260.001.0001.

[164] Linda L Hill. *Georeferencing: The geographic associations of information*. Mit Press, 2009.

[165] Heidelinde Hobel. "Extracting Semantics of Places from User Generated Content." In: *Proceedings of the 19th AGILE International Conference on Geographic Information Science*. 2016.

[166] Orland Hoeber and Xue Dong Yang. "User-oriented evaluation methods for interactive web search interfaces." In: *2007 IEEE/WIC/ACM International Conferences on Web Intelligence and Intelligent Agent Technology-Workshops*. IEEE. 2007, pp. 239–243.

[167] Johannes Hoffart, Fabian M. Suchanek, Klaus Berberich, and Gerhard Weikum. "YAGO2: A spatially and temporally enhanced knowledge base from Wikipedia." In: *Artif. Intell.* 194 (2013), pp. 28–61. DOI: 10.1016/j.artint.2012.06.001. URL: https://doi.org/10.1016/j.artint.2012.06.001.

[168] Katja Hofmann, Lihong Li, and Filip Radlinski. "Online evaluation for information retrieval." In: *Foundations and trends in information retrieval* 10.1 (2016), pp. 1–117.

[169] Livia Hollenstein and Ross Purves. "Exploring place through user-generated content: Using Flickr tags to describe city cores." In: *J. Spatial Inf. Sci.* 1.1 (2010), pp. 21–48. DOI: 10.5311/JOSIS.2010.1.3. URL: https://doi.org/10.5311/JOSIS.2010.1.3.

[170] Kathleen Stewart Hornsby and Naicong Li. "Conceptual Framework for Modeling Dynamic Paths from Natural Language Expressions." In: *Transactions in Gis* 13 (2009), pp. 27–45.

[171] Yingjie Hu. "Geospatial Semantics." In: *CoRR* abs/1707.03550 (2017). arXiv: 1707.03550.

[172] Yingjie Hu, Huina Mao, and Grant McKenzie. "A natural language processing and geospatial clustering framework for harvesting local place names from geotagged housing advertisements." In: *International Journal of Geographical Information Science* 33.4 (2019), pp. 714–738. DOI: 10.1080/13658816.2018.1458986.

[173] Hua Hua, Jochen Renz, and Xiaoyu Ge. "Qualitative representation and reasoning over direction relations across different frames of reference." In: *Sixteenth International Conference on Principles of Knowledge Representation and Reasoning*. 2018.

[174] Geoffrey Hubona, Stephanie Everett, Elaine Marsh, and Kenneth Wauchope. "Mental representations of spatial language." In: *Int. J. Hum.-Comput. Stud.* 48 (June 1998), pp. 705–728. DOI: 10.1006/ijhc.1998.0188.

[175] Drew A Hudson and Christopher D Manning. "Gqa: A new dataset for real-world visual reasoning and compositional question answering." In: *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 2019, pp. 6700–6709.

[176] G. Ritchie I. Androutsopoulos. "Database Interfaces." In: (2000), pp. 2–09–240.

[177] Nusrat Jahan, Sudha Morwal, and Deepti Chopra. "Named Entity Recognition in Indian Languages Using Gazetteer Method and Hidden Markov Model: A Hybrid Approach." In: 2012.

[178] Krzysztof Janowicz. "Sim-DL: Towards a Semantic Similarity Measurement Theory for the Description Logic *ALCNR* in Geographic Information Retrieval." In: *On the Move to Meaningful Internet Systems 2006: OTM 2006 Workshops, OTM Confederated International Workshops and Posters, AWeSOMe, CAMS, COMINF, IS, KSinBIT, MIOS-CIAO, MONET, OnToContent, ORM, PerSys, OTM Academy Doctoral Consortium, RDDS, SWWS, and SeBGIS 2006, Montpellier, France, October 29 - November 3, 2006. Proceedings, Part II*. Ed. by Robert Meersman, Zahir Tari, and Pilar Herrero. Vol. 4278. Lecture Notes in Computer Science. Springer, 2006, pp. 1681–1692. DOI: 10.1007/11915072_74. URL: https://doi.org/10.1007/11915072_74.

[179] Krzysztof Janowicz and Carsten Keßler. "The role of ontology in improving gazetteer interaction." In: *International Journal of Geographical Information Science* 22.10 (2008), pp. 1129–1157. DOI: 10.1080/13658810701851461.

[180] Krzysztof Janowicz, Sven Schade, Arne Bröring, Carsten Keßler, Patrick Maué, and Christoph Stasch. "Semantic Enablement for Spatial Data Infrastructures." In: *Trans. GIS* 14.2 (2010), pp. 111–129. DOI: 10.1111/j.1467-9671.2010.01186.x. URL: https://doi.org/10.1111/j.1467-9671.2010.01186.x.

[181] Kalervo Järvelin. "Evaluation." English. In: *Interactive Information Seeking, Behaviour and Retrieval*. Ed. by Ian Ruthven and Diane Kelly. 2011, pp. 113–138. ISBN: 978-1-85604-707-4.

[182] Jisha P Jayan, RR Rajeev, and Elizabeth Sherly. "A hybrid statistical approach for named entity recognition for malayalam language." In: *Proceedings of the 11th Workshop on Asian Language Resources*. 2013, pp. 58–63.

[183] Jing Jiang. "Information extraction from text." In: *Mining text data*. Springer, 2012, pp. 11–41.

[184] Agnew John. "Place and Politics: the geographical mediation of state and society." In: (1987). DOI: https://doi.org/10.4324/9781315756585.

[185] Mark L. Johnson. *The Body in the Mind: The Bodily Basis of Meaning, Imagination, and Reason*. University of Chicago Press, 1987.

[186] Sue E. Johnson, Pierre Jourlin, Karen Spärck Jones, and Philip C. Woodland. "Information Retrieval from Unsegmented Broadcast News Audio." In: *I. J. Speech Technology* 4.3-4 (2001), pp. 251–268.

[187] Christopher B. Jones and Ross S. Purves. "Geographical information retrieval." In: *Int. J. Geogr. Inf. Sci.* 22.3 (2008), pp. 219–228. DOI: 10.1080/13658810701626343. URL: https://doi.org/10.1080/13658810701626343.

[188] Christopher B. Jones and Ross S. Purves. "Geographical Information Retrieval." In: *Encyclopedia of Database Systems, Second Edition*. Ed. by Ling Liu and M. Tamer Özsu. Springer, 2018. DOI: 10.1007/978-1-4614-8265-9_177.

[189] Christopher B. Jones, Ross S. Purves, Paul D. Clough, and Hideo Joho. "Modelling vague places with knowledge from the Web." In: *Int. J. Geogr. Inf. Sci.* 22.10 (2008), pp. 1045–1065. DOI: 10.1080/13658810701850547. URL: https://doi.org/10.1080/13658810701850547.

[190] Christopher Jones, Alia Abdelmoty, David Finch, Gaihua Fu, and Subodh Vaid. "The SPIRIT Spatial Search Engine: Architecture, Ontologies and Spatial Indexing." In: Jan. 2005, pp. 125–139. ISBN: 978-3-540-23558-3. DOI: 10.1007/978-3-540-30231-5_9.

[191] Troy Jordan, Martin Raubal, Bryce Gartrell, and Max J. Egenhofer. *An Affordance-Based Model of Place in GIS*. 1998.

[192] Morteza Karimzadeh, Wenyi Huang, Siddhartha Banerjee, Jan Oliver Wallgrün, Frank Hardisty, Scott Pezanowski, Prasenjit Mitra, and Alan M. MacEachren. "GeoTxt: a web API to leverage place references in text." In: *Proceedings of the 7th Workshop on Geographic Information Retrieval, GIR 2013, 5th November, 2013, Orlando, Florida, USA*. Ed. by Christopher B. Jones and Ross Purves. ACM, 2013, pp. 72–73. DOI: 10.1145/2533888.2533942. URL: https://doi.org/10.1145/2533888.2533942.

[193] J. Kelleher, F. Costello, and J. Van Genabith. "Dynamically structuring, updating and interrelating representations of visual and linguistic discourse context." In: *Artificial Intelligence* 167.1–2 (2005), pp. 62–102.

[194] John D. Kelleher. "A perceptually based computational framework for the interpretation of spatial language." In: 2003.

[195] Diane Kelly. *Methods for evaluating interactive information retrieval systems with users*. Now Publishers Inc, 2009.

[196] Carsten Keßler, Krzysztof Janowicz, and Mohamed Bishr. "An Agenda for the next Generation Gazetteer: Geographic Information Contribution and Retrieval." In: *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. GIS '09. Seattle, Washington: Association for Computing Machinery, 2009, 91–100. ISBN: 9781605586496. DOI: 10.1145/1653771.1653787. URL: https://doi.org/10.1145/1653771.1653787.

[197] Jin-Dong Kim, Tomoko Ohta, Sampo Pyysalo, Yoshinobu Kano, and Jun'ichi Tsujii. "Overview of BioNLP'09 Shared Task on Event Extraction." In: *Proceedings of the BioNLP 2009 Workshop Companion Volume for Shared Task*. Boulder, Colorado: Association for Computational Linguistics, June 2009, pp. 1–9. URL: https://aclanthology.org/W09-1401.

[198] Junchul Kim, Maria Vasardani, and Stephan Winter. "From descriptions to depictions: A dynamic sketch map drawing strategy." In: *Spatial Cognition & Computation* 16.1 (2016), pp. 29–53. DOI: 10.1080/13875868.2015.1084509. eprint: https://doi.org/10.1080/13875868.2015.1084509. URL: https://doi.org/10.1080/13875868.2015.1084509.

[199] Junchul Kim, Maria Vasardani, and Stephan Winter. "Landmark Extraction from Web-Harvested Place Descriptions." In: *KI - Artificial Intelligence* 31.2 (2017), pp. 151–159. DOI: 10.1007/s13218-016-0467-3.

[200] Alexander Klippel, Sen Xu, Jinlong Yang, and Rui Li. "Spatial Event Language across Geographic Domains." In: *CEUR Workshop Proceedings* 759 (Jan. 2011).

[201] Markus Knauff, Reinhold Rauh, and Christoph Schlieder. "Preferred mental models in qualitative spatial reasoning: A cognitive assessment of Allen's calculus." In: *Proceedings of the seventeenth annual conference of the Cognitive Science Society*. 1995, pp. 200–205.

[202] Markus Knauff, Reinhold Rauh, Christoph Schlieder, and Gerhard Strube. "Mental Models in Spatial Reasoning." In: *Spatial Cognition: An Interdisciplinary Approach to Representing and Processing Spatial Knowledge*. Ed. by Christian Freksa, Christopher Habel, and Karl F. Wender. Berlin, Heidelberg: Springer Berlin Heidelberg, 1998, pp. 267–291. ISBN: 978-3-540-69342-0. DOI: 10.1007/3-540-69342-4_13. URL: https://doi.org/10.1007/3-540-69342-4_13.

[203] Markus Knauff, Gerhard Strube, Corinne Jola, Reinhold Rauh, and Christoph Schlieder. "The Psychological Validity of Qualitative Spatial Reasoning in One Dimension." In: 4.2 (2004), pp. 167–188.

[204] Donald E. Knuth. "Computer Programming as an Art." In: *Communications of the ACM* 17.12 (1974), pp. 667–673.

[205] Michal Konkol, Tomáš Brychcín, and Miloslav Konopík. "Latent semantics in Named Entity Recognition." In: *Expert Systems with Applications* 42.7 (2015), pp. 3470–3479. ISSN: 0957-4174. DOI: https://doi.org/10.1016/j.eswa.2014.12.015. URL: https://www.sciencedirect.com/science/article/pii/S0957417414007933.

[206] Natalia Konstantinova. "Review of Relation Extraction Methods: What Is New Out There?" In: *AIST*. 2014.

[207] Parisa Kordjamshidi, Martijn Otterlo, and Marie-Francine Moens. "Spatial role labeling: Towards extraction of spatial relations from natural language." In: *TSLP* 8 (Dec. 2011), p. 4. DOI: 10.1145/2050104.2050105.

[208] Parisa Kordjamshidi, Martijn Van Otterlo, and Marie-Francine Moens. "Spatial Role Labeling: Task Definition and Annotation Scheme." In: *Proceedings of the Seventh International Conference on Language Resources and Evaluation (LREC'10)*. Valletta, Malta: European Language Resources Association (ELRA), May 2010. URL: http://www.lrec-conf.org/proceedings/lrec2010/pdf/846_Paper.pdf.

[209] Marcus Kracht. "On the Semantics of Locatives." In: *Linguistics and Philosophy* 25.2 (2002), pp. 157–232. DOI: 10.1023/A:1014646826099.

[210] Christian Kray, Holger Fritze, Thore Fechner, Angela Schwering, Rui Li, and Vanessa Anacta. "Transitional Spaces: Between Indoor and Outdoor Spaces." In: vol. 8116. Sept. 2013, pp. 14–32. DOI: 10.1007/978-3-319-01790-7_2.

[211] Arne Kreutzmann, Immo Colonius, Diedrich Wolter, Frank Dylla, Lutz Frommberger, and Christian Freksa. "Temporal logic for process specification and recognition." In: *Intell. Serv. Robotics* 6.1 (2013), pp. 5–18. DOI: 10.1007/s11370-012-0122-2. URL: https://doi.org/10.1007/s11370-012-0122-2.

[212] Dilek Küçük and Adnan Yazici. "A hybrid named entity recognizer for Turkish." In: *Expert Syst. Appl.* 39.3 (2012), pp. 2733–2742. DOI: 10.1016/j.eswa.2011.08.131. URL: https://doi.org/10.1016/j.eswa.2011.08.131.

[213] Werner Kuhn. "Ontologies in support of activities in geographical space." In: *International Journal of Geographical Information Science* 15.7 (2001), pp. 613–631. DOI: 10.1080/13658810110061180.

[214] Werner Kuhn. "Core concepts of spatial information for transdisciplinary research." In: *International Journal of Geographical Information Science* 26 (12 2012), pp. 2267–2276.

[215] Benjamin Kuipers. "Modeling Spatial Knowledge." In: 2.2 (1978), pp. 129–153.

[216] Benjamin Kuipers, Dan G. Tecuci, and Brian J. Stankiewicz. "The Skeleton In The Cognitive Map: A Computational and Empirical Exploration." In: *Environment and Behavior* 35.1 (2003), pp. 81–106. DOI: 10.1177/0013916502238866. eprint: https://doi.org/10.1177/0013916502238866. URL: https://doi.org/10.1177/0013916502238866.

[217] Lars Kulik. "A geometric theory of vague boundaries based on supervaluation." In: *International Conference on Spatial Information Theory*. 2001, pp. 44–59.

[218] Abhinav Kumar and Jyoti Prakash Singh. "Location reference identification from tweets during emergencies: A deep learning approach." In: *International journal of disaster risk reduction* 33 (2019), pp. 365–375.

[219] Edward Kurwakumire. "Towards a Public Sector GIS Evaluation Methodology." In: *South African Journal of Geomatics* 3 (2014), pp. 33–52.

[220] Olivier Van Laere, Steven Schockaert, Vlad Tanasescu, Bart Dhoedt, and Christopher B. Jones. "Georeferencing Wikipedia Documents Using Data from Social Media Sources." In: *ACM Trans. Inf. Syst.* 32.3 (2014), 12:1–12:32. DOI: 10.1145/2629685. URL: https://doi.org/10.1145/2629685.

[221] George Lakoff and Mark Johnson. "Conceptual metaphor in everyday language." In: *The journal of Philosophy* 77.8 (1980), pp. 453–486.

[222] El Khadir Lamrani, El Habib Ben Lahmar, Abdelaziz Marzak, and Hammad Ballaoui. “Mixed method for extraction of domain terminology from text: Linguistic and statistical filtering.” In: *2014 Third IEEE International Colloquium in Information Science and Technology (CIST)*. 2014, pp. 291–295. DOI: 10.1109/CIST.2014.7016634.

[223] Barbara Landau. “Where’s what and what’s where: The language of objects in space.” In: *Lingua* 92 (1994), pp. 259–296.

[224] Barbara Landau and Ray Jackendoff. ““What” and “Where” in spatial language and spatial cognition.” In: *Behavioral and Brain Sciences* 16.2 (1993), pp. 217–238.

[225] Admond Lee. “why NLP is important and it will be the future-our future.” In: (2019). URL: https://towardsdatascience.com/why-nlp-is-important-and-itll-be-the-future-our-future-59d7b1600dda.

[226] Jochen L. Leidner. “Toponym resolution in text (abstract only): "which sheffield is it?"” In: *SIGIR 2004: Proceedings of the 27th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval, Sheffield, UK, July 25-29, 2004*. Ed. by Mark Sanderson, Kalervo Järvelin, James Allan, and Peter Bruza. ACM, 2004, p. 602. DOI: 10.1145/1008992.1009147. URL: https://doi.org/10.1145/1008992.1009147.

[227] Jochen L Leidner. “An evaluation dataset for the toponym resolution task.” In: *Computers, Environment and Urban Systems* 30.4 (2006), pp. 400–417.

[228] Jochen L. Leidner. “Toponym resolution in text: Annotation, evaluation and applications of spatial grounding.” In: *SIGIR Forum* 41 (Jan. 2007), pp. 124–126.

[229] Jochen L. Leidner. “Toponym resolution in text: annotation, evaluation and applications of spatial grounding.” In: *SIGIR Forum* 41.2 (2007), pp. 124–126. DOI: 10.1145/1328964.1328989.

[230] Jochen L. Leidner and Michael D. Lieberman. “Detecting geographical references in the form of place names and associated spatial natural language.” In: *SIGSPATIAL Special* 3.2 (2011), pp. 5–11. DOI: 10.1145/2047296.2047298.

[231] Johannes Leveling. “Tagging of Temporal Expressions and Geological Features in Scientific Articles.” In: *Proceedings of the 9th Workshop on Geographic Information Retrieval*. GIR ’15. Paris, France: Association for Computing Machinery, 2015. ISBN: 9781450339377. DOI: 10.1145/2837689.2837701. URL: https://doi.org/10.1145/2837689.2837701.

[232] S. Levinson. “Frames of reference and Molyneux’s question: Cross-linguistic evidence.” In: 1996.

[233] Stephen C. Levinson. *Space in Language and Cognition: Explorations in Cognitive Diversity*. Cambridge University Press, 2003.

[234] Hanjing Li, Tiejun Zhao, Sheng Li, and Jiyuan Zhao. "The Extraction of Trajectories from Real Texts Based on Linear Classification." In: *Proceedings of the 16th Nordic Conference of Computational Linguistics, NODALIDA 2007, Tartu, Estonia, May 2007*. Ed. by Joakim Nivre, Heiki-Jaan Kaalep, Kadri Muischnek, and Mare Koit. University of Tartu, Estonia, 2007, pp. 121–127. URL: https://aclanthology.org/W07-2418/.

[235] Peipei Li, Haixun Wang, Hongsong Li, and Xindong Wu. "Employing semantic context for sparse information extraction assessment." In: *ACM Transactions on Knowledge Discovery from Data (TKDD)* 12.5 (2018), pp. 1–36.

[236] Pengfei Li and Kezhi Mao. "Knowledge-oriented Convolutional Neural Network for Causal Relation Extraction from Natural Language Texts." In: *Expert Systems with Applications* 115 (Aug. 2018). DOI: 10.1016/j.eswa.2018.08.009.

[237] W. Li, Chaowei Yang, Doug Nebert, R. Raskin, Paul R. Houser, Huayi Wu, and Zhenlong Li. "Semantic-based web service discovery and chaining for building an Arctic spatial data infrastructure." In: *Comput. Geosci.* 37.11 (2011), pp. 1752–1762. DOI: 10.1016/j.cageo.2011.06.024. URL: https://doi.org/10.1016/j.cageo.2011.06.024.

[238] Wenwen Li, Robert Raskin, and Michael F. Goodchild. "Semantic similarity measurement based on knowledge mining: an artificial neural net approach." In: *International Journal of Geographical Information Science* 26.8 (2012), pp. 1415–1435. DOI: 10.1080/13658816.2011.635595.

[239] Elizabeth D. Liddy. "Enhanced Text Retrieval Using Natural Language Processing." In: *Bulletin of the American Society for Information Science and Technology* 24.4 (1998), pp. 14–16. DOI: https://doi.org/10.1002/bult.91. eprint: https://asistdl.onlinelibrary.wiley.com/doi/pdf/10.1002/bult.91. URL: https://asistdl.onlinelibrary.wiley.com/doi/abs/10.1002/bult.91.

[240] Michael D. Lieberman, Hanan Samet, and Jagan Sankaranarayanan. "Geotagging with local lexicons to build indexes for textually-specified spatial data." In: *2010 IEEE 26th International Conference on Data Engineering (ICDE 2010)*. 2010, pp. 201–212. DOI: 10.1109/ICDE.2010.5447903.

[241] Gérard Ligozat. *Qualitative Spatial and Temporal Reasoning*. 2011.

[242] Gérard Ligozat and Jochen Renz. "What Is a Qualitative Calculus? A General Framework." In: vol. 3157. LNCS. Springer, 2004, pp. 53–64.

[243] Kenneth C. Litkowski and Orin Hargraves. "SemEval-2007 Task 06: Word-Sense Disambiguation of Prepositions." In: *Proceedings of the 4th International Workshop on Semantic Evaluations, SemEval@ACL 2007, Prague, Czech Republic, June 23-24, 2007*. Ed. by Eneko Agirre, Lluís Màrquez i Villodre, and Richard Wicentowski. The Association for Computer Linguistics, 2007, pp. 24–29. URL: https://aclanthology.org/S07-1005/.

[244] Fei Liu, Maria Vasardani, and Timothy Baldwin. "Automatic Identification of Locative Expressions from Social Media Text: A Comparative Analysis." In: *Proceedings of the 4th International Workshop on Location and the Web*. LocWeb '14. Shanghai, China: Association for Computing Machinery, 2014, 9–16. ISBN: 9781450314596. DOI: 10.

1145/2663713.2664426. URL: https://doi.org/10.1145/2663713.2664426.

[245] Ji Liu and Diana Inkpen. "Estimating User Location in Social Media with Stacked Denoising Auto-encoders." In: *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing, VS@NAACL-HLT 2015, June 5, 2015, Denver, Colorado, USA*. Ed. by Phil Blunsom, Shay B. Cohen, Paramveer S. Dhillon, and Percy Liang. The Association for Computational Linguistics, 2015, pp. 201–210. DOI: 10.3115/v1/w15-1527. URL: https://doi.org/10.3115/v1/w15-1527.

[246] Kaijian Liu and Nora El-Gohary. "Ontology-based semi-supervised conditional random fields for automated information extraction from bridge inspection reports." English (US). In: *Automation in Construction* 81 (Sept. 2017), pp. 313–327. ISSN: 0926-5805. DOI: 10.1016/j.autcon.2017.02.003.

[247] Xiaohua Liu and Ming Zhou. "Two-Stage NER for Tweets with Clustering." In: *Inf. Process. Manage.* 49.1 (Jan. 2013), 264–273. ISSN: 0306-4573. DOI: 10.1016/j.ipm.2012.05.006. URL: https://doi.org/10.1016/j.ipm.2012.05.006.

[248] Yu Liu, Xiaoming Wang, Xin Jin, and Lun Wu. "On internal cardinal direction relations." In: *International Conference on Spatial Information Theory*. Springer. 2005, pp. 283–299.

[249] Richard K. Lomotey and Ralph Deters. "Topics and Terms Mining in Unstructured Data Stores." In: *2013 IEEE 16th International Conference on Computational Science and Engineering*. 2013, pp. 854–861. DOI: 10.1109/CSE.2013.129.

[250] Vitor Loureiro, Ivo Anastácio, and Bruno Martins. "Learning to Resolve Geographical and Temporal References in Text." In: *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. GIS '11. Chicago, Illinois: Association for Computing Machinery, 2011, 349–352. ISBN: 9781450310314. DOI: 10.1145/2093973.2094020. URL: https://doi.org/10.1145/2093973.2094020.

[251] Chandrika M, Parvathy J J, Srinidhi B S, Vijay Bhandi, and S Balaji. "IMPACT OF BIG DATA AND EMERGING RESEARCH TRENDS." In: *International Journal of Innovative Technology and Research* 0.0 (2015). ISSN: 2320 –5547. URL: https://www.ijitr.com/index.php/ojs/article/view/539.

[252] Kai Ma, Liang Wu, Liufeng Tao, Wenjia Li, and Zhong Xie. "Matching Descriptions to Spatial Entities Using a Siamese Hierarchical Attention Network." In: *IEEE Access* 6 (2018), pp. 28064–28072. DOI: 10.1109/ACCESS.2018.2837666.

[253] Mukta Majumder, Utsav Barman, Rahul Prasad, Kumar Saurabh, and Sujan Kumar Saha. "A Novel Technique for Name Identification from Homeopathy Diagnosis Discussion Forum." In: *Procedia Technology* 6 (2012). 2nd International Conference on Communication, Computing & Security [ICCCS-2012], pp. 379–386. ISSN: 2212-0173. DOI: https://doi.org/10.1016/j.protcy.2012.10.045. URL: https://www.sciencedirect.com/science/article/pii/S2212017312005907.

[254] Thomas Mandl. "Evaluating GIR: geography-oriented or user-oriented?" In: *ACM SIGSPATIAL Special* 3.2 (2011), pp. 42–45. DOI: 10.1145/2047296.2047306. URL: https://doi.org/10.1145/2047296.2047306.

[255] Inderjeet Mani. "SpatialML: Annotation Scheme for Marking Spatial Expressions in Natural Language." In: 2008.

[256] M. Mannai and Wahiba Abdessalem. "Bayesian information extraction network for Medline abstract." In: June 2013, pp. 1–3. ISBN: 978-1-4799-0460-0. DOI: 10.1109/WCCIT.2013.6618668.

[257] Mónica Marrero, Julián Urbano, Sonia Sánchez-Cuadrado, Jorge Morato, and Juan Miguel Gómez-Berbís. "Named entity recognition: fallacies, challenges and opportunities." In: *Computer Standards & Interfaces* 35.5 (2013), pp. 482–489.

[258] Bruno Martins. "A Supervised Machine Learning Approach for Duplicate Detection over Gazetteer Records." In: *Proceedings of the 4th International Conference on GeoSpatial Semantics*. GeoS'11. Brest, France: Springer-Verlag, 2011, 34–51. ISBN: 9783642206290.

[259] Alexey Mazalov, Bruno Martins, and David Matos. "Spatial Role Labeling with Convolutional Neural Networks." In: *Proceedings of the 9th Workshop on Geographic Information Retrieval*. GIR '15. Paris, France: Association for Computing Machinery, 2015. ISBN: 9781450339377. DOI: 10.1145/2837689.2837706. URL: https://doi.org/10.1145/2837689.2837706.

[260] Fernando Melo and Bruno Martins. "Automated Geocoding of Textual Documents: A Survey of Current Approaches." In: *Transactions in GIS* 21.1 (2017), pp. 3–38. DOI: https://doi.org/10.1111/tgis.12212. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/tgis.12212. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/tgis.12212.

[261] Jean-Baptiste Michel, Yuan Kui Shen, Aviva Presser Aiden, Adrian Veres, Matthew K Gray, Google Books Team, Joseph P Pickett, Dale Hoiberg, Dan Clancy, Peter Norvig, et al. "Quantitative analysis of culture using millions of digitized books." In: *science* 331.6014 (2011), pp. 176–182.

[262] Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. "Distributed Representations of Words and Phrases and their Compositionality." In: *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*. Ed. by Christopher J. C. Burges, Léon Bottou, Zoubin Ghahramani, and Kilian Q. Weinberger. 2013, pp. 3111–3119. URL: http://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-compositionality.

[263] George A. Miller. "WordNet: A Lexical Database for English." In: *Communications of the ACM* 38.11 (1995), pp. 39–41.

[264] Harvey Miller. "Place-Based versus People-Based Geographic Information Science." In: *Geography Compass* 1.3 (2007), pp. 503–535. DOI: https://doi.org/10.1111/j.1749-8198.2007.00025.x. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1749-8198.2007.00025.x. URL: https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1749-8198.2007.00025.x.

[265] Ludovic Moncla, Walter Renteria-Agualimpia, Javier Nogueras-Iso, and Mauro Gaio. "Geocoding for texts with fine-grain toponyms: an experiment on a geoparsed hiking descriptions corpus." In: *Proceedings of the 22nd ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems, Dallas/Fort Worth, TX, USA, November 4-7, 2014*. Ed. by Yan Huang, Markus Schneider, Michael Gertz, John Krumm, and Jagan Sankaranarayanan. ACM, 2014, pp. 183–192. DOI: 10.1145/2666310.2666386. URL: https://doi.org/10.1145/2666310.2666386.

[266] Bruno R. Monteiro, Clodoveu A. Davis Jr., and Frederico T. Fonseca. "A survey on the geographic scope of textual documents." In: *Computers & Geosciences* 96 (2016), pp. 23–34. DOI: 10.1016/j.cageo.2016.07.017.

[267] Daniel R. Montello, Alinda Friedman, and Daniel W. Phillips. "Vague cognitive regions in geography and geographic information science." In: *Int. J. Geogr. Inf. Sci.* 28.9 (2014), pp. 1802–1820. DOI: 10.1080/13658816.2014.900178. URL: https://doi.org/10.1080/13658816.2014.900178.

[268] Daniel R. Montello, Michael F. Goodchild, Jonathon Gottsegen, and Peter Fohl. "Where's Downtown? Behavioral Methods for Determining Referents of Vague Spatial Queries." In: *Spatial Cogn. Comput.* 3.2-3 (2003), pp. 185–204. DOI: 10.1080/13875868.2003.9683761. URL: https://doi.org/10.1080/13875868.2003.9683761.

[269] John Montgomery. "Making a city: Urbanity, vitality and urban design." In: *Journal of Urban Design* 3.1 (1998), pp. 93–116. DOI: 10.1080/13574809808724418. eprint: https://doi.org/10.1080/13574809808724418. URL: https://doi.org/10.1080/135748098087-24418.

[270] Vanessa Murdock. "Dynamic location models." In: *The 37th International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR '14, Gold Coast , QLD, Australia - July 06 - 11, 2014*. Ed. by Shlomo Geva, Andrew Trotman, Peter Bruza, Charles L. A. Clarke, and Kalervo Järvelin. ACM, 2014, pp. 1231–1234. DOI: 10.1145/2600428.2609552. URL: https://doi.org/10.1145/2600428.2609552.

[271] David Nadeau and Satoshi Sekine. "A Survey of Named Entity Recognition and Classification." In: *Lingvisticae Investigationes* 30 (Aug. 2007). DOI: 10.1075/li.30.1.03nad.

[272] Zorica Nedovic-Budic. "Evaluating the Effects of GIS Technology: Review of Methods." In: *Journal of Planning Literature - J PLAN LIT* 13 (Feb. 1999), pp. 284–295. DOI: 10.1177/08854129922092405.

[273] Z Nedović-Budić. "The Impact of GIS Technology." In: *Environment and Planning B: Planning and Design* 25.5 (1998), pp. 681–692. DOI: 10.1068/b250681. eprint: https://doi.org/10.1068/b250681. URL: https://doi.org/10.1068/b250681.

[274] Thien Huu Nguyen and Ralph Grishman. *Combining Neural Networks and Log-linear Models to Improve Relation Extraction*. 2015. arXiv: 1511.05926 [cs.CL].

[275] Thien Huu Nguyen and Ralph Grishman. "Relation Extraction: Perspective from Convolutional Neural Networks." In: *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*. Denver, Colorado: Association for Computational Linguistics, June 2015, pp. 39–48. DOI: 10.3115/v1/W15-1506. URL: https://aclanthology.org/W15-1506.

[276] Eric Nichols and Fadi Botros. "SpRL-CWW: Spatial Relation Classification with Independent Multi-class Models." In: Jan. 2015, pp. 895–901. DOI: 10.18653/v1/S15-2150.

[277] Joakim Nivre. *Dependency grammar and dependency parsing*. Tech. rep. Växjö University, 2005.

[278] Simon Overell and Stefan Rüger. "Using Co-Occurrence Models for Placename Disambiguation." In: *Int. J. Geogr. Inf. Sci.* 22.3 (2008), 265–287. ISSN: 1365-8816. DOI: 10.1080/13658810701626236. URL: https://doi.org/10.1080/13658810701626236.

[279] Damien Palacio, Curdin Derungs, and Ross Purves. "Development and evaluation of a geographic information retrieval system using fine grained toponyms." In: *Journal of Spatial Information Science* 2015 (Dec. 2015). DOI: 10.5311/JOSIS.2015.11.193.

[280] Damien Palacio, Guillaume Cabanac, Christian Sallaberry, and Gilles Hubert. "On the evaluation of Geographic Information Retrieval systems - Evaluation framework and case study." In: *Int. J. Digit. Libr.* 11.2 (2010), pp. 91–109. DOI: 10.1007/s00799-011-0070-z. URL: https://doi.org/10.1007/s00799-011-0070-z.

[281] Emmanuel Papadakis. "Place-based GIS: Functional Space." In: *Proceedings of the 4th AGILE PhD School, Leeds, UK, October 30 - November 02, 2017*. Ed. by Lex Comber and Nick Malleson. Vol. 2088. CEUR Workshop Proceedings. CEUR-WS.org, 2017. URL: http://ceur-ws.org/Vol-2088/paper3.pdf.

[282] Emmanuel Papadakis, Bernd Resch, and Thomas Blaschke. "Composition of place: towards a compositional view of functional space." In: *Cartography and Geographic Information Science* 47 (June 2019), pp. 1–18. DOI: 10.1080/15230406.2019.1598894.

[283] Emmanuel Papadakis, George Baryannis, Andreas Petutschnig, and Thomas Blaschke. "Function-Based Search of Place Using Theoretical, Empirical and Probabilistic Patterns." In: *ISPRS Int. J. Geo Inf.* 8.2 (2019), p. 92. DOI: 10.3390/ijgi8020092. URL: https://doi.org/10.3390/ijgi8020092.

[284] Christos Papadimitriou. "Algorithms, complexity, and the sciences." In: *Proceedings of the National Academy of Sciences* 111.45 (2014), pp. 15881–15887. ISSN: 0027-8424. DOI: 10.1073/pnas.1416954111. eprint: https://www.pnas.org/content/111/45/15881.full.pdf. URL: https://www.pnas.org/content/111/45/15881.

[285] David J.Maguire David W.Rhind Paul A.Longley Michael F. Goodchild, ed. *Geographic Information Science and Systems*. John Wiley & Sons, 2015.

[286] Ted Pedersen, Siddharth Patwardhan, and Jason Michelizzi. "WordNet: : Similarity - Measuring the Relatedness of Concepts." In: *Proceedings of the Nineteenth National Conference on Artificial Intelligence, Sixteenth Conference on Innovative Applications of Artificial Intelligence, July 25-29, 2004, San Jose, California, USA*. Ed. by Deborah L. McGuinness and George Ferguson. AAAI Press / The MIT Press, 2004, pp. 1024–1025. URL: http://www.aaai.org/Library/AAAI/2004/aaai04-160.php.

[287] Peter L. T. Pirolli. *Information Foraging Theory: Adaptive Interaction with Information*. 1st ed. USA: Oxford University Press, Inc., 2007. ISBN: 0195173325.

[288] Jakub Piskorski and Roman Yangarber. "Information extraction: Past, present and future." In: *Multi-source, multilingual information extraction and summarization*. Springer, 2013, pp. 23–49.

[289] Bruno Pouliquen et al. "Geocoding Multilingual Texts: Recognition, Disambiguation and Visualisation." In: *Proceedings of the Fifth International Conference on Language Resources and Evaluation, LREC 2006, Genoa, Italy, May 22-28, 2006*. Ed. by Nicoletta Calzolari, Khalid Choukri, Aldo Gangemi, Bente Maegaard, Joseph Mariani, Jan Odijk, and Daniel Tapias. European Language Resources Association (ELRA), 2006, pp. 53–58. URL: http://www.lrec-conf.org/proceedings/lrec2006/summaries/578.html.

[290] Gary D. Prideaux. "George Lakoff. Women, Fire, and Dangerous Things: What Categories Reveal about the Mind. Chicago: University of Chicago Press. 1987. Pp. xvii 614." In: *Canadian Journal of Linguistics/Revue canadienne de linguistique* 33.3 (1988), 288–291. DOI: 10.1017/S0008413100013050.

[291] Ross S. Purves, Stephan Winter, and Werner Kuhn. "Places in Information Science." In: *J. Assoc. Inf. Sci. Technol.* 70.11 (2019), pp. 1173–1182. DOI: 10.1002/asi.24194. URL: https://doi.org/10.1002/asi.24194.

[292] Ross S. Purves et al. "The design and implementation of SPIRIT: a spatially aware search engine for information retrieval on the Internet." In: *Int. J. Geogr. Inf. Sci.* 21.7 (2007), pp. 717–745. DOI: 10.1080/13658810601169840. URL: https://doi.org/10.1080/13658810601169840.

[293] J. Pustejovsky, Jessica L. Moszkowicz, and M. Verhagen. "Using ISO-Space for Annotating Spatial Information." In: 2011.

[294] Peiyuan Qiu, Li Yu, Jialiang Gao, and Feng Lu. "Detecting geo-relation phrases from web texts for triplet extraction of geographic knowledge: a context-enhanced method." In: *Big Earth Data* 3.3 (2019), pp. 297–314. DOI: 10.1080/20964471.2019.1657719.

[295] Qinjun Qiu, Zhong Xie, Liang Wu, and Liufeng Tao. "GNER: A Generative Model for Geological Named Entity Recognition Without Labeled Data Using Deep Learning." In: *Earth and Space Science* 6.6 (2019), pp. 931–946. DOI: https://doi.org/10.1029/2019EA000610. eprint: https://agupubs.onlinelibrary.wiley.com/doi/pdf/10.1029/2019EA000610. URL: https://agupubs.onlinelibrary.wiley.com/doi/abs/10.1029/2019EA000610.

[296] Alexandra Pomares Quimbaya, Alejandro Sierra Múnera, Rafael Andrés González Rivera, Julián Camilo Daza Rodríguez, Oscar Mauricio Muñoz Velandia, Angel Alberto Garcia Peña, and Cyril Labbé. "Named Entity Recognition Over Electronic Health Records Through a Combined Dictionary-based Approach." In: *Procedia Computer Science* 100 (2016). International Conference on ENTERprise Information Systems/International Conference on Project MANagement/International Conference on Health and Social Care Information Systems and Technologies, CENTERIS/ProjMAN / HCist 2016, pp. 55–61. ISSN: 1877-0509. DOI: https://doi.org/10.1016/j.procs.2016.09.123. URL: https://www.sciencedirect.com/science/article/pii/S1877050916322906.

[297] K.R. Rahem and Nazlia Omar. "Rule-based named entity recognition for drugrelated crime news documents." In: *Journal of Theoretical and Applied Information Technology* 77 (July 2015), pp. 229–235.

[298] David A. Randell, Zhan Cui, and Anthony G. Cohn. "A Spatial Logic based on Regions and "Connection"." In: Morgan Kaufmann, 1992, pp. 165–176.

[299] Martin Raubal and Stephan Winter. "Enriching Wayfinding Instructions with Local Landmarks." In: *Geographic Information Science, Second International Conference, GIScience 2002, Boulder, CO, USA, September 25-28, 2002, Proceedings*. Ed. by Max J. Egenhofer and David M. Mark. Vol. 2478. Lecture Notes in Computer Science. Springer, 2002, pp. 243–259. DOI: 10.1007/3-540-45799-2_17. URL: https://doi.org/10.1007/3-540-45799-2_17.

[300] Dan Remenyi, Terry White, and Michael Sherwood-Smith. *Achieving Maximum Value from Information Systems: A Process Approach*. USA: John Wiley & Sons, Inc., 1997. ISBN: 0471975001.

[301] Jochen Renz and Bernhard Nebel. "Qualitative spatial reasoning using constraint calculi." In: *Handbook of Spatial Logics*. Ed. by Marco Aiello, Ian E. Pratt-Hartmann, and Johan F.A.K. van Benthem. Springer, 2007, pp. 161–215.

[302] Gudula Retz-Schmidt. "Various Views on Spatial Prepositions." In: *AI Mag.* 9.2 (1988), pp. 95–105. URL: https://ojs.aaai.org/index.php/aimagazine/article/view/678.

[303] Shanel Reyes-Palacios, Edwin Aldana-Bobadilla, Ivan Lopez-Arevalo, and Alejandro Molina-Villegas. "Georeference Assignment of Locations based on Context." In: *Proceedings of the 1st International Conference on Geospatial Information Sciences*. Ed. by Oscar S. Siordia, José Luis Silván Cárdenas, Alejandro Molina-Villegas, Gandhi Hernandez, Pablo Lopez-Ramirez, Rodrigo Tapia-McClung, Karime González Zuccolotto, and Mario Chirinos Colunga. Vol. 13. Kalpa Publications in Computing. EasyChair, 2019, pp. 39–46. DOI: 10.29007/nxnc.

[304] Kashif Riaz. "Rule-Based Named Entity Recognition in Urdu." In: *NEWS@ACL*. 2010.

[305] Daniela Richter, Stephan Winter, Kai-Florian Richter, and Lesly Stirling. "Granuality of Locations referred to by Place descriptions." In: *Proceedings of Workshops and Posters at the 13th International Conference of Spatial Information Theory (COSIT 2017)*. computer environments and urban systems. Elsevier, 2012, pp. 88–99.

[306] C. J. Van Rijsbergen. *Information Retrieval*. 2nd. USA: Butterworth-Heinemann, 1979. ISBN: 0408709294.

[307] Kirk Roberts, Cosmin Adrian Bejan, and Sanda M. Harabagiu. "Toponym Disambiguation Using Events." In: *Proceedings of the Twenty-Third International Florida Artificial Intelligence Research Society Conference, May 19-21, 2010, Daytona Beach, Florida, USA*. Ed. by Hans W. Guesgen and R. Charles Murray. AAAI Press, 2010. URL: http://www.aaai.org/ocs/index.php/FLAIRS/2010/paper/view/1291.

[308] Stephen E. Robertson. "The methodology of information retrieval experiment." In: 1981.

[309] Tim Rocktäschel, Michael Weidlich, and Ulf Leser. "ChemSpot: a hybrid system for chemical named entity recognition." In: *Bioinformatics* 28 12 (2012), pp. 1633–40.

[310] Stephen Roller, Michael Speriosu, Sarat Rallapalli, Benjamin Wing, and Jason Baldridge. "Supervised Text-based Geolocation Using Language Models on an Adaptive Grid." In: *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning, EMNLP-CoNLL 2012, July 12-14, 2012, Jeju Island, Korea*. Ed. by Jun'ichi Tsujii, James Henderson, and Marius Pasca. ACL, 2012, pp. 1500–1510. URL: https://aclanthology.org/D12-1137/.

[311] Miguel Torres Ruiz, Rolando Quintero, Marco Moreno, and Frederico T. Fonseca. "Ontology-Driven Description of Spatial Data for Their Semantic Processing." In: *GeoSpatial Semantics, First International Conference, GeoS 2005, Mexico City, Mexico, November 29-30, 2005, Proceedings*. Ed. by M. Andrea Rodríguez, Isabel F. Cruz, Max J. Egenhofer, and Sergei Levashkin. Vol. 3799. Lecture Notes in Computer Science. Springer, 2005, pp. 242–249. DOI: 10.1007/11586180_18. URL: https://doi.org/10.1007/11586180_18.

[312] Octavian Rusu, Ionela Halcu, Oana Grigoriu, Giorgian Neculoiu, Virginia Sandulescu, Mariana Marinescu, and Viorel Marinescu. "Converting unstructured and semi-structured data into knowledge." In: *2013 11th RoEduNet International Conference*. 2013, pp. 1–4. DOI: 10.1109/RoEduNet.2013.6511736.

[313] Sujan Kumar Saha, Sanjay Chatterji, Sandipan Dandapat, West Bengal, Sudeshna Sarkar, and Pabitra Mitra. "A Hybrid Approach for Named Entity Recognition in Indian Languages." In: 2008.

[314] Sujan Kumar Saha, Shashi Narayan, Sudeshna Sarkar, and Pabitra Mitra. "A composite kernel for named entity recognition." In: *Pattern Recognition Letters* 31.12 (2010), pp. 1591–1597.

[315] Haritz Salaberri, O. Arregi, and Beñat Zapirain. "IXAGroupEHUSpaceEval: (X-Space) A WordNet-based approach towards the Automatic Recognition of Spatial Information following the ISO-Space Annotation Scheme." In: Jan. 2015, pp. 856–861. DOI: 10.18653/v1/S15-2145.

[316] Mark Sanderson. "Test Collection Based Evaluation of Information Retrieval Systems." In: *Foundations and Trends in Information Retrieval* 4 (Jan. 2010), pp. 247–375. DOI: 10.1561/1500000009.

[317] Tefko Saracevic. "Evaluation of Evaluation in Information Retrieval." In: *Proceedings of the 18th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*. SIGIR '95. Seattle, Washington, USA: Association for Computing Machinery, 1995, 138–146. ISBN: 0897917146. DOI: 10.1145/215206.215351. URL: https://doi.org/10.1145/215206.215351.

[318] Yunita Sari, Mohd Fadzil Hassan, and Norshuhani Zamin. "Rule-based pattern extractor and named entity recognition: A hybrid approach." In: *2010 International Symposium on Information Technology*. Vol. 2. 2010, pp. 563–568. DOI: 10.1109/ITSIM.2010.5561392.

[319] Siti Syakirah Sazali, Nurazzah Abdul Rahman, and Zainab Abu Bakar. "Information extraction: Evaluating named entity recognition from classical Malay documents." In: *2016 Third International Conference on Information Retrieval and Knowledge Management (CAMP)*. 2016, pp. 48–53. DOI: 10.1109/INFRKM.2016.7806333.

[320] Simon Scheider, Krzysztof Janowicz, and Werner Kuhn. "Grounding Geographic Categories in the Meaningful Environment." In: *Spatial Information Theory*. Ed. by Kathleen Stewart Hornsby, Christophe Claramunt, Michel Denis, and Gérard Ligozat. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 69–87. ISBN: 978-3-642-03832-7.

[321] Simon Scheider and Werner Kuhn. "Affordance-based categorization of road network data using a grounded theory of channel networks." In: *International Journal of Geographical Information Science* 24.8 (2010), pp. 1249–1267. DOI: 10.1080/13658810903514198. eprint: https://doi.org/10.1080/13658810903514198. URL: https://doi.org/10.1080/13658810903514198.

[322] Anne Schleicher and Diedrich Wolter. "Concept-Level Reasoning for Resolving Reference Ambiguities for Natural Language Question Understanding." In: *Proceedings of AGILE Workshop "Teaching Geospatial Technologies to All"*. 2018.

[323] Christoph Schlieder. "Reasoning About Ordering." In: *Spatial Information Theory: A Theoretical Basis for GIS, International Conference COSIT '95, Semmering, Austria, September 21-23, 1995, Proceedings*. Ed. by Andrew U. Frank and Werner Kuhn. Vol. 988. Lecture Notes in Computer Science. Springer, 1995, pp. 341–349. DOI: 10.1007/3-540-60392-1_22. URL: https://doi.org/10.1007/3-540-60392-1_22.

[324] Steven Schockaert. "Vague regions in Geographic Information Retrieval." In: *ACM SIGSPATIAL Special* 3.2 (2011), pp. 24–28. DOI: 10.1145/2047296.2047302. URL: https://doi.org/10.1145/2047296.2047302.

[325] Hinrich Schütze, Christopher D Manning, and Prabhakar Raghavan. *Introduction to information retrieval*. Vol. 39. Cambridge University Press Cambridge, 2008.

[326] Tobias Schwartz. *Design and Analysis of an Architecture for Understanding Geographical Texts*. 2020.

[327] Angela Schwering and Martin Raubal. "Spatial Relations for Semantic Similarity Measurement." In: *Proceedings of the 24th International Conference on Perspectives in Conceptual Modeling*. ER'05. Klagenfurt, Austria: Springer-Verlag, 2005, pp. 259–269. ISBN: 3-540-29395-7, 978-3-540-29395-8. DOI: 10.1007/11568346_28.

[328] D. Seamon and J. Sowers. "Place and Placelessness, Edward Relph." In: 2008.

[329] Vivek Sehgal, Lise Getoor, and Peter D Viechnicki. "Entity Resolution in Geospatial Data Integration." In: *Proceedings of the 14th Annual ACM International Symposium on Advances in Geographic Information Systems*. GIS '06. Arlington, Virginia, USA: Association for Computing Machinery, 2006, 83–90. ISBN: 1595935290. DOI: 10.1145/1183471.1183486. URL: https://doi.org/10.1145/1183471.1183486.

[330] Vassilis Serafeimidis. "The institutional dimensions of information systems evaluation." In: *Information technology evaluation methods and management*. IGI Global, 2001, pp. 99–110.

[331] Khaled F. Shaalan. "Rule-based Approach in Arabic Natural Language Processing." In: 2010.

[332] Sara Shatford. "Analyzing the Subject of a aricture: A Theoretical Approach." In: *Cataloging & Classification Quarterly* 6.3 (1986), pp. 39–62. DOI: 10.1300/J104v06n03_04. eprint: https://doi.org/10.1300/J104v06n03_04.

[333] Hyeong Jin Shin, Jeong Yeon Park, Dae Bum Yuk, and Jae Sung Lee. "BERT-based Spatial Information Extraction." In: *Proceedings of the Third International Workshop on Spatial Language Understanding*. Online: Association for Computational Linguistics, Nov. 2020, pp. 10–17. DOI: 10.18653/v1/2020.splu-1.2. URL: https://aclanthology.org/2020.splu-1.2.

[334] Yusuke Shinyama and Satoshi Sekine. "Preemptive Information Extraction using Unrestricted Relation Discovery." In: *Proceedings of the Human Language Technology Conference of the NAACL, Main Conference*. New York City, USA: Association for Computational Linguistics, June 2006, pp. 304–311. URL: https://aclanthology.org/N06-1039.

[335] Oyesh Singh, Ankur Padia, and Anupam Joshi. *Named Entity Recognition for Nepali Language*. Aug. 2019.

[336] Umrinderpal Singh, Vishal Goyal, and Gurpreet Singh Lehal. "Named Entity Recognition System for Urdu." In: *COLING*. 2012.

[337] Alisa Smirnova and Philippe Cudré-Mauroux. "Relation Extraction Using Distant Supervision: A Survey." In: *ACM Comput. Surv.* 51.5 (2018). ISSN: 0360-0300. DOI: 10.1145/3241741. URL: https://doi.org/10.1145/3241741.

[338] Barry Smith and David M. Mark. "Ontology and Geographic Kinds." In: *Proceedings of the 8th International Symposium on Spatial Data Handling*. 1998.

[339] David A. Smith and Gregory R. Crane. "Disambiguating Geographic Names in a Historical Digital Library." In: *Research and Advanced Technology for Digital Libraries, 5th European Conference, ECDL 2001, Darmstadt, Germany, September 4-9, 2001, Proceedings*. Ed. by Panos Constantopoulos and Ingeborg Sølvberg. Vol. 2163. Lecture Notes in Computer Science. Springer, 2001, pp. 127–136. DOI: 10.1007/3-540-44796-2_12. URL: https://doi.org/10.1007/3-540-44796-2_12.

[340] David A. Smith and Gregory Crane. "Disambiguating Geographic Names in a Historical Digital Library." In: *Proceedings of the 5th European Conference on Research and Advanced Technology for Digital Libraries*. ECDL '01. Berlin, Heidelberg: Springer-Verlag, 2001, 127–136. ISBN: 3540425373.

[341] David A. Smith and Gideon S. Mann. "Bootstrapping Toponym Classifiers." In: *Proceedings of the HLT-NAACL 2003 Workshop on Analysis of Geographic References - Volume 1*. HLT-NAACL-GEOREF '03. USA: Association for Computational Linguistics, 2003, 45–49. DOI: 10.3115/1119394.1119401. URL: https://doi.org/10.3115/1119394.1119401.

[342] Damires Souza, Ana Carolina Salgado, and Patricia Tedesco. "Towards a Context Ontology for Geospatial Data Integration." In: *Proceedings of the 2006 International Conference on On the Move to Meaningful Internet Systems: AWeSOMe, CAMS, COMINF, IS, KSinBIT, MIOSCIAO, MONET - Volume Part II*. OTM'06. Montpellier, France: Springer-Verlag, 2006, 1576–1585. ISBN: 3540482733. DOI: 10.1007/11915072_64. URL: https://doi.org/10.1007/11915072_64.

[343] Michael Speriosu and Jason Baldridge. "Text-Driven Toponym Resolution using Indirect Supervision." In: *Proceedings of the 51st Annual Meeting of the Association for Computational Linguistics, ACL 2013, 4-9 August 2013, Sofia, Bulgaria, Volume 1: Long Papers*. The Association for Computer Linguistics, 2013, pp. 1466–1476. URL: https://aclanthology.org/P13-1144/.

[344] Michael Speriosu, Travis Brown, Taesun Moon, Jason Baldridge, and Katrin Erk. "Connecting Language and Geography with Region-Topic Models." In: vol. 620. Jan. 2010.

[345] Andreas Spitz, Johanna Geiß, and Michael Gertz. "So far away and yet so close: augmenting toponym disambiguation and similarity with text-based networks." In: *Proceedings of the Third International ACM SIGMOD Workshop on Managing and Mining Enriched Geo-Spatial Data, GeoRich@SIGMOD 2016, San Francisco, California, USA, June 26 - July 1, 2016*. Ed. by Andreas Züfle, Benjamin Adams, and Dingming Wu. ACM, 2016, 2:1–2:6. DOI: 10.1145/2948649.2948651. URL: https://doi.org/10.1145/2948649.2948651.

[346] Kristin Stock and Mark Hall. "The Role of Context in the Interpretation of Natural Language Location Descriptions." In: Jan. 2018, pp. 245–254. ISBN: 978-3-319-63945-1. DOI: 10.1007/978-3-319-63946-8_40.

[347] Kristin Stock and Javid yousaf. "Context-aware automated interpretation of elaborate natural language descriptions of location through learning from empirical data." In: *International Journal of Geographical Information Science* 32.6 (2018), pp. 1087–1116. DOI: 10.1080/13658816.2018.1432861.

[348] Kristin Stock, Claudia Cialone, Robert Pasley, Ronald Ranjit, Amir Pourabdollah, and Michael Brown. "Multilingual Natural Language Spatial Querying using Natural Semantic Metalanguage." In: (May 2013).

[349] Louise T. Su. “A comprehensive and systematic model of user evaluation of Web search engines: I. Theory and background.” In: *Journal of the American Society for Information Science and Technology* 54.13 (2003), pp. 1175–1192. DOI: https://doi.org/10.1002/asi.10303. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/asi.10303. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/asi.10303.

[350] CLARE L. TWIGGER-ROSS and DAVID L. UZZELL. “PLACE AND IDENTITY PROCESSES.” In: *Journal of Environmental Psychology* 16.3 (1996), pp. 205–220. ISSN: 0272-4944. DOI: https://doi.org/10.1006/jevp.1996.0017. URL: https://www.sciencedirect.com/science/article/pii/S0272494496900171.

[351] L. Talmy. “Toward a Cognitive Semantics.” In: 2003.

[352] Leonard Talmy. “How Language Structures Space.” In: *Spatial Orientation: Theory, Research, and Application*. Ed. by Herbert L. Pick and Linda P. Acredolo. Boston, MA: Springer US, 1983, pp. 225–282. ISBN: 978-1-4615-9325-6. DOI: 10.1007/978-1-4615-9325-6_11. URL: https://doi.org/10.1007/978-1-4615-9325-6_11.

[353] Daniel Allen Tappan. *Knowledge-based spatial reasoning for automated scene generation from text descriptions*. New Mexico State University, 2004.

[354] Ann Taylor, Mitchell Marcus, and Beatrice Santorini. “The Penn treebank: an overview.” In: *Treebanks* (2003), pp. 5–22.

[355] Thora Tenbrink. *Space, Time, and the Use of Language: An Investigation of Relationships*. De Gruyter Mouton, 2008. ISBN: 9783110198829. DOI: doi:10.1515/9783110198829. URL: https://doi.org/10.1515/9783110198829.

[356] Paul Thomas and David Hawking. “Evaluation by Comparing Result Sets in Context.” In: *IN PROC. CIKM*. ACM Press, 2006, pp. 94–101.

[357] Peter M Todd and Thomas Trenholm Hills andTrevor W Robbins. “Cognitive search: evolution, algorithms, and the brain.” In: *Cambridge, MA : MIT Press* (2012). DOI: 10.7551/MITPRESS/9780262018098.001.0001.

[358] Jefferson Peña Torres, Raúl Gutierrez de Piñerez Reyes, and Víctor A. Bucheli. “Support Vector Machines for Semantic Relation Extraction in Spanish Language.” In: *Advances in Computing*. Ed. by Jairo E. Serrano C. and Juan Carlos Martínez-Santos. Cham: Springer International Publishing, 2018, pp. 326–337. ISBN: 978-3-319-98998-3.

[359] Stephen Tratz. “A Cross-Task Flexible Transition Model for Arabic Tokenization, Affix Detection, Affix Labeling, POS Tagging, and Dependency Parsing.” In: *Proceedings of the Fourth Workshop on Statistical Parsing of Morphologically-Rich Languages, SPMRL@EMNLP 2013, Seattle, Washington, USA, October 18, 2013*. Ed. by Yoav Goldberg, Yuval Marton, Ines Rehbein, and Yannick Versley. Association for Computational Linguistics, 2013, pp. 34–45. URL: https://aclanthology.org/W13-4904/.

[360] Ladina Tschander, Hedda Rahel Schmidtke, Carola Eschenbach, Christopher Habel, and Lars Kulik. “A Geometric Agent Following Route Instructions.” In: *Spatial Cognition*. Vol. 2685. Lecture Notes in Computer Science. Springer, 2003, pp. 89–111.

[361] Yi-fu. Tuan. "Space and place:the perspective of experience." In: *Minneapolis :University of Minnesota Press* (1977).

[362] David Tulloch and Earl Epstein. "Benefits of Community MPLIS: Efficiency, Effectiveness, and Equity." In: *T. GIS* 6 (Mar. 2002), pp. 195–211. DOI: 10.1111/1467-9671.00105.

[363] Barbara Tversky. "Cognitive Maps, Cognitive Collages, and Spatial Mental Models." In: *Spatial Information Theory: A Theoretical Basis for GIS, International Conference COSIT '93, Marciana Marina, Elba Island, Italy, September 19-22, 1993, Proceedings*. Ed. by Andrew U. Frank and Irene Campari. Vol. 716. Lecture Notes in Computer Science. Springer, 1993, pp. 14–24. DOI: 10.1007/3-540-57207-4_2. URL: https://doi.org/10.1007/3-540-57207-4_2.

[364] Olga Uryupina. "Semi-Supervised Learning of Geographical Gazetteers from the Internet." In: *Proceedings of the HLT-NAACL 2003 Workshop on Analysis of Geographic References - Volume 1*. HLT-NAACL-GEOREF '03. USA: Association for Computational Linguistics, 2003, 18–25. DOI: 10.3115/1119394.1119397. URL: https://doi.org/10.3115/1119394.1119397.

[365] Pertti Vakkari. "Evaluating Interactive Information Retrieval Systems." In: *Prisma. com* 19 (2012), pp. 55–69.

[366] Pertti Vakkari and Saila Huuskonen. "Search effort degrades search output but improves task outcome." In: *Journal of the American Society for Information Science and Technology* 63.4 (2012), pp. 657–670. DOI: https://doi.org/10.1002/asi.21683. eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/asi.21683. URL: https://onlinelibrary.wiley.com/doi/abs/10.1002/asi.21683.

[367] M. Vasardani, Lesley Stirling, and S. Winter. "The preposition at from a spatial language, cognition, and information systems perspective." In: *Semantics and Pragmatics* 10 (2017), p. 3.

[368] Maria Vasardani, Martin Tomko, and Stephan Winter. "The Cognitive Aspect of Place Properties." In: *International Conference on GIScience Short Paper Proceedings* 1 (Jan. 2016). DOI: 10.21433/B3116928W12F.

[369] Maria Vasardani and Stephan Winter. "Place Properties." In: *Advancing Geographic Information Science*. Ed. by H. J. Onsrud and W. Kuhn. GSDI Press, 2016, pp. 243–254.

[370] Maria Vasardani, Stephan Winter, and Kai-Florian Richter. "Locating place names from place descriptions." In: *International Journal of Geographical Information Science* 27.12 (2013), pp. 2509–2532.

[371] Maria Vasardani, Stephan Winter, and Kai-Florian Richter. "A Graph Database Model for Knowledge Extracted from Place Descriptions." In: *ISPRS International Journal of Geographical Information Science* 221.7 (2018).

[372] Maria Vasardani, Stephan Winter, Kai-Florian Richter, Lesley Stirling, and Daniela Richter. "Spatial interpretations of preposition "at"." In: *First ACM SIGSPATIAL International Workshop on Crowdsourced and Volunteered Geographic Information 2012, GEOCROWD 2012, Redondo Beach, CA, USA, November 6, 2012*. Ed. by Michael F. Goodchild, Dieter Pfoser, and Daniel Z. Sui. ACM, 2012, pp. 46–53. DOI: 10.1145/2442952.2442961. URL: https://doi.org/10.1145/2442952.2442961.

[373] Maria Vasardani, Sabine Timpf, Stephan Winter, and Martin Tomko. "From Descriptions to Depictions: A Conceptual Framework." In: vol. 8116. Sept. 2013, pp. 299–319. ISBN: 978-3-319-01789-1. DOI: 10.1007/978-3-319-01790-7_17.

[374] Thomas J. Vögele, Christoph Schlieder, and Ubbo Visser. "Intuitive Modelling of Place Name Regions for Spatial Information Retrieval." In: *Spatial Information Theory. Foundations of Geographic Information Science, International Conference, COSIT 2003, Ittingen, Switzerland, September 24-28, 2003, Proceedings*. Ed. by Werner Kuhn, Michael F. Worboys, and Sabine Timpf. Vol. 2825. Lecture Notes in Computer Science. Springer, 2003, pp. 239–252. DOI: 10.1007/978-3-540-39923-0_16. URL: https://doi.org/10.1007/978-3-540-39923-0_16.

[375] Jan Oliver Wallgrün, Alexander Klippel, and Morteza Karimzadeh. "Towards Contextualized Models of Spatial Relations." In: *Proceedings of Geographic Information Retrieval (GIR)*. 2015.

[376] Jan Oliver Wallgrün, Morteza Karimzadeh, Alan M. MacEachren, and Scott Pezanowski. "GeoCorpora: building a corpus to test and train microblog geoparsers." In: *International Journal of Geographical Information Science* 32.1 (2018), pp. 1–29. DOI: 10.1080/13658816.2017.1368523. eprint: https://doi.org/10.1080/13658816.2017.1368523. URL: https://doi.org/10.1080/13658816.2017.1368523.

[377] F. Wang. "Handling Grammatical Errors, Ambiguity and Impreciseness in GIS Natural Language Queries." In: *Trans. GIS* 7.1 (2003), pp. 103–121. DOI: 10.1111/1467-9671.00132. URL: https://doi.org/10.1111/1467-9671.00132.

[378] Fangju Wang. "A fuzzy grammar and possibility theory-based natural language user interface for spatial queries." In: *Fuzzy Sets Syst.* 113.1 (2000), pp. 147–159. DOI: 10.1016/S0165-0114(99)00018-4. URL: https://doi.org/10.1016/S0165-0114(99)00018-4.

[379] Hongmei Wang, Alan MacEachren, and Guoray Cai. "Design of Human-GIS Dialogue for Communication of Vague Spatial Concepts Based on Human Communication Framework." In: (Jan. 2004).

[380] Jing Wang, Yongchuan Yu, Jianzhuo Yan, Jianhui Chen, Zhongcheng Zhao, and Dongsheng Wang. "A Probabilistic Method for Linking BI Provenances to Open Knowledge Base." In: *BIH*. 2016.

[381] Kaiqiang Wang and Yijie Shi. "User information extraction in big data environment." In: *2017 3rd IEEE International Conference on Computer and Communications (ICCC)*. 2017, pp. 2315–2318. DOI: 10.1109/CompComm.2017.8322948.

[382] Pengwei Wang, Tianyong Hao, Jun Yan, and Lianwen Jin. "Large-Scale Extraction of Drug-Disease Pairs from the Medical Literature." In: *J. Assoc. Inf. Sci. Technol.* 68.11 (2017), 2649–2661. ISSN: 2330-1635. DOI: 10.1002/asi.23876. URL: https://doi.org/10.1002/asi.23876.

[383] Yichuan Wang, LeeAnn Kung, and Terry Anthony Byrd. "Big data analytics: Understanding its capabilities and potential benefits for healthcare organizations." In: *Technological Forecasting and Social Change* 126 (2018), pp. 3–13. ISSN: 0040-1625. DOI: https://doi.org/10.1016/j.techfore.2015.12.019. URL: https://www.sciencedirect.com/science/article/pii/S0040162516000500.

[384] William H. Warren, Daniel B. Rothman, Benjamin H. Schnapp, and Jonathan D. Ericson. “Wormholes in Virtual Space: From Cognitive Maps to Cognitive Graphs.” In: *Cognition* 166 (2017), pp. 152–163. DOI: 10.1016/j.cognition.2017.05.020.

[385] Demi van Weerdenburg, Simon Scheider, Benjamin Adams, Bas Spierings, and Egbert van der Zee. “Where to go and what to do: Extracting leisure activity potentials from Web data on urban space.” In: *Comput. Environ. Urban Syst.* 73 (2019), pp. 143–156. DOI: 10.1016/j.compenvurbsys.2018.09.005. URL: https://doi.org/10.1016/j.compenvurbsys.2018.09.005.

[386] Ryen W White. *Interactions with search systems*. Cambridge University Press, 2016.

[387] John Wieczorek, Qinghua Guo, and Robert Hijmans. “The point-radius method for georeferencing locality descriptions and calculating associated uncertainty.” In: *Int. J. Geogr. Inf. Sci.* 18.8 (2004), pp. 745–767. DOI: 10.1080/13658810412331280211. URL: https://doi.org/10.1080/13658810412331280211.

[388] Lahbib Wiem, Ibrahim Bounhas, Bilel Elayeb, and Yahya Slimani. “A hybrid approach for arabic semantic relation extraction.” In: Jan. 2013.

[389] Brian C. Williams and Johan de Kleer. “Qualitative Reasoning About Physical Systems: a Return to Roots.” In: 51.1-3 (1991), pp. 1–9.

[390] Max L. Wilson. “Search User Interface Design.” In: *Search User Interface Design*. 2011.

[391] Benjamin Wing and Jason Baldridge. “Simple supervised document geolocation with geodesic grids.” In: *The 49th Annual Meeting of the Association for Computational Linguistics: Human Language Technologies, Proceedings of the Conference, 19-24 June, 2011, Portland, Oregon, USA*. Ed. by Dekang Lin, Yuji Matsumoto, and Rada Mihalcea. The Association for Computer Linguistics, 2011, pp. 955–964. URL: https://aclanthology.org/P11-1096/.

[392] Benjamin Wing and Jason Baldridge. “Hierarchical Discriminative Classification for Text-Based Geolocation.” In: *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing, EMNLP 2014, October 25-29, 2014, Doha, Qatar, A meeting of SIGDAT, a Special Interest Group of the ACL*. Ed. by Alessandro Moschitti, Bo Pang, and Walter Daelemans. ACL, 2014, pp. 336–348. DOI: 10.3115/v1/d14-1039. URL: https://doi.org/10.3115/v1/d14-1039.

[393] S. Winter, R.M. Bennett, and ... [et al.] “Spatially enabling place information.” English. In: *GSDI 12 world conference proceedings : realising spatially enabled societies, Singapore, Japan, 19-22 October 2010. 14 p.* 2010.

[394] Stephan Winter and Christian Freksa. “Approaching the notion of place by contrast.” In: *J. Spatial Information Science* 5.1 (2012), pp. 31–50. DOI: 10.5311/JOSIS.2012.5.90.

[395] Stephan Winter and M. Truelove. “Talking about place where it matters.” In: *Las Navas 20th anniversary meeting on cognitive and linguistic aspects of geographic space, Las Navas, Spain*. Ed. by Andrew Frank and David Mark. Springer, 2010, pp. 121–139.

[396] Stephan Winter and Yunhui Wu. "Intelligent Spatial Communication." In: *Research Trends in Geographic Information Science*. Ed. by Gerhard Navratil. Berlin: Springer, 2009, pp. 235–250.

[397] Diedrich Wolter. "SparQ - A Spatial Reasoning Toolbox." In: *Benchmarking of Qualitative Spatial and Temporal Reasoning Systems, Papers from the 2009 AAAI Spring Symposium, Technical Report SS-09-02, Stanford, California, USA, March 23-25, 2009*. AAAI, 2009, p. 53. URL: http://www.aaai.org/Library/Symposia/Spring/2009/ss09-02-012.php.

[398] Diedrich Wolter and Madiha Yousaf. "Context and Vagueness in Automated Interpretation of Place Description: A Computational Model." In: *Proceedings of Workshops and Posters at the 13th International Conference on Spatial Information Theory, COSIT 2017, September 4-8, 2017, L'Aquila, Italy*. 2017, pp. 137–142. DOI: 10.1007/978-3-319-63946-8_27.

[399] Prof.Dr.Diedrich Wolter. "Context-Sensitive Qualitative Spatial Reasoning for Interpreting Vague Place Descriptions." In: (2016-2020). DOI: 10.1177/1461444814523872. URL: https://www.uni-bamberg.de/en/sme/research/qsr4vgi/.

[400] Michael F. Worboys, Matt Duckham, and Michael F. Goodchild, eds. *Foundations of Geographic Information Science*. Taylor & Francis, 2003. ISBN: 978-0-415-30726-0. DOI: 10.1201/9780203009543. URL: https://doi.org/10.1201/9780203009543.

[401] Fei Wu and Daniel S. Weld. "Open Information Extraction Using Wikipedia." In: *ACL 2010, Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics, July 11-16, 2010, Uppsala, Sweden*. Ed. by Jan Hajic, Sandra Carberry, and Stephen Clark. The Association for Computer Linguistics, 2010, pp. 118–127. URL: https://aclanthology.org/P10-1013/.

[402] Clarissa Castellã Xavier, Vera Lúcia Strube de Lima, and Marlo Souza. "Open Information Extraction Based on Lexical-Syntactic Patterns." In: *Brazilian Conference on Intelligent Systems, BRACIS 2013, Fortaleza, CE, Brazil, 19-24 October, 2013*. IEEE Computer Society, 2013, pp. 189–194. DOI: 10.1109/BRACIS.2013.39. URL: https://doi.org/10.1109/BRACIS.2013.39.

[403] Xiaobai Yao and Jean-Claude Thill. "How Far Is Too Far? - A Statistical Approach to Context-contingent Proximity Modeling." In: *Trans. GIS* 9.2 (2005), pp. 157–178.

[404] Madiha Yousaf and Diedrich Wolter. "Core Computations for Supporting Question Answering by Spatio-Ontological Reason- ing." In: *Core Computations on Spatial Infor- mation of GIS*. 2018.

[405] Madiha Yousaf and Diedrich Wolter. "Spatial Information Extraction from Text Using Spatio-Ontological Reasoning (Short Paper)." In: *10th International Conference on Geographic Information Science (GIScience 2018)*. Ed. by Stephan Winter, Amy Griffin, and Monika Sester. Vol. 114. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum für Informatik, 2018, 71:1–71:6.

[406] Madiha Yousaf and Diedrich Wolter. "How to Identify Appropriate Key-Value Pairs for Querying OSM." In: *Proceedings of the 13th Workshop on Geographic Information Retrieval*. GIR '19. Lyon, France: Association for Computing Machinery, 2019. ISBN: 9781450372602. DOI: 10.1145/3371140.3371147. URL: https://doi.org/10.1145/3371140.3371147.

[407] Madiha Yousaf and Diedrich Wolter. "A reasoning model for georeferencing named and unnamed spatial entities in natural language place descriptions." In: *Spatial Cognition & Computation* 0.0 (2021), pp. 1–39. DOI: 10.1080/13875868.2021.2002872. eprint: https://doi.org/10.1080/13875868.2021.2002872. URL: https://doi.org/10.1080/13875868.2021.2002872.

[408] Yecheng Yuan. "Extracting spatial relations from document for geographic information retrieval." In: *2011 19th International Conference on Geoinformatics*. 2011, pp. 1–5. DOI: 10.1109/GeoInformatics.2011.5980797.

[409] Paul A. Zandbergen. "A comparison of address point, parcel and street geocoding techniques." In: *Computers, Environment and Urban Systems* 32.3 (2008). Discrete Global Grids, pp. 214–232. ISSN: 0198-9715. DOI: https://doi.org/10.1016/j.compenvurbsys.2007.11.006. URL: https://www.sciencedirect.com/science/article/pii/S0198971507000890.

[410] Dmitry Zelenko, Chinatsu Aone, and Anthony Richardella. "Kernel Methods for Relation Extraction." In: *J. Mach. Learn. Res.* 3.null (2003), 1083–1106. ISSN: 1532-4435.

[411] Guanming Zeng, Chuang Zhang, Bo Xiao, and Zhiqing Lin. "CRFs-Based Chinese Named Entity Recognition with Improved Tag Set." In: Jan. 2009, pp. 519–522. DOI: 10.1109/CSIE.2009.551.

[412] Chuanrong Zhang, Tian Zhao, and Weidong Li. "Automatic search of geospatial features for disaster and emergency management." In: *Int. J. Appl. Earth Obs. Geoinformation* 12.6 (2010), pp. 409–418. DOI: 10.1016/j.jag.2010.05.004. URL: https://doi.org/10.1016/j.jag.2010.05.004.

[413] Shaodian Zhang and Noémie Elhadad. "Unsupervised biomedical named entity recognition: Experiments with clinical and biological texts." In: *Journal of biomedical informatics* 46.6 (2013), pp. 1088–1098.

[414] Xiao Zhang, Baojun Qiu, Prasenjit Mitra, Sen Xu, Alexander Klippel, and Alan M. MacEachren. "Disambiguating Road Names in Text Route Descriptions using Exact-All-Hop Shortest Path Algorithm." In: *ECAI 2012 - 20th European Conference on Artificial Intelligence. Including Prestigious Applications of Artificial Intelligence (PAIS-2012) System Demonstrations Track, Montpellier, France, August 27-31 , 2012*. Ed. by Luc De Raedt, Christian Bessiere, Didier Dubois, Patrick Doherty, Paolo Frasconi, Fredrik Heintz, and Peter J. F. Lucas. Vol. 242. Frontiers in Artificial Intelligence and Applications. IOS Press, 2012, pp. 876–881. DOI: 10.3233/978-1-61499-098-7-876. URL: https://doi.org/10.3233/978-1-61499-098-7-876.

[415] Yanfei Zhong, Wu Siqi, and Bei Zhao. "Scene Semantic Understanding Based on the Spatial Context Relations of Multiple Objects." In: *Remote Sensing* 9 (Oct. 2017), p. 1030. DOI: 10.3390/rs9101030.

[416] Liguo Zhou, Xuezhi Feng, Jiangfeng She, and Shunping Xie. "An approach for natural language understanding in GIS based on ontology." In: *Geoinformatics 2007: Geospatial Information Science*. Ed. by Jingming Chen and Yingxia Pu. Vol. 6753. International Society for Optics and Photonics. SPIE, 2007, pp. 357 –367. DOI: 10.1117/12.761760. URL: https://doi.org/10.1117/12.761760.

[417] Fei Zhu and Bairong Shen. "Combined SVM-CRFs for biological named entity recognition with maximal bidirectional squeezing." In: *PloS one* 7.6 (2012), e39230.

[418] Kai Zimmermann and Christian Freksa. "Qualitative Spatial Reasoning Using Orientation, Distance, and Path Knowledge." In: *Appl. Intell.* 6.1 (1996), pp. 49–58. DOI: 10.1007/BF00117601. URL: https://doi.org/10.1007/BF00117601.

[419] J. Zlatev. "Spatial Semantics." In: *The Oxford Handbook of Cognitive Linguistics* (Jan. 2012). DOI: 10.1093/oxfordhb/9780199738632.013.0013.

[420] J. Zwarts. "Prepositional Aspect and the Algebra of Paths." In: *Linguistics and Philosophy* 28 (Dec. 2005), pp. 739–779. DOI: 10.1007/s10988-005-2466-y.

[421] Joost Zwarts and Yoad Winter. "Vector space semantics: A model-theoretic analysis of locative prepositions." In: *Journal of logic, language and information* 9.2 (2000), pp. 169–211.

[422] Shairoz sohail. "Geospatial Natural Language processing." In: (2020). URL: https://medium.com/geoai/geospatial-natural-language-processing-ee3fc6ea6939.

[423] Geospatial world. "In focus: how mature is your GIS?" In: (2013). URL: https://www.geospatialworld.net/article/in-focus-how-mature-is-your-gis.