University Library

A gateway to Melbourne's research publications

Minerva Access is the Institutional Repository of The University of Melbourne

Author/s:
Hamzei, Ehsan

Title:
Place-related question answering: From questions to relevant answers

Date:
2021

Persistent Link:
https://hdl.handle.net/11343/273216

Terms and Conditions:
Terms and Conditions: Copyright in works deposited in Minerva Access is retained by the copyright owner. The work may not be altered without permission from the copyright owner. Readers may only download, print and save electronic copies of whole works for their own personal non-commercial use. Any use that exceeds these limits requires permission from the copyright owner. Attribution is essential when quoting or paraphrasing from these works.

PLACE-RELATED QUESTION ANSWERING: FROM QUESTIONS TO RELEVANT ANSWERS

EHSAN HAMZEI

A thesis submitted in total fulfillment of the requirements for the degree of Doctor of Philosophy

Feb 2021
Department of Infrastructure Engineering
The University of Melbourne

Ehsan Hamzei: *Place-related Question Answering: From Questions to Relevant Answers,* © Feb 2021

There are no shortcuts to any place worth going.

— Beverly Sills

Dedicated to the loving memory of my grandma.

1943 – 2021

ABSTRACT

In everyday communications, people talk about space by referring to places. While the common sense notion of place is understandable to humans, formalising place in a computational model remains a challenging issue. The strong context dependency, diverse metaphorical uses, indeterminacy of boundaries, and vernacular reference use are major challenges in making place knowledge digestible for computers. This research aims to utilise domain knowledge to study place-related questions and their corresponding answers, and to develop models and methods to answer the questions.

In the context of place-related question answering, this study investigates what people expect from computers to understand about places, and how these place-related questions are answered in human-generated responses. First, a place model is designed for the question answering purpose using the collective domain knowledge extracted from literature. Later, the model is used to characterize the platial information in place-related questions and their human-generated answers. In the next step, the natural language questions are translated to GeoSPARQL queries to enable the spatial analysis for answering place-related questions. Finally, templates for answering where-questions are proposed to generate relevant responses similar to human-generated answers.

The results of this study show that *domain knowledge* can be used to improve current methods of place-related question answering. Using domain knowledge, an encoding method is devised that can characterise large question answering corpora with minimal supervision. The encoding results are used to identify descriptive patterns inside the questions and answers.

In the next step, a novel approach is designed using domain knowledge and object-based conceptualization of place to translate natural language questions to GeoSPARQL queries. The novelty of the approach is mainly to (1) use domain knowledge and avoid reinventing new terms, and (2) utilise First Order Logic (FOL) statements as the intermediate representation which can be later translated not only to GeoSPARQL but any other formal query languages with minimal efforts. The method is tested using the Geospatial Gold Standard dataset, and the results show significant improvements in extracting information and translating questions to queries in comparison to the state-of-the-art approaches.

Finally, the relevance of answers to where-questions is investigated using templates of generic information (i.e., type, scale

and prominence). The results show that generic representations can be used to characterise answers in a few frequent patterns and also to study relevance of answers to the questions. Moreover, the extracted knowledge can be captured using sequence prediction methods in a machine digestible manner. The results of this study can be used to test the relevance of machine-generated responses or to generate automatic responses similar to human-generated answers.

Overall, this thesis contributes to the domain of geographic question answering with a focus on geographic places. The results of this study can be used in question answering systems to analyse and classify the questions, generate queries and formulate relevant responses. The results of this study show the importance of domain knowledge in improving the performance of existing question answering systems, and also provide useful insights about human answering behaviour.

DECLARATION

I confirm that:

1. this thesis and the work presented in it are my own, and only includes of the results of my own original contributions;
2. the main content of some chapters are revised based on published manuscripts of which I am the first author, and my contributions have been explained individually.
3. this thesis is fewer than 100,000 words excluding tables, bibliographies and appendices.

Ehsan Hamzei
Melbourne, May 1, 2021

PREFACE

This thesis presents the work that I carried out during my PhD candidature in the Department of Infrastructure Engineering, the University of Melbourne. The research described in this thesis was supervised by *Prof. Stephan Winter* and *Dr. Martin Tomko*.

Under their supervision, I conducted the research presented in all chapters of this thesis, and I am the primary author for the publications that are mainly used in its content. My specific contributions for Chapters 3, 4, 5, and 6 are approximately 70%, 60%, 70% and 60%, respectively. My tasks in these studies are summarised as developing research ideas, designing and implementing methods and writing the papers. My supervisors helped me in every stage of the research including (1) planning the overall development of the work, (2) participating in defining the research problems, (3) helping me to understand research materials necessary to tackle the problems, (4) proofreading and paper polishing, and (5) technical support.

Most of the ideas and figures in Chapters 3, 4 and 6 of this thesis are derived from my publications (Hamzei, Li, Vasardani, Baldwin, Winter, and Tomko, 2020; Hamzei, Winter, and Tomko, 2019, 2020, 2021). Section 5.2.4 in Chapter 5 is an extended version of my original contributions to a conference paper that I coauthored (Xu, Hamzei, Nyamsuren, Kruiger, Winter, Tomko, and Scheider, 2020). Moreover, Section 5.2.5 of Chapter 5 and Section 4.4 of Chapter 3 are based on my contributions to two papers under review (Kuhn, Hamzei, Tomko, Vasardani, Winter, and Li, 2021, Under Review; Li, Hamzei, Majic, Hua, Renz, Tomko, Vasardani, Winter, and Baldwin, 2021, Under Review). The detailed information of my publications during this PhD study is provided below (only the aforementioned publications are used in this thesis):

Hamzei, Ehsan, Hao Chen, Maria Vasardani, Martin Tomko, and Stephan Winter (2018). "Deriving place graphs from spatial databases." In: *Proceedings of the 5th Annual Academic Research Stream at the Locate Conference, Research@Locate 2018, Adelaide, Australia, April 9-11, 2018*. Ed. by Stefan Peters and Kourosh Khoshelham. Vol. 2087. CEUR Workshop Proceedings. CEUR-WS.org, pp. 25–32.

Hamzei, Ehsan, Haonan Li, Maria Vasardani, Timothy Baldwin, Stephan Winter, and Martin Tomko (2020). "Place Questions and Human-Generated Answers: A Data Analysis Approach." In: *Geospatial Technologies for Local and Regional Development*. Ed. by Phaedon Kyriakidis, Diofantos Hadjimitsis, Dimitrios Skarlatos,

and Ali Mansourian. Cham: Springer International Publishing, pp. 3–19. ISBN: 978-3-030-14745-7. DOI: https://doi.org/10.1007/978-3-030-14745-7_1.

Hamzei, Ehsan, Stephan Winter, and Martin Tomko (2019). "Initial Analysis of Simple Where-Questions and Human-Generated Answers (Short Paper)." In: *14th International Conference on Spatial Information Theory (COSIT 2019)*. Ed. by Sabine Timpf, Christoph Schlieder, Markus Kattenbeck, Bernd Ludwig, and Kathleen Stewart. Vol. 142. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 12:1–12:8. ISBN: 978-3-95977-115-3. DOI: 10.4230/LIPIcs.COSIT.2019.12.

— (2020). "Place facets: a systematic literature review." In: *Spatial Cognition & Computation* 20.1, pp. 33–81. DOI: 10.1080/13875868.2019.1688332.

— (2021). "Templates of generic geographic information for answering where-questions." In: *International Journal of Geographical Information Science*, pp. 1–27. DOI: 10.1080/13658816.2020.1869977.

Winter, Stephan, Ehsan Hamzei, Nico Van de Weghe, and Kristien Ooms (2018). "A Graph Representation for Verbal Indoor Route Descriptions." In: *Spatial Cognition XI*. Cham: Springer International Publishing, pp. 77–91. ISBN: 978-3-319-96385-3.

Xu, Haiqi., Ehsan. Hamzei, Enkhbold. Nyamsuren, Han. Kruiger, Stephan. Winter, Martin. Tomko, and Simon. Scheider (2020). "Extracting interrogative intents and concepts from geo-analytic questions." In: *AGILE: GIScience Series* 1, pp. 1 –23. DOI: 10.5194/agile-giss-1-23-2020.

The under-review publications that I coauthored are listed below:

Kuhn, Werner, Ehsan Hamzei, Martin Tomko, Maria Vasardani, Stephan Winter, and Haonan Li (2021, Under Review). "Neural Geospatial Question Answering." In:

Li, Haonan, Ehsan Hamzei, Ivan Majic, Hua Hua, Jochen Renz, Martin Tomko, Maria Vasardani, Stephan Winter, and Timothy Baldwin (2021, Under Review). "Neural Geospatial Question Answering." In:

When I started counting my blessings,
my whole life turned around.

— Willie Nelson

ACKNOWLEDGMENTS

Looking back to 2017, when I decided to study abroad, I had a naive idea about the PhD journey. During these years, I learned a lot, I failed and I succeeded. I experienced the PhD roller coaster, the ups and downs that I believe many PhD students have experienced before me. However, without my seniors to guide me, my friends to help me and my far-away family to support and cherish me, it would have been impossible to finish this journey.

First and foremost, I would like to express my deepest gratitude to my supervisors, Prof. *Stephan Winter* and Dr. *Martin Tomko* who offer me the chance to learn from them and work with them. Without their help, guidance, patience, and encouragements, the learning process and research would have been a frustrating experience for me. They provide critical feedback for my numerous research proposals, and they always encourage me in research during my PhD study. I sincerely like to thank them for their availability, support, time and efforts during the last three and a half years.

Also, I want to thank the Australian Research Council for funding my PhD research as a part of a discovery research project, *making human place knowledge digestible for computers* (DP170100109). I like to thank the senior members of the ARC project team, Prof. *Werner Kuhn*, Dr. *Maria Vasardani*, Prof. *Tim Baldwin*, Prof. *Jochen Renz* who provided feedback to my work during our annual meetings. I am grateful for the collaborations with the RHD members of the team, *Haonan Li*, *Hua Hua* and *Hao Chen*, who were actively participated and shared their viewpoints in our heckfest events.

I would like to thank my friends in the Spatial Cognitive group, at University of Melbourne, specially *Masoud*, *Kamal*, *Kimia*, *Ivan*, *David*, *Surabhi*, *Hao*, *Wenyan*, *Xiaonan* and *Kristen* for their true friendship. Studying at the University of Melbourne was not only about research for me, and I found my best friends here, *Masoud*, *Aref*, *Shima*, *Mohammad Masoud*, *Shadi*, and *Mostafa*.

I am grateful to my parents and siblings for their unconditional support throughout my life. They always encourage me to think big, to be diligent and to enjoy life, and no matter what was the result of my works, they were always proud of me. And finally, to the most adorable angels of my life, my niece and my nephew, *Bahar* and *Pouyan*. Thank you all.

CONTENTS

1 INTRODUCTION 1
1.1 Background and Motivation 1
1.2 Towards Place-Related Question Answering . . . 3
1.3 Research Objectives and Hypothesis 5
1.4 Contributions and Significance 7
1.4.1 Place Modelling 7
1.4.2 Question Analysis 8
1.4.3 Question Formalisation 9
1.4.4 Generating Relevant Answers 9
1.4.5 Contribution Summary 9
1.5 Thesis Structure 10
2 QUESTION ANSWERING IN THE GEOGRAPHIC DOMAIN 13
2.1 Question Answering 13
2.1.1 Human Question Answering 14
2.1.2 Automatic Question Answering Systems . 16
2.1.3 Human QA versus Automatic QA 20
2.2 Geographic Question Answering 21
2.2.1 Geographic Information Retrieval 21
2.2.2 Analysing Geographic Questions 26
2.2.3 Modelling Geographic Questions 28
2.2.4 Answering Geographic Questions 29
2.3 Summary 30
3 A MULTI-FACETED MODEL OF PLACE 33
3.1 Introduction 33
3.2 Methodology for a Systematic Literature Review . 35
3.2.1 Selection Phase 36
3.2.2 Extraction Phase 38
3.2.3 Execution phase 38
3.3 Results of the Systematic Literature Review 39
3.3.1 Place-Related Research: Results of selection 40
3.3.2 Place Facets: Results of Extraction 42
3.3.3 Categories: Results of Execution 43
3.3.4 Individual Groupings 43
3.3.5 Hierarchical Clustering 45
3.4 A Multi-Faceted Model of Place 46
3.4.1 Primitive Facets 49
3.4.2 Derived Facets 57
3.4.3 Linguistic Facets 58
3.5 Discussion 60
3.6 Conclusion 61
4 UNDERSTANDING PLACE-RELATED QUESTIONS . . . 63

4.1 Introduction . . . 63
4.2 Data-driven Analysis . . . 65
4.2.1 Dataset . . . 65
4.2.2 Encoding Method . . . 66
4.3 Data-driven Results . . . 70
4.3.1 Preliminary Analysis . . . 71
4.3.2 Frequent Patterns . . . 73
4.3.3 Categorizations of Place-Related Questions/Answers . . . 75
4.3.4 Question/Answer Relationships . . . 77
4.4 Theory-driven Classification of Place-Related Questions . . . 78
4.4.1 Core Concept of Spatial Information . . . 79
4.4.2 Classification of Place-Related Questions . 80
4.5 Discussion and Conclusions . . . 84
5 FROM PLACE-RELATED QUESTIONS TO QUERIES . . 87
5.1 Introduction . . . 87
5.2 Methodology . . . 89
5.2.1 Dataset . . . 89
5.2.2 Method . . . 89
5.2.3 Encoding Extraction . . . 90
5.2.4 Constituency and Dependency Parsing . . 92
5.2.5 Concept Identification . . . 98
5.2.6 Generating First-Order Logic Statements . 98
5.2.7 Generating GeoSPARQL Queries . . . 100
5.3 Results and Discussion . . . 103
5.3.1 Extraction Results and Logical Statements 103
5.3.2 Query Generation Results . . . 105
5.4 Conclusions . . . 107
6 ANSWERING WHERE QUESTIONS . . . 109
6.1 Introduction . . . 109
6.2 Methodology . . . 112
6.2.1 Specific-Generic Translation . . . 112
6.2.2 Type, Scale and Prominence (TSP) Encoding113
6.3 Implementation . . . 115
6.3.1 Data . . . 116
6.3.2 Extraction . . . 117
6.3.3 Encoding . . . 118
6.3.4 Distribution Analysis and Rule Mining . . 118
6.3.5 Prediction . . . 119
6.4 Results . . . 121
6.4.1 Extraction and Encoding . . . 121
6.4.2 Distributions . . . 122
6.4.3 Extracted Rules . . . 124
6.4.4 Predicting the Generic Form of Answers . 128
6.5 Demonstration: From Generic to Specific . . . 132

6.6 Discussion . . . 134
6.7 Conclusions . . . 135
7 DISCUSSION AND CONCLUSIONS . . . 137
7.1 Discussions on Major Findings and Limitations . 138
7.1.1 Place Facets . . . 138
7.1.2 Question Analysis . . . 139
7.1.3 From Question to Queries . . . 141
7.1.4 Answering Questions . . . 142
7.2 Contributions of This Thesis . . . 144
7.3 Future Directions . . . 146
7.3.1 A Computational Model of Place . . . 146
7.3.2 Questions to SPATIAL SQL: A Step Towards Natural Language Interfaces to GIS . . . 147
7.3.3 Generic Geographic Information to Specific References . . . 147
7.3.4 Benchmark Dataset for Place-Related Question Answering . . . 148
7.3.5 Place Recommendation: Moving from Content to Context . . . 148

I APPENDIX 151

A APPENDIX A . . . 153
A.1 Extracted Facets . . . 153
A.2 Facet Definitions . . . 158
B APPENDIX B . . . 181
B.1 Part-of-Speech Tags . . . 181
B.2 Typed Dependencies . . . 181
C APPENDIX C . . . 185
C.1 Implementation . . . 185
C.1.1 Scale Encoding . . . 185
C.1.2 Prominence Encoding . . . 185
C.2 Comparing Simple-Where Questions and Detailed Where-Questions . . . 185
C.3 Sequence Distributions and Frequent Patterns . . 188
C.3.1 Sequence Distributions . . . 188
C.3.2 Frequent Patterns . . . 190
C.4 Demonstration Examples . . . 192
C.5 Feature Codes . . . 197

BIBLIOGRAPHY . . . 199

LIST OF FIGURES

Figure 1 A question and the response from Google Search (accessed 13/07/2020) 4
Figure 2 Logical flow of the main contributions . . 7
Figure 3 The general architecture of QA systems . 17
Figure 4 QA architecture using machine comprehension 19
Figure 5 Selection and screening process. 37
Figure 6 Process of synthesising place facets in categories. 39
Figure 7 Publication date analysis of the selected publications, by decade. 41
Figure 8 Citations of the papers selected for this literature review (log scale). 42
Figure 9 Pairwise comparison of grouping by first and second participants. 44
Figure 10 Pairwise comparison of grouping by first and third participants. 44
Figure 11 Pairwise comparison of grouping by second and third participants. 46
Figure 12 Ordered similarity matrix and its relation to the proposed categorisation. 47
Figure 13 Hierarchical categories of place facets . . 48
Figure 14 The result of hierarchical clustering of primitive facets. 49
Figure 15 Geographical distribution of resolved toponyms in questions and answers . . . 74
Figure 16 The proposed method to translate place-related question to GeoSPARQL queries. 90
Figure 17 Constituency parsing 93
Figure 18 Dependency parsing 94
Figure 19 Labelled constituency parse tree 95
Figure 20 Dependency parse tree after preprocessing 97
Figure 21 Specific-generic translation approach . . . 112
Figure 22 The proposed implementation approach . 116
Figure 23 Distribution of place types in the questions and in the answers. 122
Figure 24 Distribution of levels of scale in all toponym-based where questions and answers. . . . 122
Figure 25 Distribution of prominence levels in the questions versus answers 123

Figure 26 Distribution of types in the SWQs versus DWQs. 123
Figure 27 Distribution of the levels of scale in the SWQs versus DWQs. 124
Figure 28 Distribution of prominence levels in the SWQs versus DWQs 125
Figure 29 Histogram of importance values and boundaries of prominence levels. 186
Figure 30 Distribution of types in the SWQs versus DWQs. 186
Figure 31 Distribution of the levels of scale in the SWQs versus DWQs. 187
Figure 32 Distribution of prominence levels in the SWQs versus DWQs 188
Figure 33 Sequence distributions of the questions and answers 189
Figure 34 Top ten frequent pattern in type-sequences of the answers 191
Figure 35 Top ten frequent pattern in scale-sequences of the answers 191
Figure 36 Top ten frequent pattern in prominence-sequences of the answers . . 192

LIST OF TABLES

Table 1 Classification schema of geographic questions (Punjani et al., 2018) . . . 27
Table 2 Selected journals . . . 37
Table 3 Hierarchical clustering distance/similarity matrix . . . 40
Table 4 Number of selected papers in each iteration 40
Table 5 Type of publications. . . . 40
Table 6 Participants' groupings of the extracted facets . . . 45
Table 7 Number of answers per location question in the combined dataset . . . 66
Table 8 Semantic representation encoding . . . 68
Table 9 Evaluation of encoding extraction . . . 71
Table 10 Top-five frequent place-related semantics extracted from the dataset. Frequency in brackets. . . . 72
Table 11 Statistics of extraction and disambiguation process . . . 73
Table 12 Contingency table of ambiguous/resolved and unanswered/answered questions . . . 74
Table 13 Frequent encoding patterns in questions . 75
Table 14 Frequent encoding patterns in answers . 75
Table 15 Types of questions . . . 76
Table 16 Types of answers . . . 77
Table 17 Contingency table of linking the clusters 78
Table 18 Frequent patterns in concatenated clusters 79
Table 19 Question templates using emotive, functional and physical facets . . . 82
Table 20 Question templates using spatial facets . 83
Table 21 Question templates using linguistic facets 83
Table 22 Question templates for places in conjunction with events and fields . . . 84
Table 23 Semantic representation encoding . . . 91
Table 24 Extraction performance . . . 104
Table 25 FOL statements evaluation . . . 105
Table 26 GeoSPARQL queries evaluation . . . 106
Table 27 Sequence prediction methods . . . 120
Table 28 Extraction evaluation . . . 121
Table 29 Encoding results . . . 121
Table 30 Comparison of levels of scale in the answers compared to their questions . . . 125

Table 31 Comparison of prominence levels in the answers compared to their questions . . . 126
Table 32 Extracted rules from type sequences . . . 126
Table 33 Extracted rules from scale sequences . . . 127
Table 34 Extracted rules from prominence sequences . . . 127
Table 35 Prediction accuracy for type sequences . 130
Table 36 Prediction accuracy for scale sequences . 131
Table 37 Prediction accuracy of prominence sequences . . . 131
Table 38 Accuracy improvement using sequence prediction compared to the baselines . . . 132
Table 39 RMSE of sequence prediction methods . . 133
Table 40 SPARQL results to find specific form of the answer . . . 135
Table 41 Extracted facets from the literature 153
Table 42 Extracted facets from the literature with their definitions . . . 158
Table 43 A selected list of part-of-speech tags . . . 182
Table 44 A selected list dependency types 183
Table 45 Scale mapping from OSM values to proposed schema . . . 185
Table 46 Comparison of levels of scale in the answers compared to their questions 187
Table 47 Comparison of prominence levels in the answers compared to their questions . . . 189
Table 48 SPARQL results to Example 1 193
Table 49 SPARQL results to Example 2 194
Table 50 SPARQL results to Example 3 196
Table 51 SPARQL results to Example 4 197
Table 52 Feature codes used in Chapter 6 (extracted from Geonames documentation) 197

ACRONYMS

GIScience Geographic Information Science
HCI Human Computer Interaction
QA Question Answering
GIR Geographic Information Retrieval
MS MARCO Microsoft Machine Comprehension
NLP Natural Language Processing
NL natural language
IR Information Retrieval
GeoQA Geographic Question Answering
AI Artificial Intelligence
MC Machine Comprehension
GIS Geographic Information System
GIA Geographic Information Access
TREC Text Retrieval Conference
SLR Systematic Literature Review
SVM Support Vector Machine
FOL First Order Logic
NER Named Entity Recognition

1 INTRODUCTION

1.1 BACKGROUND AND MOTIVATION

In everyday communications, people talk about space by referring to places (Vasardani and Winter, 2016). While the common sense notion of place is understandable to humans (Cresswell, 2004; Purves, Winter, and Kuhn, 2019), formalising place in a computational model remains a challenging issue (Blaschke et al., 2018; Sui and Goodchild, 2011). The strong context dependency (Winter and Freksa, 2012) and cultural dependency of place (Ballatore, 2016), diverse metaphorical uses of place names (Agnew, 2011), indeterminacy of boundaries (Montello et al., 2003), and complex and vernacular reference use (Hollenstein and Purves, 2010) are just a few examples of the challenges in making place knowledge digestible for computers. Consequently, the current computational models for capturing human place knowledge (e.g., gazetteers) are limited to a simplistic view on the elusive concept of place (Adams, 2015).

Based on the core concepts of spatial information (Kuhn, 2012), places are conceptualised as *spatial objects* in information systems (Purves, Winter, and Kuhn, 2019). However, a computational model of place which can bridge the gap between human place knowledge and spatial information stored in computers is still an open topic in Geographic Information Science (GIScience) (Blaschke et al., 2018). Such a computational model is hypothesised to support Human Computer Interaction (HCI) and facilitate the communication between people and computers about places. Modelling place is thus a crucial step in different HCI scenarios such as Question Answering (QA) (Li and Roth, 2006), and Information Retrieval (IR) (Cheema, Zhang, and Chang, 2016; Sanderson and Kohler, 2004).

Due to the complexity of modelling places, narrowing down the scope to a specific scenario (e.g., QA) is essential to make the modelling feasible (Goodchild and Li, 2012). Hence, this PhD research confines its investigation to *place-related question answering* to study the notion of place using questions and answers about *geographic places*. Moreover, this work is focused on *English-language* questions and answers.

The static context of communication in QA compared to the dynamically established context in human-computer dialogue helps to focus on the domain (i.e., place modelling) and to skip dealing with unwanted complexity of HCI scenarios. Yet, questions and answers are basic constructs of dialogues (Varges, Weng, and

Pon-Barry, 2006) and consequently a place model for QA is expected to be applicable to responses in dialogue systems. Narrowing down the scope to geographic places allows to avoid dealing with special characteristics of virtual (e.g., *a chat room*) and mystic (e.g., *Heaven*) places. The pragmatic importance of geographic places compared to the other kinds in real-world place-related Web search and question answering satisfies this narrower scope (Hamzei et al., 2020).

Several obstacles have to be addressed to bridge the gap between human place knowledge and the computer models in the question answering scenario.

- First, a detailed analysis of *place references* is required to characterize and capture information inside place-related questions and their answers.
- Next, a *representative classification* of place-related questions is required to design a place-related question answering system.
- *Formalising questions* is another crucial step since without such formalisation computers are not able to retrieve or generate the answers.
- Finally, generating *useful responses* (i.e., in the same quality to human-generated answers) to place-related questions is a necessary step to bridge the gap.

The analysis of place references in natural language is not only a fundamental step toward place-related question answering but also can be used in different applications such as Geographic Information Retrieval (GIR) (Scheider and Purves, 2013). The classification of place-related questions is crucial to find out what people expect of computers in terms of storing and analyzing place knowledge (Henrich and Luedecke, 2007). The formalisation of place-related questions enables the usage of knowledge bases and databases in answering the questions (Punjani et al., 2018). Without such formalisation, the answering process is restricted to retrieving information from textual documents.

Analysing the human-generated answers is a step toward intelligent place-related question answering. Winter (2009) defined spatial intelligence based on the landmark paper of Turing (1950), *Computing Machinery and Intelligence*. The intelligence of a machine is defined through the machine's communication behaviour. More specifically, intelligence is achieved when people cannot differentiate whether they are communicating with a machine or a human (Turing, 1950). The key to such intelligence communication is to enable machines to imitate human answering behaviour (Winter, 2009). Consequently, analysing human-generated answers to place-related questions is the first step toward a natural and successful communication between people and computers.

Place-related question answering enables a natural communication between humans and computers about geographic places. This will directly impact different applications defined through the communication of people and machines. Web search would benefit the most through enabling place-related question answering. As Ed Parsons, Geospatial Chief Technologist of Google, indicated, "about 1 in 3 of queries that people just type into a standard Google search bar are about places"[1]. Enabling a natural communication about places in Web search would results into finding answers directly instead of searching, selecting documents, skimming and exploring paragraphs to hopefully find the answers. The application is not bound to place-related Web search, it can be used in any communication between people and machine where place plays a role. The intelligent assistants in smart phones, smart cars and smart houses are just examples of current technologies which would benefits from progress in place-related question answering.

1.2 TOWARDS PLACE-RELATED QUESTION ANSWERING

The question "*Where can I exercise during COVID-19?*" is a typical example of a place-related question aiming to find a suitable (i.e., in this case, 'safe') place for an activity. To show how current computing systems deal with such questions, let us explore the response to this question from Google Search Engine[2] (for a user in Melbourne, Australia in 2020).

Figure 1 shows the results of submitting this question to Google. The top results are Web pages about rules, regulations and advice (based on the user's location, Australia) related to the intended activity during the event. While the system makes use of the user's location to retrieve geographically relevant documents, the answer cannot be found there. Let us explore the example to find out the reasons behind the failure in this communication between the user and the system.

First, the system is unable to capture the *intent* of the question. The intent of this where-question is to find out the *location* of places which are safe for exercise during the pandemic. Consequently, if the system could recognise the intent, the response should be either a map or a place description of such places. Here, the important word in the question which is not considered is '*where*'.

Next, the system fails in applying a proper reasoning over the *place reference* – i.e., a safe place to exercise during COVID-19 near to the user. This is the result of modelling the place-related question without carefully analysing and reasoning over its content. Based on

1 http://youtu.be/ucYiMBfyNfo
2 http://www.google.com

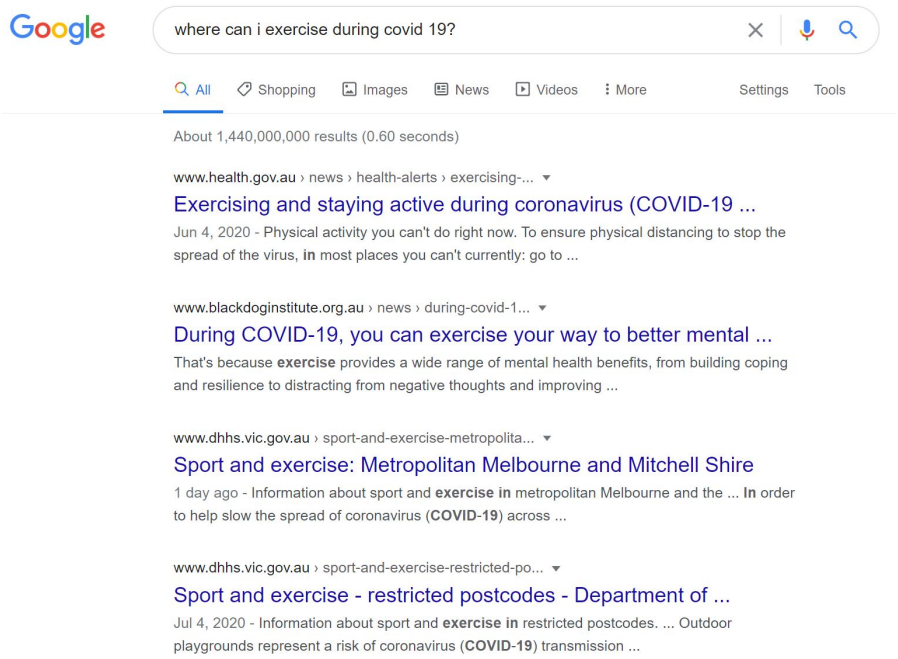

Figure 1: A question and the response from Google Search (accessed 13/07/2020)

the top-rank results, the words 'exercise' and 'COVID-19' are biasing the system to retrieve Web pages about (1) health guidance during the pandemic and (2) current situations in suburban areas. Without linking 'exercise' and 'COVID-19' to 'where', the system fails to interpret that these words are part of a reference to some geographic places.

Considering the top-rank results, the system was only able to use the indexed Web documents to provide a response to the inquirer. The structured geographic information (in the Google Map database and third party knowledge bases) is completely neglected during the answering process. This is where the lack of proper *formalisation* of the question leads to irrelevant responses. The formalised version of the question could be used to retrieve information from structured databases or semi-structured knowledge-bases.

Finally, the presentation of the response is not suitable for this question. Even if the documents were relevant to the question, the user still has to read, interpret, combine and assess the documents to find an answer. A proper answering method could save the user's time by extracting or even generating (i.e., combining relevant information units from different sources) a relevant response. For this question, such a response could be a map or a place description of relevant places.

1.3 RESEARCH OBJECTIVES AND HYPOTHESIS

The *overall hypothesis* of this thesis is that **an analysis of specific (e.g., toponyms) and generic (e.g., place types) references to geographic places can be used to characterise and imitate human answering behavior to place-related questions**. Here, the *overall objective* is **to propose models and methods suitable for characterising, formalising and answering place-related questions**. In this thesis, place-related question answering is investigated in four stages. In what follows, each stage is briefly described based on the corresponding, hypothesis, objective and research questions, and linked to the structure of the thesis.

1. **A multi-faceted model of place (Chapter 3)**: **The *hypothesis* is that a multi-faceted model of place can be derived by categorisation of extracted place properties from a systematic literature review**. Here, the objective is to collect place properties which are described in the literature and to categorise the collected properties into a multi-faceted model. This model is the foundation of the subsequent stages. In addition, it can be used in other place-related applications – i.e., not only for the question answering scenario. To test this hypothesis, the following research questions are addressed:
   (a) How can the place-related properties be extracted and systematized from an interdisciplinary review?
   (b) How can the extracted place-related properties be categorized to derive a multi-faceted model of place?

2. **Classification of place-related questions (Chapter 4)**: **The *hypothesis* of this study is that place-related questions can be classified using structural patterns derived through Natural Language Processing (NLP) and information extraction in a bottom-up approach, or can be derived from a theoretical conceptualisation of place in information system in a top-down approach**. Classifying place-related questions using both top-down and bottom-up approaches is the goal of this study. The top-down approach is theoretically grounded on the object-based conceptualisation of place in information systems (Purves, Winter, and Kuhn, 2019). Here, the multi-faceted model of place is used to derive the classification of place-related questions. In the bottom-up approach, the Microsoft Machine Comprehension (MS MARCO) dataset (Nguyen et al., 2016) is analysed to derive classification of place-related questions based on structural patterns. The following research questions are investigated to achieve the goal:

(a) How can the structural patterns from place-related questions be extracted from MS MARCO dataset?

(b) What are the frequent structural patterns in place-related questions?

(c) How can the question templates be derived using the multi-faceted model of place?

3. **Formalising place-related questions (Chapter 5)**: **The *hypothesis* is that formal query representations of place-related questions can be derived from a pipeline process of information extraction, textual parsing and template-based query generation.** Here, the objective is to translate natural language (NL) place-related questions to a formal computer-interpretable query language. The target query language is GeoSPARQL (Perry and Herring, 2012) developed by Open Geospatial Consortium[3] as an standard formal language to retrieve geospatial data from linked data knowledge bases and the Semantic Web. To test the hypothesis, the following research questions are addressed:

   (a) How can the intent of questions be identified using NL parsing methods?

   (b) How can a geographic parser be designed to translate NL place-related questions to the intermediate representation?

   (c) How can the syntactically and semantically correct GeoSPARQL queries be formulated based on the parsing results?

4. **Answering place-related questions (Chapter 6)**: This study focuses on characterising human-generated answers to place-related questions, specifically *where-questions*. **This study hypothesizes that generic information can be used to characterize human answering behavior and ultimately to generate templates for answering where-questions.** The objective is to study and imitate human answering behaviour to where-question using generic representation (e.g., place types) of geographic places. The following research questions are investigated to test this hypothesis:

   (a) How can the characterizing patterns of the human-generated answers to where-questions be derived from a large question answering corpus?

   (b) How does generic geographic information in where-questions relate to the generic information in their human-generated answers?

3 http://www.ogc.org/

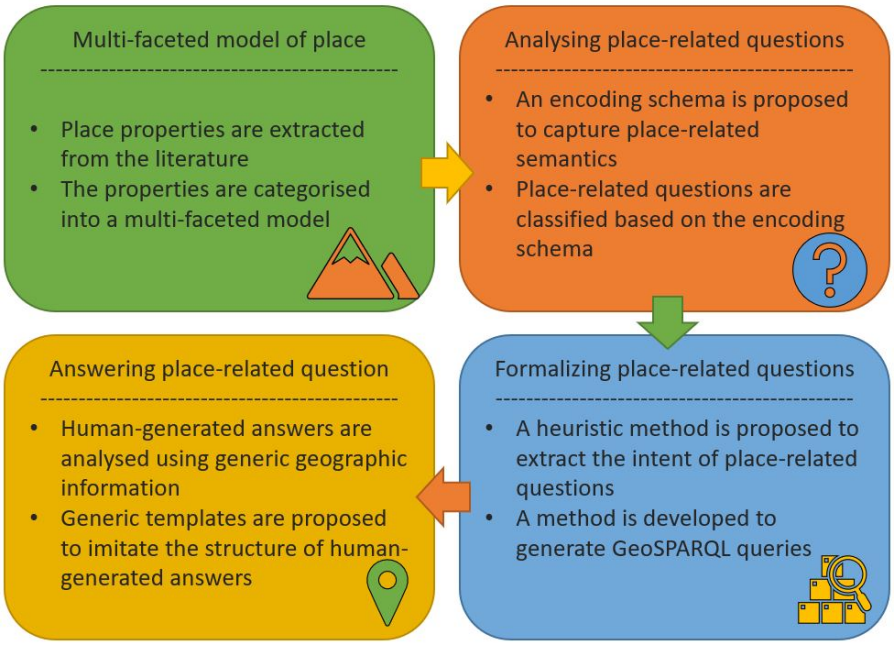

Figure 2: Logical flow of the main contributions

(c) How can the structure of human-generated responses to where-questions be imitated automatically?

All the aforementioned research questions correspond to the overall research topic – i.e., *place-related question answering*. To successfully communicate about geographic places, computers should be able to interpret and formalise the questions, and ultimately be able generate relevant responses. This thesis makes a step forward in analysing and formalising place-related question using the multi-faceted model of place. Finally, the special case of answering where-questions is thoroughly investigated to imitate human answering behaviour.

1.4 CONTRIBUTIONS AND SIGNIFICANCE

Figure 2 shows the logical flow this study and its major contributions. This thesis contributes toward different phases of place-related question answering including place modelling, analysing, formalising and finally answering place-related questions. In the following subsections, each phase and the corresponding contributions of this thesis are discussed.

1.4.1 *Place Modelling*

The first contribution is a model of place derived from the extraction and categorisation of the properties, aspects and characteristics of

the notion of place from an interdisciplinary literature review. Place is an important concept in different disciplines including *Human Geography*, *GIScience*, *Architecture* and *Environmental Psychology*. Here, a systematic approach is used to select relevant publications, extract place properties and synthesise them into place facets. The systematic approach is based on the proposed guideline by Okoli (2015) which maximises the reproducibility of the review and its results.

This contribution underpins the extraction of place-related properties, the heterogeneity resolution of diverse definitions, and the synthesis of the properties into place-related facets. While this multi-faceted model is later used to provide schema and methods for analysing and answering place-related questions, it can be used in other place-related applications including information retrieval (IR), and database and knowledge-base design.

1.4.2 *Question Analysis*

The second contribution is to extract, analyse and characterise place-related questions of MS MARCO dataset (Nguyen et al., 2016). Here, linguistic facets (i.e., one category of place model that captures references to places) are used to extract and characterise place-related questions. An encoding approach is developed to extract *place names*, *types*, *activities*, *situations*, *qualities* and *spatial relationships*. Then the encoded questions are categorised into geographical, and non-geographical questions. Finally, the geographical questions are categorised into spatial and non-spatial categories.

The results are evaluated by comparison of the encoding categorisation to categorisation based on contextualised word embedding model (ELMo) (Peters et al., 2018). The word embedding models provide a high dimensional vector representations of words. Such a model is trained over millions of documents using a deep neural network structure. However, this study shows that a simple encoding model based on domain (i.e., geographic) knowledge can derive comparable results.

Finally, using the multi-faceted model of place and core concept of spatial information (Kuhn, 2012), a top-down approach is devised to classify the place-related questions. This classification is based on the core questions that are derived from the core concepts. Based on the conceptualisation of place in information system (Purves, Winter, and Kuhn, 2019), core question for *objects* are used and enriched by place facets to derive templates for place-related questions. The results of this top-down classification is compared to the bottom-up encoding-based categorisations. This comparison shows the incompleteness of bottom-up categorisation due to the biases to the

source dataset. Hence, the theoretically-driven (top-down) schema captures more diverse types of place-related questions.

1.4.3 *Question Formalisation*

The third contribution is a method to translate NL place-related questions into GeoSPARQL queries. First, an intermediate logical model is devised to parse place-related questions. The intermediate model is a FOL representation of questions' structure. Later, the intermediate representations of questions are translated into formal query representations.

This formalisation is tested using the Geospatial Gold Standard benchmark called *GeoQuestion201* (Punjani et al., 2018). The results are compared with existing approaches. The results of this study shows the effectiveness of intermediate model of questions in translating questions to syntactically and semantically correct queries.

1.4.4 *Generating Relevant Answers*

Lastly, the fourth contribution is a method for analysing the structure of human-generated answers. Here, a detailed analysis is conducted to study human-generated answers to *where-questions*. The human-generated answers to where-questions are studied through generic representations (e.g., place types) of place names. A new approach called *Specific to Generic Translation* is devised and proposed to study the human-generated answers. This study shows that a few patterns based on generic representations can describe hundreds of where-questions and their answers. These patterns are used to enable computers to imitate human answering behaviour to where-questions. The results shows the usefulness of generic references to imitate the structure of human-generated answers.

1.4.5 *Contribution Summary*

In summary, the major contributions of this PhD study are listed below:

- A multi-faceted model of place that captures place-related properties which are extracted from multidisciplinary literature review. The model is further used in this study to show its application in place-related question answering;
- An encoding approach to characterise and categorise place-related questions;

- A top-down classification of place-related questions using the multi-faceted model of place;
- A heuristic method to extract the intent of place-related questions using constituency parsing;
- A conceptual model to parse place-related questions into an intermediate logical representation;
- A template-based method to translate parsing results into GeoSPARQL queries;
- A new approach to study and imitate human answering behaviour to where-questions using generic representations of places.

This thesis contributes to different phases of place-related question answering including modelling, encoding, formalising and also generating useful answers. This work demonstrates that extracting platial information and analysing generic references can provide useful insights for characterising and answering place-related questions. The results are justified with strong patterns extracted from MS MARCO corpus using the encoding schema and generic references. In addition, this thesis shows that using generic representations of places human answering behaviour can be imitated by computers.

1.5 THESIS STRUCTURE

The rest of this thesis is organized as follows.

Chapter 2 presents a detailed literature review of previous works in geographic question answering. This chapter covers the current models and methods for answering geographic questions. Finally, the research gaps are identified and explained. Chapter 2 establishes a foundation for its subsequent chapters.

Chapter 3 describes the systematic method to extract place-related properties from relevant multidisciplinary publications and synthesise the extracted properties into a multi-faceted model of place.

Chapter 4 investigates an encoding approach to characterise place-related questions. Then, the encoding method is used to analyse and compare different corpora of geographic questions. Finally, the top-down approach to derive templates for place-related questions is presented.

Chapter 5 first describes the intermediate representation of place-related questions. Next, the template-based approach to formulate GeoSPARQL queries from the intermediate representations is discussed. The implementation and evaluation of translating questions to GeoSPARQL queries are provided.

Chapter 6 presents a detailed analysis of human-generated answers to where-questions. The proposed Specific to Generic method is described to study and imitate human answering behaviour to where-questions. Finally, the implementation and results of generating templates for answering where-questions are discussed.

The last chapter provides an overall discussion of findings of this study, limitations of the proposed methods, and finally conclusions and future works (Chapter 7).

2

QUESTION ANSWERING IN THE GEOGRAPHIC DOMAIN

> *The answering of questions is not a simple retrieval and response of stored information; rather the process is embedded in a general structural framework containing knowledge of the questioner, the question, and the world around it.*
>
> — Norman (1973)

While *place-related question answering* is not specifically studied before, question answering in general (i. e., open-domain) and Geographic Question Answering (GeoQA) have been studied and investigated in literature. Place-related QA, as defined in this study with a pure focus on geographical places, is a branch of GeoQA. While the subject of questions in GeoQA can be other geographic concepts and phenomena such as *events* and *processes*, place-related QA is only focused on the concept of *place*. To present related works and to provide the background for the rest of this thesis, this chapter focuses on relevant scientific studies in the scope of QA. First, open-domain QA approaches are briefly reviewed to provide a background information about human QA and automatic QA. Then a detailed review of GeoQA literature is presented, and the identified research gaps and their relevance to this PhD study are discussed. Finally, the contributions of this thesis are highlighted and justified based on the literature.

2.1 QUESTION ANSWERING

The study of QA is of interest of different disciplines including linguistics, psychology and computer science. Linguists are mainly concerned about how language is syntactically, semantically and pragmatically used in communication (Guerzoni, 2003). Cognitive psychologists are focused on how people answer questions (Norman, 1973), and computer scientists devise methods and algorithms to automatically retrieve responses to satisfy inquirers' information needs (Dwivedi and Singh, 2013). In the following subsections, first previous works in linguistics and psychology to study human QA are reviewed. Next, the datasets, models and algorithms for automatic QA are discussed. Finally, a discussion of the identified gaps between human QA and automatic QA is presented.

2.1.1 *Human Question Answering*

The studies in human question answering behaviour can be classified based on the level of investigation to *syntactic*, *semantic* and *pragmatic* analysis of the questions and answers (Pomerantz, 2005; Searle, Parret, and Verschueren, 1992). Previous works on the syntax provide an understanding of the structure of questions and their corresponding answers. Using semantic analysis, the meaning of questions and the relevance of answers are investigated without considering how it might be interpreted differently in various contexts (Karttunen, 1977). Finally, pragmatic studies focused on the contextual aspects in which the meaning of questions may differ. Consequently, the validity and relevance of the corresponding answers are judged based on their relevance to the meaning of questions in the context of communication (Krifka, 1996; Malamud, 2006).

The study of syntax in human QA and dialogue is based on the following assumption: *The subsequent utterance (in QA, the answer) is based on its previous utterances (i.e., the question) and this dependency can be captured through detecting syntactic patterns* (Du Bois, 2014). Syntactic analysis in human QA resulted in (1) different taxonomies of questions using *question words* (Robinson and Rackstraw, 1972), (2) detecting *part-of-speech patterns* from utterances (Du Bois, 2014), and (3) analysing structure of questions and answers using *linguistic parsing* (Bouma, 2006).

Using the semantic level of linguistic analysis, several taxonomies are proposed to classify questions based on their subjects (e.g., McClure et al., 2002; Mooney, 1991). In addition, semantic parsing and semantic role tagging are proposed to extract useful information from the questions, and to model questions into a logical form (Delmonte, 2007; Karttunen, 1977).

In pragmatic analysis, the questions can be classified based on (1) the forms of expected answers, and (2) sources where the answers can be found (Pomerantz, 2005). The form of answer determines the characteristics of a useful answer for an inquirer in a specific context. These characteristics are listed below:

- **All versus some**: In one context, a useful response should provide a complete list of possible answers while in another context some of the possible answers (e.g., only the best ones) are sufficient (Rooy, 2004). Considering the question, *where to exercise during COVID-19?*, if the context is for writing a blog post to inform other people, the answer should be a complete list of possible places in the neighbourhood. In the context of a user searching in the Web, showing only a few places, the best candidates, are enough.

- **Specificity**: A specific answer or a generic one? Specificity determines the form of answers in terms of generic or specific responses (Groenendijk and Stokhof, 1980). For example, *where did you meet her?* can be answered through generic or specific references of the intended place. In some contexts a generic answer such as *in a cafe* is useful, while in a different context the specific answer, *at Cafe Central*, is demanded.

- **Domain**: The domain of answers could be restricted in one context and not in another context (Chierchia, 2006). For example, *what are the best places to visit?* In a wide-domain interpretation, it could be any tourist attraction in the world while in a local domain, the place should be near to the inquirer.

- **Granularity**: The level of details with which the answer is useful depends on the context (Malamud, 2006). A useful answer to a typical where-question such as *where are you from?* depends on the context. In some context, a detailed (fine level of granularity) response is needed – e.g., *Malayer* (a small town in Iran). While in another context, the useful answer is provided in coarse granularity level – e.g., *Iran*.

- **Polarity**: When asking about a property, sometimes a lower bound of the property is a useful response and sometimes the upper bound is needed (Seuren, 1984). For example, *how cold is the Melbourne weather in winter?* If the person is asking to know if she should bring a jacket, then a useful answer is a lower bound of temperature, while if she is unsure whether she can use her hoodies, then the upper bound response is useful to her.

- **Gradability**: Some criteria are gradable and they can be interpreted maximally or minimally. Thus, the usefulness of answer is related to this context-dependent interpretation (Gazdar, 1977). Let us consider the question *is Brunetti open during COVID-19?* Here, the *maximal interpretation* considers affordance of every services including eat-in and takeaway. While, *the minimal interpretation* relates to any services that can be afforded in this place – e.g., takeaway.

- **Plurals**: In case of plural references (e.g., bars) the interpretation of questions' criteria can be either universal (all should satisfy the criteria) or existential (one or more satisfy the criteria) (Krifka, 1996). For example, *are bars open during COVID-19?* Here, the *universal interpretation* rewrites the question as *are all bars open during COVID-19?* while using the *existential interpretation* the question can be rewritten as *is there a bar that is open during COVID-19?* Consequently, this

context-dependant interpretation determines what should be the answer.

The pragmatic study of questions and answers can be viewed using higher level theories of communication. Grice (1975) developed his *conversational implicature theory* with four maxims known as *quantity*, *quality*, *relation* and *manner*. Quantity describes that the information should be neither more nor less than what is required for a successful communication. Quality describes that the response should be accurate and truthful. Relation is being appropriate to an immediate information need at each stage of communication. Finally, manner describes that the response should be clearly executed to be understandable to the inquirer.

Wilson and Sperber (2002) investigated the notion of *relevance* in their seminal *relevance theory of communication*. According to this theory, answers to questions in human communication are *simple*, *short*, *selective* and *cognitively informative* responses. The common question addressed by relevant studies of QA is: *What is relevant to be included in the answers to questions?* Wilson and Sperber's relevance theory is grounded in two principles: (1) the greater cognitive effect by a piece of information, the greater its relevance (appraisal of informative responses), and (2) the greater processing effort of a piece of information, the lesser its relevance (appraisal of short responses) (Wilson and Sperber, 2002).

2.1.2 *Automatic Question Answering Systems*

Automatic QA is defined as a process of finding a correct answer to the question posed in natural language using information stored in documents, databases and knowledge bases (Diefenbach et al., 2018; Gupta and Gupta, 2012). Two major aspects of QA systems are the *content* and the *context* analysis of questions (Marx et al., 2014). Analysing the content determines *what is asked in the question*, while the context is related to *who has asked the question* or *when the question is posed*. Yet, both aspects are necessary for a successful communication between users and QA systems.

Automatic QA systems can be classified based on their approach for information access (Ferrés Domènech, 2017) or the method they use to find the answers (Dwivedi and Singh, 2013). Based on the information sources, QA systems are classified into IR-based QA, database-oriented QA and hybrid QA systems (Dimitrakis, Sgontzos, and Tzitzikas, 2019). In IR-based QA, unstructured documents are used as a source to retrieve answers, while in database-oriented methods structured information in databases and knowledge bases are the basis to answer the questions. The hybrid methods utilize both structured and unstructured information.

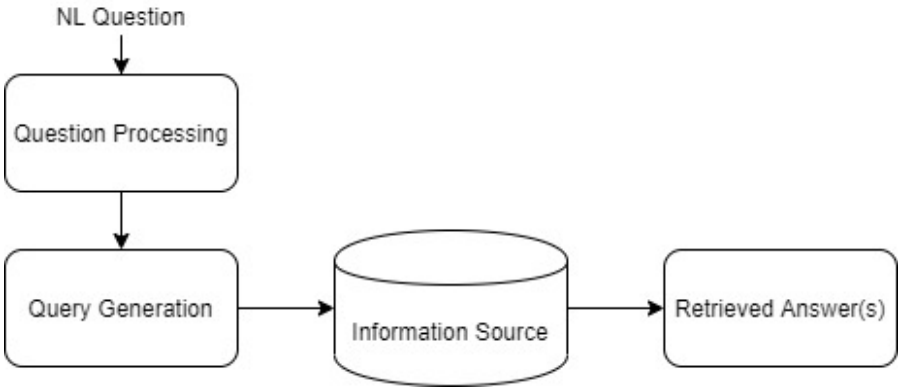

Figure 3: The general architecture of QA systems

Based on the methodology, the QA systems classified into (1) linguistic approaches, (2) statistical methods and (3) pattern matching (Diefenbach et al., 2018; Dwivedi and Singh, 2013). Linguistic approaches include diverse methods for answering questions using NLP (e.g., Chung et al., 2004; Mishra, Mishra, and Agrawal, 2010), and Artificial Intelligence (AI) (e.g., Sharma and Gupta, 2018; Xie et al., 2016). Statistical methods are based on statistical learning models such as Support Vector Machine (SVM), Bayesian classifiers or maximum entropy models (e.g., Bordes, Weston, and Usunier, 2014a; Yen et al., 2013). Finally, pattern matching methods are focused on utilizing templates and regular expressions for finding answers to questions (e.g., Ravichandran and Hovy, 2002; Sneiders, 2002). The pattern matching methods have a limited capacity for real-world QA and can only answer limited types of questions.

2.1.2.1 *QA Architecture*

The general architecture of a QA system includes three main components, (1) question processing, (2) query generation, and (3) finding the answers (Gupta and Gupta, 2012). Figure 3 shows the general architecture of QA systems.

The first step, question processing, includes several sub-tasks such as preprocessing, question classification, intent recognition, relation extraction and question reformulation (Dimitrakis, Sgontzos, and Tzitzikas, 2019). The preprocessing step includes normalising the NL questions into a standard format by removing unwanted tokens (e.g., a, an), string matching to remove typos and misspelling, and standardization of complex references (Saquete et al., 2009). The goal of question classification is to determine the subject of questions which later can be used to find a solution for answering the question (Madabushi and Lee, 2016). Intent recognition or answer type classification determines the type of relevant responses (Di Gennaro et al., 2020). Relation extraction aims at finding how named entities

(e.g., names for people and places) are related to each other in the questions (Li et al., 2019). Finally, question reformulation or question expansion is a task to map questions into a semantically equivalent form that is more explicit and descriptive (Buck et al., 2018).

Query generation is a process of translating the analysed question into a representation suitable for finding the answers. This representation depends on the classification of QA system. For example, IR-based QA systems transform the questions into *vector representations* (e.g., bag-of-words) (Bordes, Weston, and Usunier, 2014b) while database-oriented QA systems translate the questions into *SQL* or *SPARQL* queries (e.g., Huilin et al., 2019; Wang et al., 2020). Later, these representations are used to search information source for finding the answers.

Generating formal queries from NL questions is extensively studied in the last decades (Kim et al., 2020). Kim et al. (2020) describe this task as a set of sequential steps including (1) input enrichment, (2) translation to intermediate (sometimes logical) representation, (3) post-translation and (4) machine learning based or rule-based query generation. The trend for query generation tasks starts with handcrafted rules (Zelle and Mooney, 1996) and lead to development of recent models based on deep neural networks (Dong and Lapata, 2018; Xu et al., 2019). Dong and Lapata (2018) presented a method to first derive a sketch representation that omit the details of the translation task and then generate detailed version of the representation using neural network approach. Xu et al. (2019) presented an encoder-decoder neural network model that can also utilise a database schema to identify out-of-the-scope questions and reject them in the processing step. However, these methods are open-domain approaches that are not necessarily taking the geospatial analysis into consideration.

Finding answers and presenting the response is the last phase of automatic QA. In IR-based QA systems, finding relevant passages (e.g., Cohen and Croft, 2018; Cohen, Yang, and Croft, 2018b), and summarizing the passage or finding sentences or phrases (e.g., Song et al., 2017; Yulianti et al., 2016) are the sub-tasks of this phase (Dimitrakis, Sgontzos, and Tzitzikas, 2019). In database-oriented QA systems, this phase involves finding the answer by submitting the query to knowledge base or database (e.g., Zou et al., 2014), inferring and reasoning (e.g., Mountantonakis and Tzitzikas, 2016) and selecting an appropriate representation (e.g., Xu et al., 2016) to present the answers (Dimitrakis, Sgontzos, and Tzitzikas, 2019).

A new method which is used recently in QA is called Machine Comprehension (MC). This method utilizes deep learning approaches for answering questions (Richardson, Burges, and Renshaw, 2013). The power of deep learning approaches is their high performance results compared to traditional methods, specially with

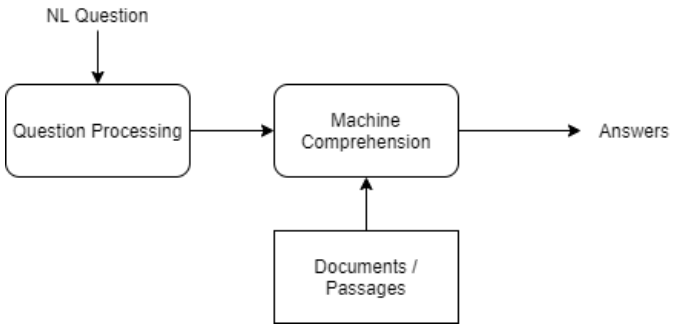

Figure 4: QA architecture using machine comprehension

large amount of data. This approach is a part of IR-based QA but with a simplified architecture. The architecture of QA using MC approaches is shown in Figure 4. Here, the query generation and answer selection phases (including their sub-tasks) are generalised into a deep learning architecture called MC.

2.1.2.2 *QA Datasets*

This subsection reviews the currently available datasets for automatic QA. For textual QA, the datasets can be divided to database-oriented datasets and IR-based datasets (including MC datasets).

In database-oriented datasets, the schema of the knowledge base or database is available. These datasets provide questions and the valid formalised version of the questions (i.e., queries) based on the available schema (e.g., Chen et al., 2019, 2020; Dubey et al., 2019; Trivedi et al., 2017; Yu et al., 2018). Some database-oriented datasets are generated using the schema of well-known knowledge bases such as DBpedia, Wikidata, and Freebase and the valid answers for questions are also included in those datasets (e.g., Dubey et al., 2019; Trivedi et al., 2017).

IR-based datasets include essential information such as a set of questions, answers sentences, and retrieved and selected passages (Yang, Yih, and Meek, 2015). Some datasets are focused on a special task of IR-based QA such as passage retrieval (e.g., Cohen, Yang, and Croft, 2018a) or answer extraction (e.g., Wang, Smith, and Mitamura, 2007). The early IR-based QA datasets provided by Text Retrieval Conference (TREC)[1] (Voorhees, 1999) contain editor-generated questions (e.g., TREC 8-13) while the newer datasets are generated in a more natural process by extracting questions asked by Web search users from query logs (e.g., Yang, Yih, and Meek, 2015).

The MC datasets provide questions, retrieved passages, and the extracted answers as essential information (Bajgar, Kadlec, and Kleindienst, 2016; Hermann et al., 2015; Hill et al., 2015; Rajpurkar et al., 2016; Richardson, Burges, and Renshaw, 2013). Some of these

1 https://trec.nist.gov/data/qa.html

datasets provide additional information such as well-formed answers and question types (e.g., Nguyen et al., 2016). Lakshmi and Arivuchelvan (2019) present a detailed review of several MC datasets and compare their characteristics. This review shows that MS MARCO dataset (Nguyen et al., 2016) has the potential to investigate human answering behaviour due to its answering policy in the data generation. The answers in the MS MARCO could be spanned over several passages. Consequently, the answers in MS MARCO are more similar to responses in human QA due to its free-form answering policy. In addition, providing a large set of questions and answers (i.e., more than 1 million records) and tagging question types makes MS MARCO a suitable dataset to this PhD study. Most of the data analysis parts in this thesis are conducted using the MS MARCO dataset (see Chapters 4 and 6).

2.1.3 *Human QA versus Automatic QA*

Current automatic QA methods are based on a simplistic model of human QA. Human-generated responses to questions differ from the answers *retrieved* from documents and knowledge bases (Mendes and Coheur, 2013). The retrieved responses often contain information irrelevant in the context of the question or missing some essential information needed to satisfy inquirers' information needs. Moreover, the flow of information in the content of the retrieved responses is not adapted to the question and its context due to the retrieval and non-generative answering methods. This is in contrast to the human-generated answers that are specifically *generated* with respect to the content of the question and the context of the communication.

To generate answers automatically, the structure of human-generated answers should be learned in a computer digestible model. However, in open-domain QA, the structure of answers are complex due to the various forms of questions, and cannot be captured in a unified model. Hence, focusing on a specific domain is the solution for designing a structural model of human-generated answers.

One of the contributions of this thesis is to study and imitate the structure of human-generated answers to *where-questions* (see Chapter 6). The syntactic and semantic level of the human-generated answers are investigated using generic geographic information. To minimise the impact of pragmatic interpretations, a dataset that is relatively generated in the same context (i.e., MS MARCO) is used for this study. This dataset includes questions that are submitted a search engine, and answered without any information about the inquirers.

2.2 GEOGRAPHIC QUESTION ANSWERING

GeoQA is defined as methods and algorithms that help inquirers to satisfy their information need by answering their geographic questions (Ferrés Domènech, 2017). In GeoQA, answering geographic questions can be based on diverse information sources such as textual information (Ferrés and Rodríguez, 2006; Mishra, Mishra, and Agrawal, 2010), geodatabases (Chen et al., 2013), and spatially-enabled knowledge bases (Ferrés and Rodríguez, 2010). The answer to these questions can be represented in natural language, structured into tables and graphs, or visualised as maps (Chen, 2014; Ferrés Domènech, 2017; Scheider et al., 2020b). GeoQA architectures typically resolve three tasks: (a) question classification and intent analysis, (b) finding relevant sources, and (c) extracting answers from the sources (Ferrés Domènech, 2017).

Geographic questions include a wide range of questions about different concepts such as geographic *objects* (e.g., places), *fields* (e.g., temperature), *events* (e.g., earthquakes) and *networks* (e.g., transportation). These questions can be posed in everyday communication, searching the Web or using a Geographic Information System (GIS) (Scheider et al., 2020a).

In the following subsections, first the shared tasks between GIR and GeoQA are reviewed. Next, relevant studies to each task of GeoQA are provided and the identified gaps are presented.

2.2.1 *Geographic Information Retrieval*

The goal of GIR is to develop spatially-aware search tools to support users' geographical information needs. The core tasks of a GIR system are extracting and localising place names, qualifying spatial relationships, indexing documents, matching geographic queries to documents by inferring their geographical relevance, and ranking the matching results (Purves et al., 2018).

GIR and GeoQA are both Geographic Information Access (GIA) methods aiming at finding relevant information to satisfy a geographical information need. Despite the differences in their inputs (i.e., queries versus questions) and outputs (i.e., documents versus answers), several tasks are shared between GIR and GeoQA. In other words, several tasks originally developed for GIR are still used in GeoQA. The basic shared tasks in GIR and GeoQA includes toponym identification (aka, geoparsing), toponym resolution (aka, geocoding), and reasoning over spatial relationships (Ferrés Domènech, 2017). Later, these basic tasks are used to analyse geographic queries and documents.

Geoparsing methods are designed to extract toponyms from textual information. Geoparsing is a step to resolve *Geo/Non-Geo*

ambiguity (Purves et al., 2018), – i.e., whether a noun phrase refer to a geographic place or a non-geographic entity. Three main approaches for geoparsing are (1) gazetteers lookup, (2) rule-based methods, and (3) machine learning (Leidner and Lieberman, 2011). Using gazetteers lookup, noun phrases are considered as place names if they match to a place name in predefined gazetteers. This approach is not able to find place names which are not captured in gazetteers, and suffer from mismatches when the noun phrases are in gazetteers but actually refer to non-platial entities (e.g., Washington, when it refers to a person) (Purves et al., 2018). The rule-based and machine learning methods extract toponyms using the surrounding words and the structure of the sentences. The former utilise handcrafted rules (Bilhaut et al., 2003) while the latter learn these patterns from a labelled training dataset (Curran and Clark, 2003).

Geocoding or toponym resolution is the process of disambiguation of the extracted toponyms to their geographic locations (Buscaldi, 2011). In this task, the *Geo/Geo ambiguity* is resolved by assigning a unique location to an extracted toponym (Purves et al., 2018). Buscaldi (2011) categorised the geocoding approaches into (1) map-based, (2) knowledge-based, and (3) machine learning methods. In map-based methods, the spatial configurations of ambigous toponyms in a given context are used to design a heuristic approach for resolving the ambiguity. For example, the *minimum spatial context heuristic* (Leidner, Sinclair, and Webber, 2003) resolves toponym ambiguity by selecting the set of locations (a unique location for each toponym in the text) that covers the smallest geographic area (e.g., bounding box or convex hull area) compared to other possible location sets. Knowledge-based approaches utilises the information available in gazetteers and knowledge bases (e.g., population) to find the best candidate and resolve toponym ambiguity (Rauch, Bukatin, and Baker, 2003). Machine learning methods use labelled datasets and learning models for finding the suitable location for toponyms. These methods are able to use other geographic information such as events for toponym resolution (Roberts, Bejan, and Harabagiu, 2010).

Extracting spatial relationships from NL texts is one of the basic tasks of GIR. Spatial relationships are often expressed through prepositions in natural language (Bennett, 1975). The interpretation of spatial preposition is a challenge due to their context dependency. For example, the spatial preposition 'at' can have different meaning depending on the scale level and classification type of the anchor places (Vasardani et al., 2012). Here, the goal is to model the extracted spatial relationships and enable GIR systems to perform spatial reasoning for analysing geographic queries and documents (Sallaberry et al., 2008).

2.2.1.1 *Analysing Geographic Queries*

In GIR literature, different methods are proposed to classify geographic queries (e.g., Sanderson and Kohler, 2004), to extract intents of the queries (e.g., Henrich and Luedecke, 2007) and to model the queries based on their structure and semantics (e.g., Purves et al., 2007).

GIR from Web sources focuses primarily on: (1) geographic Web queries (e.g., Aloteibi and Sanderson, 2014; Sanderson and Kohler, 2004), and (2) geographic context of Web search and its spatial variation (e.g., Backstrom et al., 2008). Geographical Web queries were first described by Sanderson and Kohler (2004) who focused on *explicit* geographic queries, those containing place names (i.e., toponyms). Later, *implicit* geographic queries that relate to generic nouns designating locations known to a subgroup of people (e.g., a family sharing a location known to them as home) have also been identified (Wang et al., 2005). The manual labeling of such geographic queries was described by Gan et al. (2008).

In contrast, Jones et al. (2008) analyzed the intent behind spatial search executed through Yahoo. They studied the distances between toponyms in the query and the location of the user indicated by their IP address. They then classified the spatial search focus by distance into granular categories: same-city, same-state, same-country and different-country, and correlated search distance distributions with the topics searched for. A more recent study of the geolocated search logs of Bing (Lee, Craswell, and Murdock, 2015) profited from the higher accuracy of user positioning, supporting a finer grained analysis of the search location to searched entity distances, and correlated it with the distribution of entities of different types in space. They found that users have a strong preference for the closest available match for the search conditions, although these preferences are weaker for certain types of business categories. The preferred search results issued by users from familiar and unfamiliar locations are distinct, as shown by Kravi et al. (2015). Similar findings have been reported in differences in search queries by frequent versus one-time visitors to shopping centres (Ren et al., 2017).

Henrich and Luedecke (2007) developed a bottom-up approach to categorize geographic web queries into four main groups, *finding accommodation*, *searching for habitation*, *seeking information about places*, and *finding places that afford leisure time activities*. They proposed a taxonomy of information needs for geographic questions, consisting of four classes: (1) *to there*: wayfinding purpose; (2) *about there*: information about a place; (3) *there*: activity inside a place; and (4) *from there*: to get something from a specific place (Henrich and Luedecke, 2007).

Purves et al. (2007) proposed a model for geographic queries based on the thematic and spatial aspects of the queries. This model

is a triplet constructed from three constituents – i.e., <*theme, spatial relationship, toponym*>. This model is used to build a geographic search engine called *SPIRIT* which enable geographic search through a structured query interface (Purves et al., 2007). Hamzei and Hakimpour (2017) developed an iterative approach to extract toponyms and spatial relationships from unstructured keyword-based queries. This approach is used to analyse geographic queries which are constructed based on toponyms and spatial relationships, i.e., <*toponym, spatial relationship, toponym*>.

2.2.1.2 *Analysing Geographic Text*

The task of analysing NL texts in GIR has two purposes: (1) extracting geographic information from NL texts to build knowledge bases and (2) determining geographic scope of Web documents (Santos and Chaves, 2006). The former lead to deriving and populating geographic ontologies and the latter is a basis for matching and retrieving relevant documents to geographic queries.

Kordjamshidi, Otterlo, and Moens (2011) proposed a triplet model to extract spatial information from NL texts. The triplet is composed of *trajector*, *landmark* and *spatial indicator* parts. The trajector is the spatial object which the sentence localises. Landmark is the anchor object which its location is taken for granted and spatial indicator is a spatial relationship which describes the location of trajector based on the landmark location. Later, Vasardani et al. (2013) developed a similar triplet model which includes *locatum*, *relatum* and *spatial relationship*. The difference was that Vasardani et al. (2013) proposed a conceptual framework to not only extract these triplets (Khan, Vasardani, and Winter, 2013) but to merge them (Kim, Vasardani, and Winter, 2017) to from a property graph which is called *place graph*. Later, this model extended in N-plet model to capture nary relationships and enriched but essential place-related information such as place type and affordance (Chen et al., 2018). The thematic aspects of geographic information are also investigated in literature. McKenzie and Adams (2017) proposed a topic modelling approach to extract thematic aspects of places from NL texts.

Andogah, Bouma, and Nerbonne (2012) defined geographic scope of a NL text as a region or an area which the document is about. Geographic scope enable GIR systems to index documents based on their geographic scope and retrieve them in an efficient manner using spatial indexing methods (Monteiro, Davis Jr, and Fonseca, 2016). The task of determining geographic scope of documents is based on geoparsing and geocoding of their textual content. Often the dominant location or a weighted mean of the extracted locations are considered as the geographic scope of the text (Purves et al., 2018). Some studies suggest that a bounding box of extracted

location can be useful for GIR purposes (Brun, Dominguès, and Van Damme, 2015).

2.2.1.3 *From GIR to GeoQA*

The transition from GIR to GeoQA, and in general from IR to QA, is the results of new communication channels (e.g., voice search), and new expectations from progressive developments of IR systems (Guy, 2018; Zhao et al., 2019). Analysing query logs show emerging patterns in queries that are submitted to search engines. These search queries are more similar to NL questions than traditional keyword-based queries (Guy, 2016; White, Richardson, and Yih, 2015). Consequently, researchers claim that a new design to enable natural interaction between users and IR systems is a necessary step ahead (Zhou, 2018).

Zhou (2018) described *Natural Search* as the future of IR systems. The naturalness of interaction between users and IR systems can be achieved by enabling natural language QA. However, to enable natural search new methods are needed for analysing questions and generating relevant answers.

The difference between queries and questions is not only about the syntax or presentation but also about the semantics and intents. In terms of syntax, the NL questions are often longer (based on number of tokens), and avoid abbreviations (Guy, 2016). NL questions have clear and specific intents which are interpretable by people (Guy, 2018) while queries are often vaguely defined. Consequently, the task of analysing questions is to devise models and methods for enabling computer to digest and interpret NL questions.

Several search engines, including Google, ASK[2] and START[3], initiated their question answering capabilities recently (Zhao et al., 2019). The answers are retrieved from knowledge bases (e.g., Google Knowledge Graph) or directly extracted from a selected document. These answers are shown in answer box on top of the ranked document list whenever the system can retrieve answers to the questions. Zhao et al. (2019) experimented the answering capability of Google Search Engine and they found that the current results are limited in several ways. While the system relatively perform well for person-related questions (e.g., celebrities), other questions such as event-related and organization-related are often remained unanswered. In addition, the current answering capability is only limited to simple factoid questions and unable to answer more sophisticated types of questions (Zhao et al., 2019).

Purves (2014) described answering geographic questions as one of the neglected aspects in the current GIR systems. Currently, the

2 http://www.ask.com
3 http://start.csail.mit.edu

geographic questions are treated equally as other types of questions despite their fundamental differences such as high context dependency, vagueness and the need for spatial reasoning and analysis to properly answer them (Mai et al., 2020).

2.2.2 *Analysing Geographic Questions*

Analysing geographic questions in GeoQA includes two major sub-tasks, (1) question classification and (2) intent recognition. The former defines a type for each question which is later used for finding a solution strategy (Ferrés Domènech, 2017), while the latter describes what is the purpose of question and what should be presented in the answers (Jang et al., 2010).

Several classifications are proposed in literature to annotate geographic questions. Li and Roth (2002) proposed a hierarchical classification of open-domain question. The first level of the hierarchy includes six types of questions – i.e., abbreviation, entity, description, human, location and numeric. The second level of the hierarchical classes include detailed classification which some of these classes are related to geographic questions. The location category is divided into questions about *city*, *state*, *country*, *mountain* and *other places*. Some of the numeric questions such as *distance*, *temperature* and *count* can be also related to geographical concepts. Here, the questions are classified based on some properties and types of geographic places.

Ferrés and Rodríguez (2006) proposed a classification for geographic questions based on their corresponding answer types. This schema includes nine types of questions: (1) *count-objects*, (2) *how-many-people*, (3) *what-area*, (4) *what-flow*, (5) *what-height*, (6) *what-length*, (7) *where-action*, (8) *where-location*, and (9) *where-quality*. This classification capture what/where/how-many questions about geographic places but only for a few popular properties such as length, height and actions.

Chen et al. (2013) developed a GeoQA system to answer basic types of geographic questions. These question types are derived from simple spatial analysis such as *locating places*, *calculation of distance*, *determining proximity buffers*. The templates for these questions are (1) where is X?, (2) how far is X from Y?, (3) which Y is the nearest to X?, and (4) what Y are within a specific distance from X?

Punjani et al. (2018) categorized geographic questions into question templates based on location, type, spatial relationships and thematic attributes of places. They identified seven categories of questions which are listed in Table 1. This classification purely captures the syntactic structure of geographic questions and biased to only a specific types of questions.

Table 1: Classification schema of geographic questions (Punjani et al., 2018)

Question Type	Example
location of place X	Where is Melbourne?
place X, spatial relation R, place Y	Is Melbourne south of Sydney?
place type T, spatial relation R, place Y	What are the cities near to Melbourne?
place type T1, spatial relation R, place type T2	Which restaurants near to bars?
place type T, spatial relation R, place type T2 in a given geographic context	Which restaurants near to bars in Melbourne?
thematic properties P of place X	What is the population of Melbourne?
quantities and aggregates	Which is the highest mountain in Australia?

Jang et al. (2010) described sixteen categories of intents for geographic queries. They studied geographic information in textual documents to propose this categorisation of intents for geographic questions. This categorisation includes (1) *direction*, (2) *price*, (3) *quality*, (4) *landmark*, (5) *location*, (6) *time*, (7) *reason*, (8) *product*, (9) *opinion*, (10) *event*, (11) *accessibility*, (12) *mood*, (13) *list*, (14) *experience*, (15) *activity* and (16) *number of participants* (Jang et al., 2010). This categorisation can be simplified into *navigation* between places, *location* of a place, *getting information* about places and experiences related to *activities* and *sense of place*.

The aforementioned classification and analysis of geographic questions are biased to expert opinions to establish the classification schema. Consequently, only the popular types of geographic questions are expressed in these schemas. A bottom-up approach over a large corpora of question can be used to find more diverse types of questions. In addition, a theoretically grounded top-down approach can be used to derive a classification schema with less biases to experts' opinions. In chapter 4, this gap is resolved using both top-down and bottom-up approaches to classify and analyse geographic questions.

Another gap which is identified here is related to intent recognition. While the list of intent categories provide a detail description of question intents, it does not provide a solution to extract the intent from the questions automatically. In Chapter 5, a heuristic method is used to extract the intent of questions based on their grammatical structure.

2.2.3 *Modelling Geographic Questions*

The major task of modelling geographic questions is to translate NL questions into a formal representation that can be used for retrieving or generating the answers. This formal representation can be a bag-of-words vector representation (e.g., Ferrés and Rodríguez, 2006) in IR-based GeoQA or a formal query in SQL (e.g., Chen, 2014) or SPARQL (e.g., Punjani et al., 2018) when the question is submitted to a database or a knowledge-base, respectively.

Zelle and Mooney (1996) and Tang and Mooney (2001) designed a logical form to translate geographic questions into database queries. This logical form includes factual statements extracted from the questions in terms of objects and predicates. The objects are limited to a few geographic concepts such as *city*, *country*, *state*, *river*, and *place*. Inferring objects from questions are achieved through a dictionary lookup of place names for the United States. The predicates are relations and operations that are applied to objects. The predicate list includes twenty functions for defining objects, getting their properties such as population and area, defining spatial and logical relationships such as capital-of and equal (Tang and Mooney, 2001; Zelle and Mooney, 1996).

Chen (2014) devised a method to translate a limited types of geographic questions into *Spatial SQL* queries. This method is only developed for four types of questions: (1) finding coordinates of a place, (2) finding distance between places, (3) finding closest place to another place and (3) finding places in a predefined neighbourhood of another place. The method is based on annotated questions and predefined Spatial SQL templates. For each type of questions, a Spatial SQL template is designed and later these templates are parameterised using annotated questions (Chen, 2014). This method uses standard functions of Spatial SQL such as *ST_Distance*, *ST_Buffer* and *ST_WITHIN* to perform spatial analysis for finding answers.

Punjani et al. (2018) developed a template-based approach to generate GeoSPARQL queries from NL geographic questions. Their approach includes two step of natural language processing for information extraction and resolution, and using a query generator to parameterise GeoSPARQL templates. In the first step, toponyms, place types, spatial relationships and properties are identified and extracted from NL questions. Later, for each question type (see Table 1) a predefined GeoSPARQL template is designed. Finally, the standard spatial functions (e.g., topological and metric functions) implemented in GeoSPARQL are used to test and evaluate the resulted GeoSPARQL queries.

Modelling questions to be answered by GIS is a complicated process of extracting information, finding relevant GIS layers, and

performing spatial analysis using diverse set of toolboxes (Scheider et al., 2020b). Scheider et al. (2020a) described a process of extracting information from NL geo-analytical questions as finding textual references related to core concepts of spatial information (Kuhn, 2012). Then, these annotated questions are translated into a spatial analysis workflow that links the data layers to the result layer using spatial analysis tools (Scheider et al., 2020b).

Several studies are focused on rewriting questions when the geographic questions are either complex or unanswerable. Hartrumpf and Leveling (2010) designed a recursive method to rewrite complex geographic questions into simpler set of questions which can be answered using GeoQA methods. Mai et al. (2020) proposed a relaxation method for unanswered geographic questions. They have used distance decay to rewrite query form of unanswerable geographic questions into a relaxed query which can be answered by information stored in a knowledge base.

While several studies are conducted to model geographic questions, a proper intermediate representation of the questions is still missing. The intermediate representation should be a model that capture semantics of questions without considering technical aspects of formal query languages. In Chapter 5, a tree-based representation of question is presented to capture semantics of place-related questions. In addition, current templates are mainly devised for limited types of questions. Chapter 5 shows a dynamic construction approach which is developed to fill this gap.

2.2.4 *Answering Geographic Questions*

The last phase of GeoQA is to formulate answers. This task can be divided into finding and retrieving relevant information, and presenting the results in an appropriate form. The former results in relevant content for the answers and the latter focused on how the content should be expressed to the inquirers. The aim of this phase is to present correct, useful and non-misleading responses which are known as cooperative answers (Mendes and Coheur, 2013).

The understanding of the content of questions (*available information inside the question*) and the context of the conversations (e.g., *information about the inquirers and their state of knowledge*) are essential aspects of QA (Marx et al., 2014). Geographic questions contain clues about what would be relevant to be mentioned in the answers (Behrangi et al., 2007). In addition, the inquirers' physical location is noted as an important factor for providing relevant answers to geographic questions (Shanon, 1983).

Methods for selecting relevant responses depend on the information sources used for answering the questions. In IR-based GeoQA often a single span of a document or a snippet is extracted as

answers to geographic questions (e.g., Mai et al., 2018; Mishra, Mishra, and Agrawal, 2010). These methods used word representation models and NLP techniques to train a model over annotated datasets for extracting the answers. In database-oriented GeoQA the answers are a partial view of a database tables (e.g., Chen, 2014) or knowledge graphs (e.g., Punjani et al., 2018; Stadler et al., 2012) which are selected via a formal query. In the special case of geo-analytical questions, the answers are the results of a spatial analysis workflow which are not retrieved but generated (Scheider et al., 2020b). In this case, the answers are extracted by running a workflow of predefined GIS analysis tools over a selected set of GIS layers.

In terms of question types, GeoQA studies are mostly focused on what/which questions about geographic places (e.g., Scheider et al., 2020b; Vahedi, Kuhn, and Ballatore, 2016). In answering complex questions such as *where-questions* and *how-to-get-to questions*, the task is either simplified into retrieving stored coordinates (Luque et al., 2006; Stadler et al., 2012), or selecting a part of textual information without any adaptation to QA context (e.g., Buscaldi et al., 2006). However, a study on human-generated answers to *where-questions* and *how-to-get-to questions* are needed to generate proper place and route descriptions as answers. In Chapter 6, a detailed analysis of human-generated answers to where-questions is presented.

The answers to geographic questions can be expressed in different communication channel including visual and auditory channels. They can be expressed in unstructured natural language form (Mishra, Mishra, and Agrawal, 2010), or presented in structured forms such as tables and graphs (Chen, 2014; Punjani et al., 2018). Moreover, the answers can be shown in a map visualisation (Scheider et al., 2020b). The presentation style of answers is the part that is mainly missing in GeoQA. Specially for NL responses to *where-questions* and *how-to-get-to* questions, this task is particularly challenging. To fill the gap, a proper structural analysis is needed to understand and imitate the structure of human-generated answers. In Chapter 6, I investigate the structure of human-generated answers to where-questions. Chapter 6 describes generic templates to find structural patterns and to use these patterns to imitate human answering behaviour to where-questions.

2.3 SUMMARY

This chapter described a review of QA and GeoQA literature to provide a background for the following chapters. First human QA and automatic QA are reviewed and compared. The results show that human QA is more complex and the current retrieval approaches in automatic QA have a simplistic view to natural

question answering. Later, relevant GIR and GeoQA publications were scrutinised to identify the current challenges and gaps. The identified research gaps can be summarised as:

- The proposed schemas for analysing geographic questions are heterogeneous and biased to experts' opinions. Hence, a proper classification schema for geographic question is still missing (see Section 2.2.2).
- An automated method for extracting and annotating information from geographic questions calls for a proper syntactic and semantic analysis question contents (see Section 2.2.2).
- Extracting the intent of geographic questions currently performed through human (i.e., expert) interpretations and developing a method to recognise the intents is a necessary step toward automatic GeoQA (see Section 2.2.2).
- Current methods for translating geographic questions to formal query representations are limited to a handful types of questions. Hence, a method to dynamically generate queries is required to support diverse question types (see Section 2.2.3).
- Retrieving responses instead of generating useful answers is a major difference between the machine-generated and human-generated answers. A detailed structural analysis of human-generated answers is required to learn and imitate human answering behaviour (see Sections 2.1.3 and 2.2.4).

In this thesis, the scope of GeoQA is narrowed down to place-related QA and the identified research gaps are investigated only for this subset of GeoQA. To systematise the study of analysing, modelling and answering place-related questions, the notion of place should investigated and modelled. In other words, a model of place is a prerequisite of place-related QA that should be addressed in the first step. Such a place model is a theoretical foundation to develop schema for annotating and analysing geographic questions, modelling questions into intermediate representations and designing a structural model for analysing and imitating human-generated answers.

3

A MULTI-FACETED MODEL OF PLACE

They are all spaces which people have made meaningful.
They are all spaces people are attached to in one way or another.
This is the most straightforward and common definition of place
– a meaningful location.

— Cresswell (2014)

3.1 INTRODUCTION

Place has been defined as space overlaid with meaning (Relph, 1976; Tuan, 1977) and attachment to individuals (Cresswell, 2004, 2014) or groups of people (Turner and Turner, 2006). The difference between the notions of place and space has been studied in geography and captured in the space-place continuum, with space and place at the respective extremes (Couclelis, 1992; Edwardes, 2007). Space is an abstract concept (Couclelis, 1992) suitable to be formalised in a computer-based environment, while place is related to human experiences of the world in a common sense manner (Cresswell, 2004) and its formal modelling is thus challenging (Purves, Winter, and Kuhn, 2019). Due to multi-disciplinary interests in the notions of space and place, several definitions of these terms have been proposed in the literature (Edwardes and Purves, 2007a; Turner and Turner, 2006). For example, Harrison and Dourish (1996) described the difference between space and place similar to the difference between a house (the abstract) and a home (the personal), while Tuan (1977) has called this difference equivalent to the difference between freedom (openness) and security (stability). Alternatively, in architecture and urban planning, space is regarded as the raw material which turns to place after design is introduced to it (Hillier, 2007).

In addition to these diverse definitions, the notion of place has also been conceptualised differently, depending on the emphasis of

o This chapter is based on the publication ***"Place facets: A systematic literature review"*** published in *Spatial Cognition & Computation* by Hamzei, Winter, and Tomko (2020). My contribution as a first author of this work includes proposing the hypothesis and research problem, designing a systematic literature review method, analysing the results and presenting the discussions and conclusion. This work was supervised by *Prof. Stephan Winter* (my primary supervisor) and *Dr. Martin Tomko* (my co-supervisor) who provided feedback and suggestions in every stage of this study.

the distinct perspectives on place (Turner and Turner, 2006). In particular, phenomenological (Relph, 1976), sociological (Gustafson, 2001), and psychological (Canter, 1977) perspectives have contributed to the fundamental conceptualisations of the notion of place in human geography, social science, and environmental psychology, respectively (Turner and Turner, 2006). These conceptualisations differ in their shifting emphasis between the social (e.g., functional differentiation (Canter, 1997)) and individual meaning of place (e.g., sense of place (Agnew, 2011)), as well as between objective (e.g., location (Agnew, 2011)) and subjective aspects of place (e.g., attachment (Relph, 1976)). These conceptualisations are complementary, yet manifest also unresolved conceptual overlaps.

Apart from the conceptualisations of the notion of place, several attempts have been made for modelling information about (specific) places (ElGindy and Abdelmoty, 2014a; Papadakis, Resch, and Blaschke, 2016; Scheider and Janowicz, 2014), to the end of capturing this information in (geographic) information systems (Jonietz, 2016). While the conceptualisations are not primarily interested in formalisation, the information models are presented to capture information about places in a formalised manner. Two major and divergent views on modelling places in GIScience view places *as objects* with identities (Purves, Winter, and Kuhn, 2019; Scheider and Janowicz, 2014; Scheider and Purves, 2013), or in contrast, *as emergent aggregates* of spatial units with similar themes (Ballatore, 2014; Jonietz, 2016; Papadakis, Resch, and Blaschke, 2016). In addition to the conceptual differences between these views, the individual models differ in the types of information stored. For example, gazetteers mainly capture names, types, and locations of geographical objects (Hill, Frew, and Zheng, 1999). More recently, refined models were presented to capture also the affordance and equipment of places (Goldberg, Wilson, and Knoblock, 2009; Scheider and Janowicz, 2014). However, what should be included in a place information model itself and what is domain-dependent information remains open (Ballatore, 2016).

The diversity of the multidisciplinary perspectives on definitions, conceptualisations and models of places motivates this study. Its primary aim is to report on the current state of place-related research, by systematising these perspectives into a multi-faceted model of place. As described in Chapter 1, place is narrowed down to geographic places from a room-level scale to the scale of the Earth itself in this PhD thesis. Other types of places, such as virtual places (e.g., a Web-based chat room), imaginary places (e.g., cloud-cuckoo-land), smaller places (e.g., table-tops) and mystic places (e.g., Heaven) are not within the scope of this study.

In order to extract, categorise, and report on the results of previous studies, this study focuses on facets of place, rather than an entire conceptualisation or information model. A facet of place is here defined as *a particular type of information about (geographic) place that has been defined, described, or formalised in the literature and at the same time can be used to differentiate places from each other*. Here, facet is defined and used as a general and inclusive term that supersedes terms such as properties (e.g., good shape which is a property of place), attribute and characteristics (e.g., type), and aspects (e.g., social aspect, or functional aspect). For example, *location* can be considered as a place facet because it is a type of information about places, and places may be differentiated based on their locations. Facets can thus be considered as units with which the models and conceptualisations of place can be constructed.

The term *facet*, as it is defined in facet theory (Canter, 2012), has an additional criterion of being mutually exclusive from each other. The facets of place extracted from the literature are, however, not necessarily mutually exclusive. Therefore, one of the contributions of this study is the attempts to systematise the facets in order to remove duplications and deriving categories of facets which are mutually exclusive as far as possible. These categories form a model of place that capture and systematise place properties defined in previous studies. As such, the mutual exclusion criterion is only aspirational in this study and treated in a relaxed manner using hierarchical categorisation of facets. Consequently, the deeper the level in the hierarchical categorisation, the less valid are the categories in terms of mutual exclusivity.

Providing a systematic review of place in the scientific literature enables to capture what is already known about the notion of place, and what the similarities and differences in previous work are. Diversity in terminologies and different levels of formalisms in place-related research contribute to confusions in understanding place and in scholarly communication about the notion of place. In order to address these issues, this chapter bridges the different perspectives by categorising the identified facets of place.

3.2 METHODOLOGY FOR A SYSTEMATIC LITERATURE REVIEW

This study realises a Systematic Literature Review (SLR) aiming at reviewing, synthesising, and reporting results in a reproducible manner (Okoli, 2015). A SLR is grounded in four phases as defined by Okoli (2015):

(1) A **planning phase** includes defining the purpose of the review, and the design of a protocol determining the procedure for applying selection, extraction and execution phases. In the

planning phase of this study, the purpose of the review is defined through the objectives mentioned above (Section 3.1).

(2) The **selection phase** of a SLR defines how papers are selected and how screening is done for excluding publications that are out of the scope of the SLR. The forthcoming section (Section 3.2.1) describes the selection method of this study to support a reproducible review procedure.

(3) The **extraction phase** includes an appraisal of the quality and the extraction of data from the selected papers. Here, the place facets are what is extracted from each selected paper. The extraction of facets is presented in Section 3.2.2.

(4) The **execution phase** defines how the extracted data are synthesised and how the results are reported. In this study, the execution phase includes card-sorting and hierarchical clustering methods to derive the multi-faceted model of place (Section 3.2.3).

3.2.1 *Selection Phase*

Figure 5 shows the process of selection and screening used in this study. First, a set of journals in different disciplines are subjectively selected for finding *seed papers* that are chosen based on their relevance to facets of place (Table 2).

The seed journals are the high-quality journals (based on impact factors) in Human Geography, GIScience and Environmental Psychology. The focus on these disciplines sources from the purpose of this study to serve the community of information science for enriching current conceptualisations and models of place. The title, abstracts, authors, and citations of all papers published in the selected journals are extracted using the Scopus API, limiting the study to English publications between 2008 and 2018. Keyword-based criteria are not used in this selection process due to the possible interchangeable usages of the term *place* with other related terms in the literature, such as *landscape*, *region*, and even *location* (Bennett and Agarwal, 2007; Cresswell, 2004).

The selection of a refined list of publications is accomplished by exclusion of irrelevant publications, in three iterations: based on the *title*, the *abstract*, and in some cases based on the content of the publications themselves (using subject headings and the conclusion sections). Only publications with the main scope of *place*, *region*, and *human geographical knowledge*, or with a focus on *a specific place facet* (e.g., affordance) are retained for analysis and review.

During the selection process, the irrelevant publications are (manually) removed from the selection list due to the out-of-scope

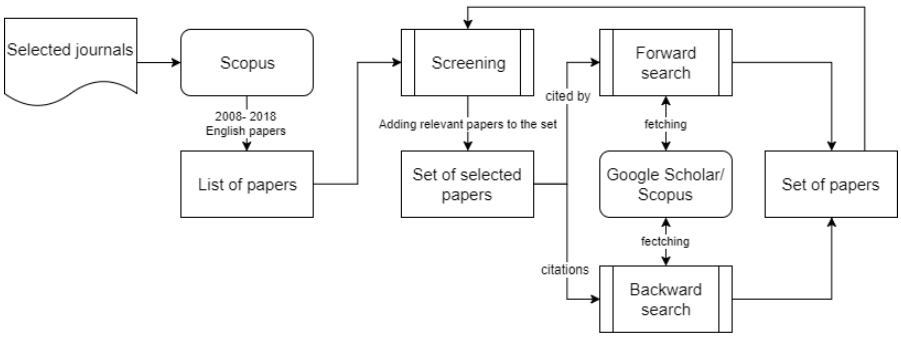

Figure 5: Selection and screening process.

Table 2: Selected journals

Journal name	Field of study
Progress in Human Geography	Human Geography
International Journal of Geographical Information Science	GIScience
Journal of Spatial Information Science	GIScience
Spatial Cognition and Computation	GIScience; Psychology
Journal of Environmental Psychology	Psychology

reason. Specifically, three generic rules are followed in exclusions of unwanted publications.

(1) If the publication is not about place in any sense; it is not relevant to the purpose of this review – e.g., miscellaneous topics such as spatial data infrastructure or spatial algorithms.

(2) If the publication is related to place, but place is only a case study or of a secondary interest then the publication is unwanted – e.g., investigation of political issues in different places.

(3) If the publication is about place but it is not directly related to facets of place, then it is considered as out-of-scope and, consequently removed from the list – e.g., publications in 'placemaking'.

This selection process is designed to avoid possible bias to a specific view to the notion of place. It enables to select publications with different objectives, from proposals of formal place models (e.g., Scheider and Janowicz, 2014) to explorations of the concept of place by discussion of examples (e.g., Capineri, 2016). In addition, publications about an individual facet of place (e.g., Galton, 2010) to a conceptualisation (e.g., Agnew, 2011), an ontology (e.g., Jones,

Alani, and Tudhope, 2001) or a model of place (e.g., Jordan et al., 1998) are included in the review process.

After selecting this set of seed papers from the seed journals, two iterations of forward and backward searches are undertaken. In the forward and backward searches the *cited-by* and *citations* of the selected seed papers are analyzed with the same screening criteria, using Scopus and Google Scholar. Books and dissertations are excluded from the backward and forward search (e.g., Cresswell, 2004; Tuan, 1977) due to their broad scope and large numbers of cited-by and citations. The reason to include Google Scholar as a tool for finding new papers stems from the lack of sufficient coverage of a number of relevant journals, conferences, and books in Scopus.

3.2.2 *Extraction Phase*

In this phase, the selected publications are read, and place facets are extracted from text. For each facet, its name, definition, and examples are collected from text. The facets with the same name and similar definition are merged into one facet. In addition, the publications that the facets are extracted from their content are recorded. The results of this phase address the first objective of this study. In order to present a comprehensive list of facets, in addition to explicitly-defined facets, implicitly-defined facets are extracted as well. This is done based on the best judgment of the author, using the following definitions:

- **Explicitly-defined facets** are facets named and defined in a relatively formal manner in the source paper(s). Some of these facets are part of a conceptualisation, a formal ontology or a model (e.g., *affordance* in the model presented by Scheider and Janowicz (2014)), and the others are defined separately in the literature (e.g., *spirit of place* (Vanclay, 2008)).
- **Implicitly-defined facets** are those that are not defined formally in the sourced paper(s), but their informal definitions may still be found (e.g., *saliency* as defined by Winter and Freksa (2012)). These facets may sometimes not be named but are discussed through examples (e.g., *spatial reference* as described by Winter and Freksa (2012)).

3.2.3 *Execution phase*

As shown in Figure 6, the final categories of place facets (i.e., the multi-faceted model) are the results of synthesis of the extracted facets using card sorting (Spencer, 2009) and hierarchical clustering (Johnson, 1967) techniques. The extracted facets – names, definitions, and examples – had been written on cards, and handed out to three researchers to independently organise them.

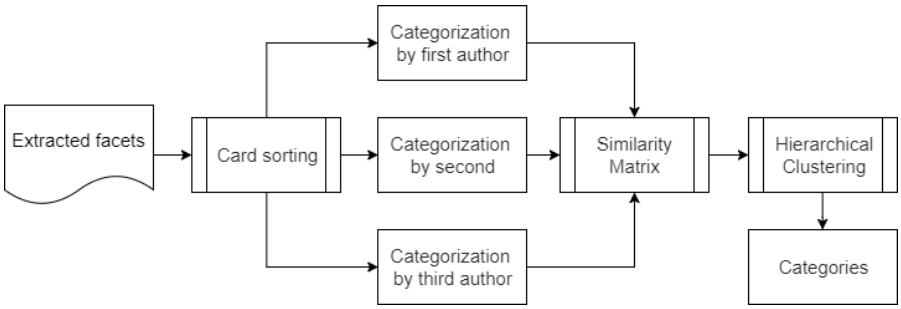

Figure 6: Process of synthesising place facets in categories.

The open card sorting strategy (Spencer, 2009) is used by the participants to establish their own subjective groupings of facets assisted by the definitions extracted in the previous step. In open card sorting, no predefined set of categories are available, and the participants are free to make and label their own groupings. Hence, the results usually contain different numbers of groupings and meanings (group labels).

A convenient way to combine the results of card sorting experiments is to use hierarchical clustering (Spencer, 2009). To utilise hierarchical clustering, a similarity matrix is constructed, capturing agreements where a pair of facets belongs to same group. The results of the card sorting by the participants are thus combined in a similarity matrix capturing whether facets are categorised under the same category in individual groupings. In other words, the similarity matrix captures the count of groupings where two facets belong to the same category (Table 3). Agglomerative hierarchical clustering (Johnson, 1967; Maechler, 2018) is then applied on the similarity matrix and different cut-offs of the categorisation are derived to extract a small number of meaningful higher-level groups. The cut-offs of hierarchical clustering determine the categories of facets derived by different groupings, and the results form a hierarchical multi-faceted model of place. The relation between the derived categories of different cut-offs is thus a parent-child relation. Finally, the results of hierarchical clustering are interpreted and reported.

3.3 RESULTS OF THE SYSTEMATIC LITERATURE REVIEW

In this section, the results of the selection, extraction, and execution phases of the systematic literature review are reported and discussed. The selected publications, their type, publication dates, and citations are investigated in the first section. Next, the extracted facets, and derived categories of facets are discussed.

Table 3: Hierarchical clustering distance/similarity matrix

Distance	Similarity	Description
3	0	if the pair of facets are not in same category in any of individual groupings;
2	1	if the pair of facets are in same category only in one of the individual groupings;
1	2	if the pair of facets are in same category in two individual groupings;
0	3	if the pair of facets are in same category in all (three) individual groupings

3.3.1 *Place-Related Research: Results of selection*

Based on the selection and screening criteria, 72 publications were selected for identifying and extracting place facets. Table 4 shows the number of selected papers in each iteration. The list of selected publications includes books, book sections, conference papers, journal papers, and theses (Table 5).

Table 4: Number of selected papers in each iteration

Iteration	#Unique Publications[1]	#Selected
Seed papers	2560	13
1st iteration (backward/forward)	1005	38
2nd iteration (backward/forward)	489	21
All	4056	72

Table 5: Type of publications.

Type	Number of publications
Books	7
Book sections[2]	7
Conference papers	20
Journal papers	34
Theses and dissertations	4

One way to analyze the selected publications is based on the publication date. Figure 7 shows aggregated counts of place-related publications across disciplines from the 1970s to present. In the selected publications, a trend of increasing scientific interest in place facets is observable. In a discipline-based interpretation, the early

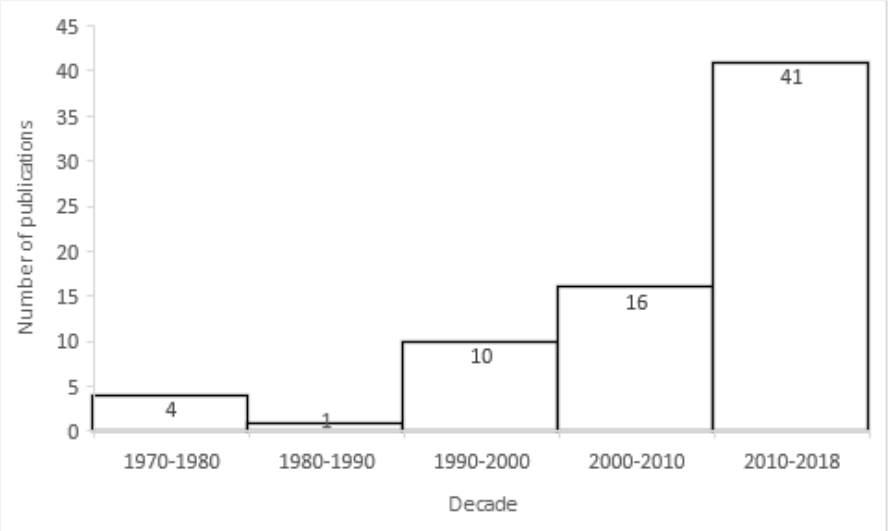

Figure 7: Publication date analysis of the selected publications, by decade.

research – in the seventies to the nineties – are landmark publications in human geography (e.g., Buttimer, 1976; Relph, 1976; Tuan, 1977) and environmental psychology (e.g., Canter, 1977; Gibson, 1979). Later, several publications are selected from sociologists (e.g., Gieryn, 2000; Gustafson, 2001) in the list, and recently, attempts to formalise and to model the notion of place in the information science community (e.g., Adams, 2015; Papadakis, Resch, and Blaschke, 2016; Vasardani and Winter, 2016). Hence, following the screening, a salient sequence of scientific work are noted progressing from initial explanations, through conceptualisations to computational models and formalisations.

While a citation analysis of scientific publications cannot comprehensively evaluate the quality of the academic works (MacRoberts and MacRoberts, 1989), citation counts are often used as a rough quantitative measure to evaluate the impact of publications. Figure 8 shows the seventy-two selected publications, ordered by citations counts [3]. Four of the top-five most cited publications (Cresswell, 2004; Gibson, 1979; Massey, 1994; Relph, 1976; Tuan, 1977) are all books, with the only exception (Cresswell, 2004), representing an accessible summary work written a relatively long time ago. On the other hand, the five publications without any citation, (Almuzaini, 2017; Calafiore, 2016; Mennis and Mason, 2016; Papadakis, Resch, and Blaschke, 2016; Vasardani, Tomko, and Winter, 2016) are all conference papers and a thesis that have been written very recently. Hence, the relatively large differences (in orders of magnitude) between the citation frequencies of publications stem from "variation in citation rate with type of publication, nationality, time period, and size and type specialty" as

3 The citations are collected from Google Scholar on 09/07/2018

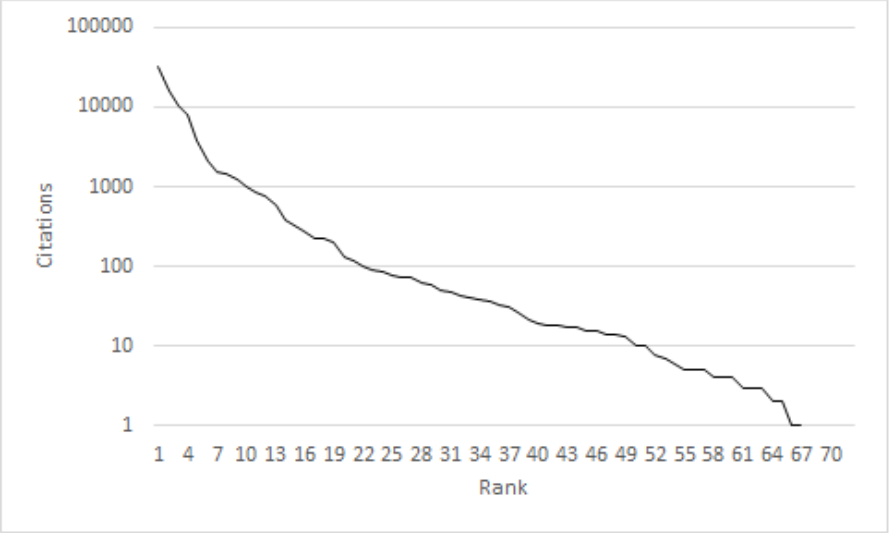

Figure 8: Citations of the papers selected for this literature review (log scale).

outline in (MacRoberts and MacRoberts, 1989), in addition to the differences in importance of their scientific contributions.

3.3.2 *Place Facets: Results of Extraction*

Using the extraction strategy, *116 place facets*[4] are extracted from the selected literature. For each facet the facet name, the facet's definition/description, an example of the facet use, and a list of publications where the facet is referenced are collected. The analysis of the reference list of extracted facets shows that *name*, *type*, and *location* are mentioned in all of the selected publications. Beyond the aforementioned three facets, the top-five frequently mentioned facets are *affordance* (e.g., Gibson, 1979; Jordan et al., 1998), *activity* (e.g., ElGindy and Abdelmoty, 2014b; Purves, Edwardes, and Wood, 2011), *meaning* (e.g., Relph, 1976; Tuan, 1977), *place identity* (e.g., Cresswell, 2004, 2014; Entrikin, 1991), and *sense of place* (e.g., Buttimer, 1976; Massey, 1994).

In several cases place facets with similar or even same names are defined with irreconcilable definitions. For example, *sense of place* is defined by Agnew (2011) as a part of a place conceptualisation, while Vanclay (2008) defined sense of place as personal feeling, relating to individuals rather than places themselves. Vanclay (2008) suggested to use the term of *spirit of place* instead of sense of place for describing how a place can evoke feelings. Hence, in the extracted list of facets two different facets are listed as sense of place. The same issue exists for *place identity*. While Relph (1976) defined place identity as "persistent sameness and unity which allows that

4 The list extracted facets are provided in Appendix A

thing to be differentiated from others", Vanclay (2008) has considered place identity as a part of his notion of sense of place, which is still a part of individuals' thought about a specific place. The issue of facets with the same name and different definitions is only observed for *sense of place*, *place identity*, *place attachment*, and *qualities*.

3.3.3 *Categories: Results of Execution*

In this section, the results of the execution phase are reported. First, individual categorisations are discussed and compared in a pairwise manner. Next, the similarity matrix derived from combining individual categorisations and the final categories are presented.

3.3.4 *Individual Groupings*

Table 6 shows the results of individual groupings. As shown in the table, the participants have different strategies to categorise the collected place facets. While second and third participants identified a top-level ontology from the collected facets, the first participant tried to integrate these facets in a group labelled as identity and purpose. The third participant's categories are more detailed compared to the other categorisations. Yet the similarity between these categories can be noticed through shared opinions about the spatial, physical and social groups of facets.

Pairwise comparisons of participants' groupings using the *Dice similarity coefficient* (Equation 1) are shown in Figure 9, Figure 10, and Figure 11. The comparisons are based on the groups derived from open card sorting, and consequently there are *one-to-one*, *one-to-many*, and *many-to-many* relations between the groups. Hence, in case that cardinality of the relationship is greater than one, the Dice similarity coefficient would be lower compared to one-to-one relationship case. However, this issue sources from the granularity in different groupings, and not a disagreement between the participants. For example, in pairwise comparison of first and second participants' categorisations, several one-to-many relationships are noticeable such as: first participant: {physical character} -related to- second participant: {form, location in space, boundaries/path, access}; first participant: {place in relation to human} -related to- second participant: {social, purpose/affordance}.

$$\mathtt{Similarity}(X, Y) = \frac{2 \times |X \cap Y|}{|X| + |Y|} \tag{1}$$

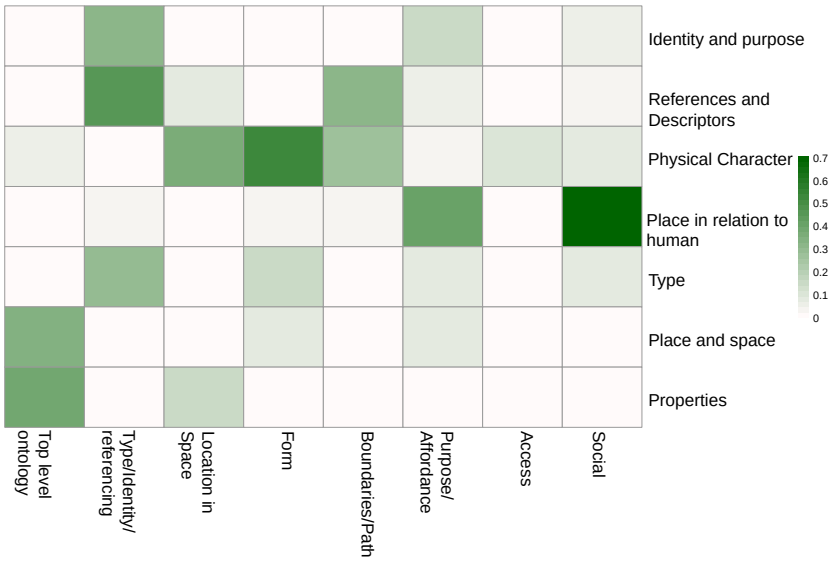

Figure 9: Pairwise comparison of grouping by first and second participants.

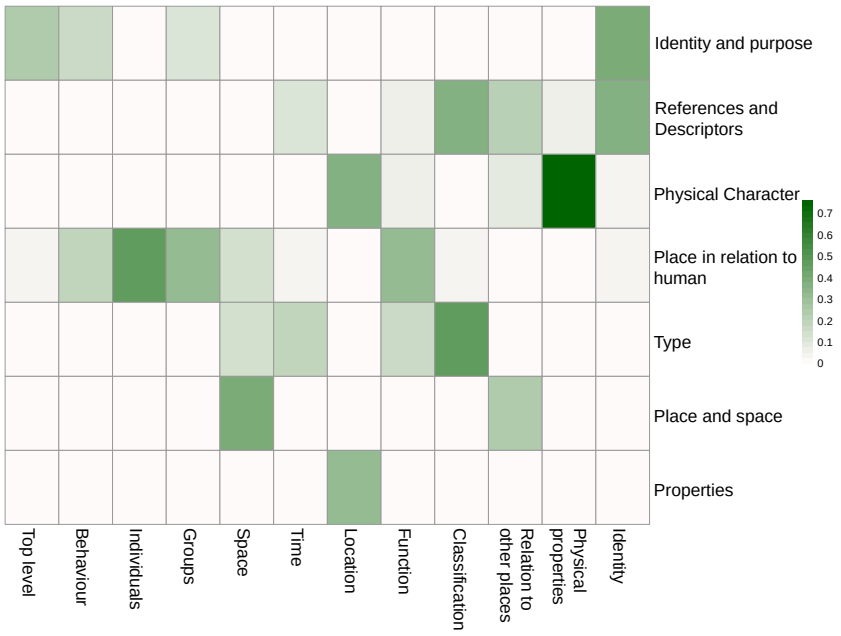

Figure 10: Pairwise comparison of grouping by first and third participants.

Table 6: Participants' groupings of the extracted facets

Participants	Group Name (Frequency)
First participant	Identity and Purpose (6)
	References and Descriptors (13)
	Physical Character (38)
	Place in Relation to Human (46)
	Type (8)
	Place and Space (3)
	Properties (2)
Second participant	Top Level Ontology (3)
	Type/Identity/Referencing (13)
	Location in Space (11)
	Form (19)
	Boundaries/Path (12)
	Purpose/Affordance (20)
	Access (2)
	Social (36)
Third participant	Top Level (2)
	Behavior (6)
	Individuals (14)
	Groups (10)
	Space (7)
	Time (3)
	Location (11)
	Function (15)
	Classification (9)
	Relation to Other Places (5)
	Physical Properties (25)
	Identity (9)

3.3.5 *Hierarchical Clustering*

Figure 12 shows the reordered similarity matrix based on categorisation agreement. The categorisation agreement is derived using the similarity values presented in Table 3. The agreement value shows how many participants agreed that i-th and j-th facets of the matrix belong to the same category. The reordered similarity matrix shows a strong clusters of facets that can form categories of

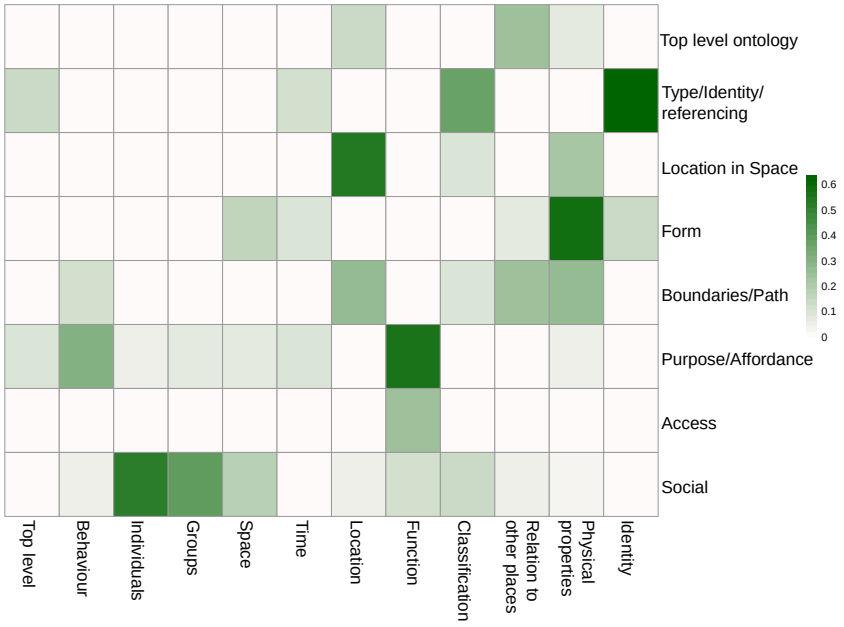

Figure 11: Pairwise comparison of grouping by second and third participants.

facets based on the participant opinions. Later, this similarity matrix is used to derive hierarchical clustering that combines participant opinions.

Using the agglomerative clustering, a set of integrated hierarchical categories are derived. Figure 13 shows the hierarchical categories derived from the dendrogram of the hierarchical clustering.

3.4 A MULTI-FACETED MODEL OF PLACE

The multi-faceted model of place is the results of the interpretation of the hierarchical clustering (as shown in Figure 13). This model of place differentiates three types of facets in the coarsest level of hierarchy: *primitive facets*, *derived facets*, and *linguistic facets*.

The rationale behind this typology is the fact that while 'primitive' facets are about particular aspects of place (e.g., its relation to people), 'derived' facets have mixed meanings in a way that they can be derived from a combination of primitive facets. On the other hand, linguistic facets such as *place names* have different objectives (i.e., referring and describing) compared to primitive and derived facets. Consequently, due to their differences as references to places compared to properties of places, they should be categorised in a different category. In the reminder of this section, the primitive, derived, and linguistic facets are discussed.

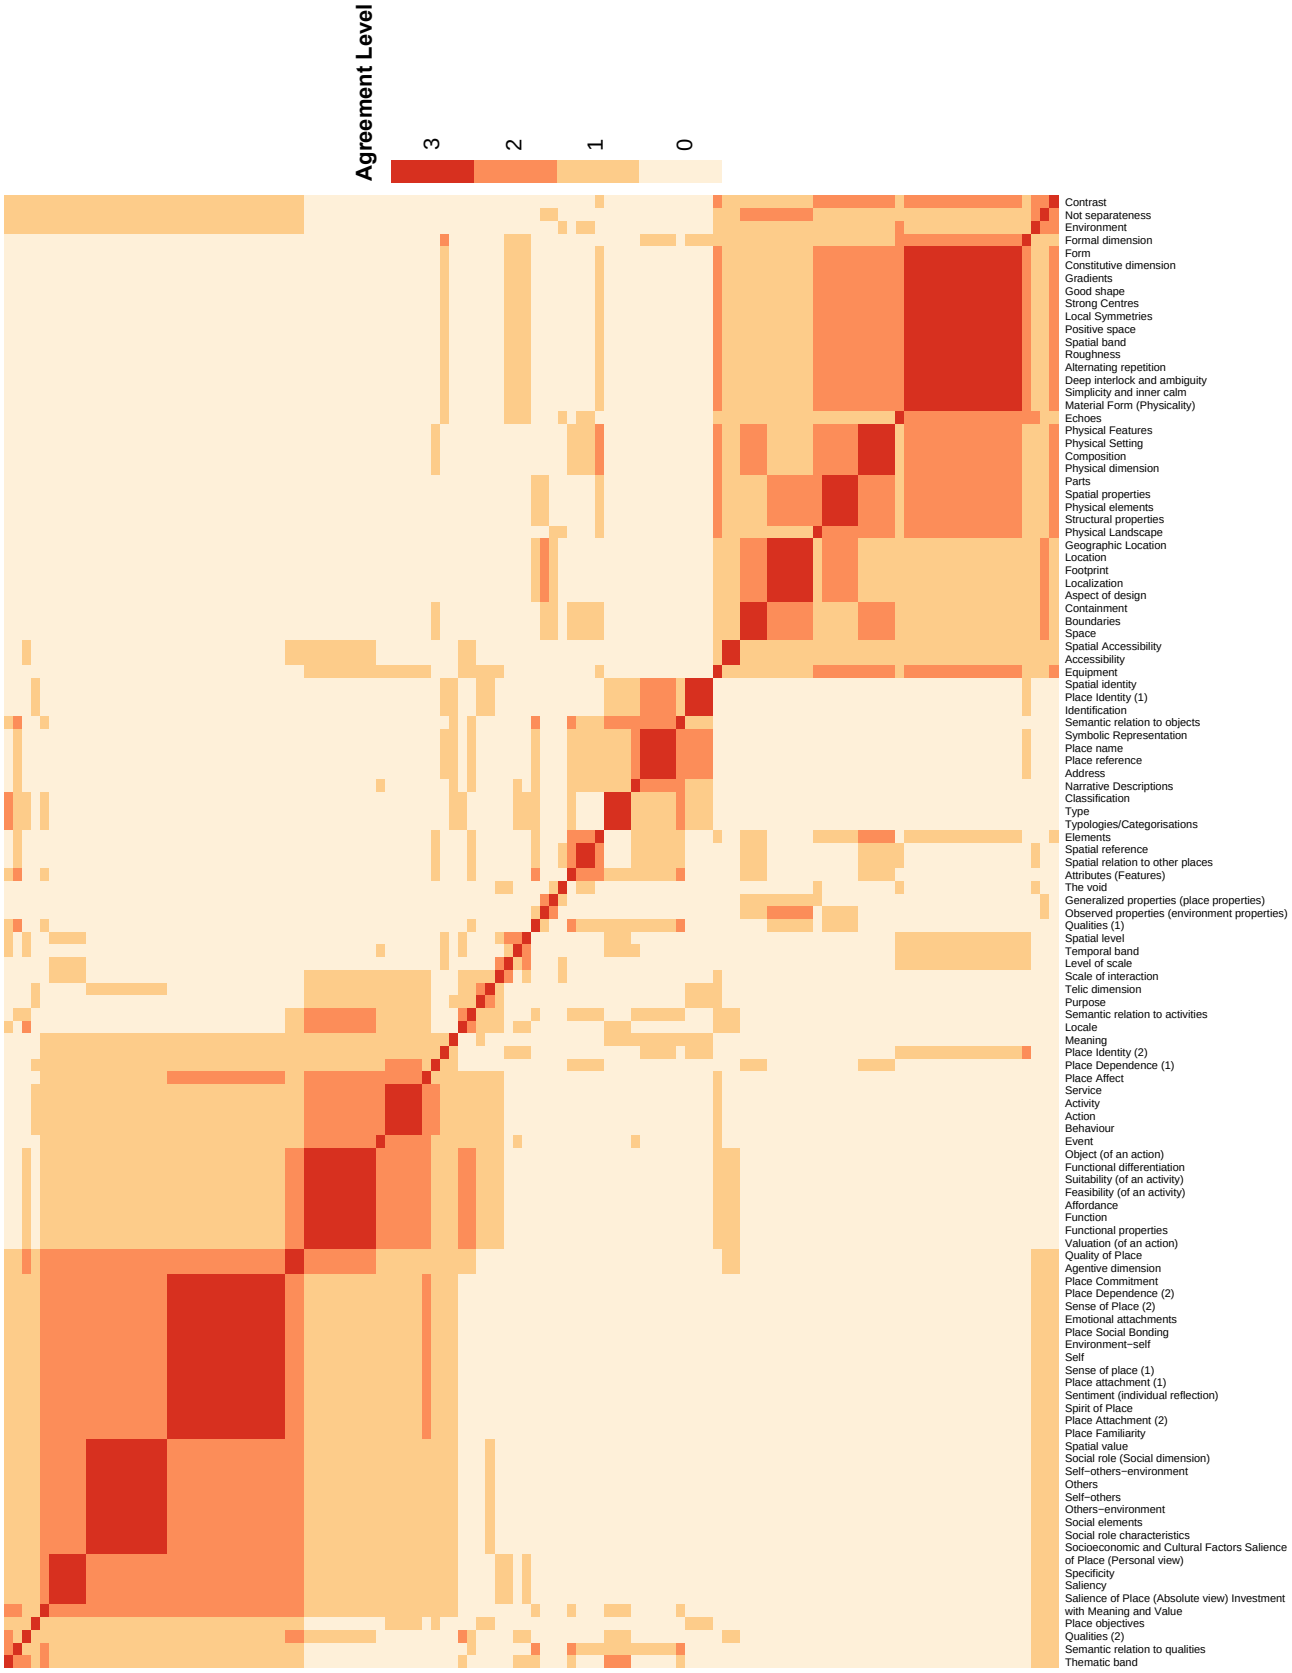

Figure 12: Ordered similarity matrix and its relation to the proposed categorisation.

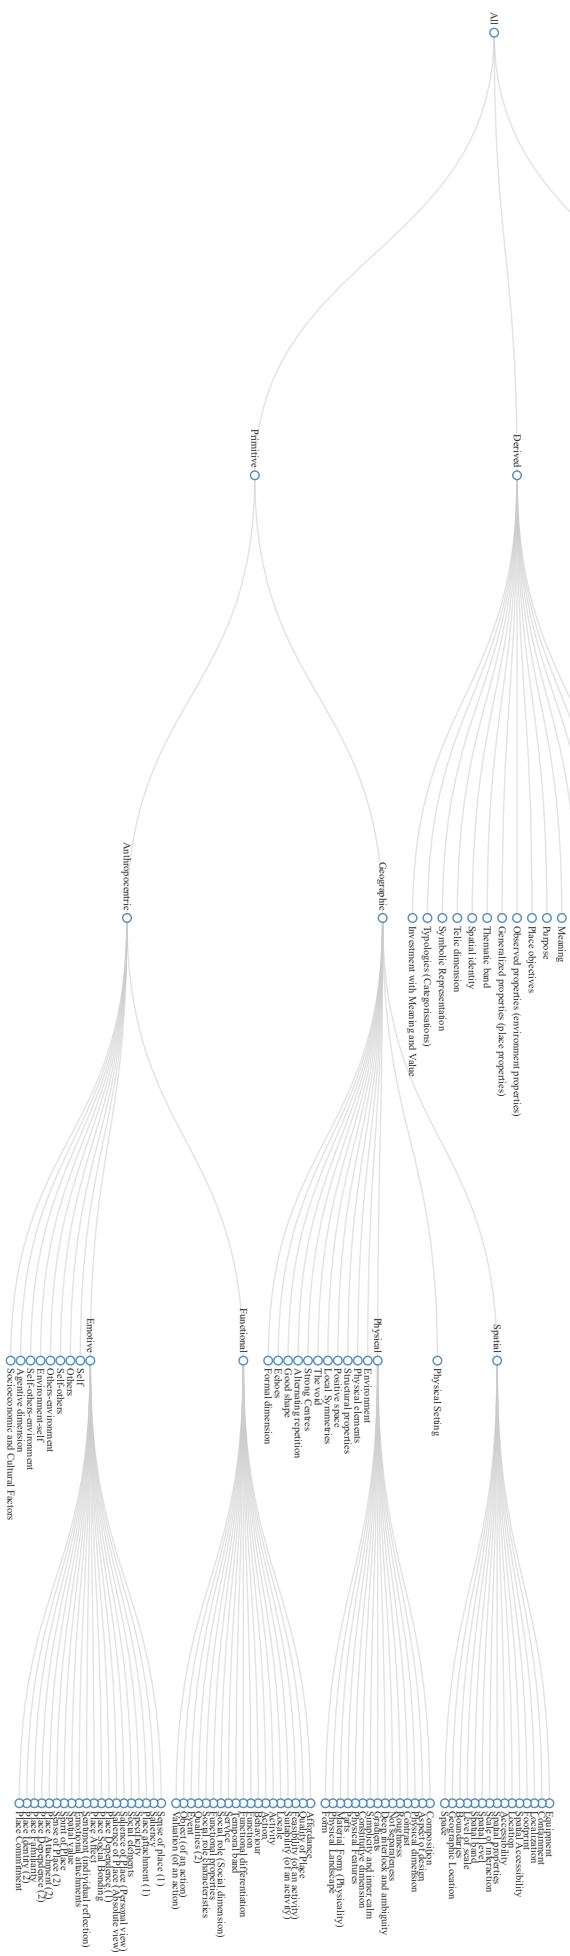

Figure 13: Hierarchical categories of place facets

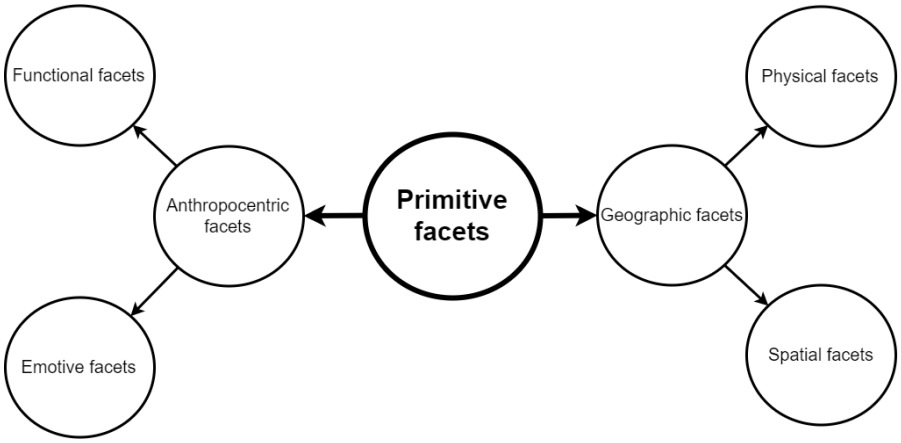

Figure 14: The result of hierarchical clustering of primitive facets.

3.4.1 *Primitive Facets*

The two upper levels of categorisation for primitive facets are shown in Figure 14. In the top level, the primitive facets are divided into *anthropocentric* and *geographic* groups. Then the former group is divided into *emotive* and *functional* groups and the latter is split to *physical* and *spatial* groups.

The defining distinction between *anthropocentric facets* and *geographic facets* (Figure 14) arises from the perspective on the notion of place. The anthropocentric group contains facets that capture the relationships of individuals, or groups of people with places; while the geographic group captures facets describing spatial and physical properties of places. In addition, this division reflects the discipline-specific perspectives on place: while selected publications in human geography and social science literature are usually primarily studying facets in the anthropocentric group, publications selected from architecture and geospatial science primarily (but not exclusively) focus on facets in the geographic group.

3.4.1.1 *Anthropocentric Facets*

The relation between people and places underpins this grouping of facets. In her ontologies of geographic information (Couclelis, 2010), Couclelis named the facets of an agent-environment relationship as agentive dimension. In Gustafson (2001)'s conceptualisation, *self*, *other*, *self-other*, *self-environment*, *other-environment*, and *self-other-environment* are all related to the facets in this category. According to Gustafson's definition, the difference between the *self* and the *other* is the difference between individual's feelings and thoughts, and social roles and functions. Here, the sub-division between the individual and the social perspective is made after differentiating anthropocentric facets into functional facets, and

emotive facets. Functional and emotive facets are differentiated by the distinction between *what potential activities a place can afford* as opposed to *what emotional attachments and feelings it triggers amongst people*.

ANTHROPOCENTRIC FUNCTIONAL FACETS This category can be further sub-divided into two sub-categories: affordance and activity, and function. The contrast rule that is used for structuring this category to sub-categories stems from differentiating between the relation of people and places and (social) groups of people and places. While affordances are referring to individual-based answers to what a place can afford; function is defined as a less subjective answer to this same question (Papadakis, Resch, and Blaschke, 2016).

The term *affordance* was first introduced by Gibson (1979) to describe the relation between an agent and an environment, defining affordance as a perceivable action offered by an environment to an agent. While the concept of affordance was conceptualised for generic notions of environment and agent; it fits well for describing the potential interactions of a person and a place (Jonietz and Timpf, 2015). Affordance is thus one of the ways to explore the relation of people and places – i.e., investigating a place as somewhere in the world where specific *actions* and *behaviors* are afforded. The term affordance has been used as a component of place models (e.g., Alazzawi, Abdelmoty, and Jones, 2012; Jordan et al., 1998) and formal ontologies (e.g., Ballatore, 2016; Scheider and Janowicz, 2014). The affordance itself has been investigated and modelled, and ontologies have been proposed to capture affordance information (e.g., Galton, 2010; Ortmann and Kuhn, 2010).

Activity and *locale* are terms with relatively similar meaning to affordance in the context of place. While these terms have been observed less frequently in the literature compared to affordance, they are still amongst the top-ten frequently mentioned facets. While locale and affordance have different meanings in common language, 'locale' is defined and used interchangeably with affordance by some scholars (Agnew, 2011; Purves and Derungs, 2015). Unlike locale, the meaning of 'activity' is differentiated from affordance in the literature. Activity and affordance are different, reflecting what action is *actually performed* in a place, as opposed to what actions are *possible or supported* by a place. Researchers argue that the difference between activity and affordance is related to whether the actions are performed or only potentially achievable (e.g., Ortmann and Kuhn, 2010; Scheider and Janowicz, 2014). Note, however, that these terms are often used interchangeably and readers are encouraged to critically evaluate what authors mean (e.g., Alazzawi, Abdelmoty, and Jones, 2012). Models and ontologies designed for activity (e.g., Das and Winter, 2016; Kuhn, 2001) often manifest similar design

patterns with models and ontologies of affordance, e.g., complex affordances or activities are modelled as a combination of more primitive affordances or activities (Das and Winter, 2016; Kuhn, 2001; Ortmann and Kuhn, 2010).

Action and *behavior* belong to the sub-category of affordance within the functional facets group. Jonietz (2016) discussed the relation of activity, action, and behavior based on definitions from Werlen (2003). The term *behavior* denotes activities that are related to mechanical movement, while actions are the activities that are intentional and goal-oriented (Jonietz, 2016; Werlen, 2003). Furthermore, Jonietz (2016) formalised activities by considering *feasibility* of an activity in a place in conjunction with taking the *suitability* of the place for the activity into account. Moreover, he argued that *quality of a place* is directly influenced by feasibility and suitability of place for different activities (Jonietz, 2016; Jonietz and Timpf, 2015). Analogically to Jonietz (2016), Hockenberry (2006) formalised the notion of place by modelling activities, by considering *valuation* and *object of activities* for capturing information about places. For example, eating as an activity which can be performed in a restaurant is supported by objects such as meal or food, and the valuation is modelled as a rating which defines how well the activity is performed in the restaurant from a subjective view of a person (Hockenberry, 2006).

Function (Markus, 1987; Papadakis, Resch, and Blaschke, 2016; Stefano De, Stefano, and Tumasch, 2015) is a facet of place describing what a place can afford to a group of people (Papadakis, Resch, and Blaschke, 2016). Hence, function, as a facet of place, conveys shared meaning compared to activities and affordances. Several scholars have used other terms, such as *service* (Alazzawi, Abdelmoty, and Jones, 2012; Winter and Truelove, 2013), *functional properties* (Vasardani, Tomko, and Winter, 2016), *functional differentiation* (Canter, 1997), and *event* (Roche, 2015) with relatively equivalent meanings to function. *Social role*, and *social dimension* of place are other terms used to describe impacts and contributions of places on social groups (Ballatore, 2016; Calafiore, 2016). For example, a shopping mall as an urban place provides a social role of trading to people in the society (Calafiore, 2016).

One way to model functions and social roles of places is to investigate the interaction between people and place in terms of time (Adams and Janowicz, 2015; Janowicz, 2012; Janowicz, Scheider, and Adams, 2013; Ye et al., 2011). The *temporal band*, as defined by Ye et al. (2011), relates to temporal patterns of visiting specific places for particular activities. Accordingly, the functional differentiation of places can be derived from temporal patterns, e.g., the temporal pattern of visiting a restaurant to eat is different from the temporal

pattern of going to a bar based on days of the week, and hours of the day (McKenzie and Adams, 2017; Ye et al., 2011).

ANTHROPOCENTRIC EMOTIVE FACETS The shared meanings of numerous facets relating places to people's feelings and emotions underpins the differentiation of the group of *emotive facets*. By studying the definitions of these facets, a noticeable differentiation is identified among them – i.e., whether the purpose of the facet is to capture place-related emotions as a purely subjective (individual) relationship, or to capture a more shared, objectivisable relationship between the place and a group of people. Hence, similarly to the division of functional facets, emotive facets can be sub-divided into two groups: 1- attachment and sense of place, capturing the subjective emotional relationship between a single person and a place (e.g., home for the owner); and 2- theme, social bonding, and cultural attachment, with its facets capturing emotional bonding of a group of people to a place (e.g., country for citizens).

Sense of place and *attachment* are the main facets of the subjective subgroup of the emotive category. While sense of place captures *how* a place feels to an individual (Agnew, 2011; Cresswell, 2004), attachment is related to *positive feelings and bonding* of a person to a particular place (Scannell and Gifford, 2010). Hence, attachment and sense of place contrast: attachment captures positive feelings, while sense of place describes a less structured range of emotions from positive to negative through neutral or hard to categorise ones. This underspecification of these two facets leads to a lack of solid agreement on the relationship between these two facets. For example, the definition of sense of place by Vanclay (2008) encompasses *attachment*, *dependence*, *familiarity*, *commitment*, and even *place identity*; while an overlap of meanings is observable between attachment and sense of place by Scannell and Gifford (2010).

In the literature, alternative terms have been proposed for describing emotional relationships between a person and a place such as *spirit of place* (Vanclay, 2008), and *sentiment* (ElGindy and Abdelmoty, 2014a,b) related to sense of place, and *topophilia* (Tuan, 1990), *emotional attachment* (Mennis and Mason, 2016), *bonding* (Alazzawi, Abdelmoty, and Jones, 2012; Kyle, Graefe, and Manning, 2005), and *place affect* (Kyle, Graefe, and Manning, 2005) similar to the attachment. Consequently, the emotional bonding and attachment of a person to a place can be related to *saliency* in a personal view (Almuzaini, 2017). The impact of attachment in saliency, as the importance of a place to an individual, have been studied and reported in several studies (e.g., Jorgensen and Stedman, 2001; Kyle et al., 2004).

Several scholars have investigated the emotional relations between society and places by generalising the facets of attachment and sense

of place from individuals to higher social groupings. They thus investigated shared feelings between a group of people to specific places. The terms proposed are *spatial value* (Quesnot and Roche, 2015), as the value of a location to a society, *social elements* (Scannell and Gifford, 2010), and *specificity* (Scannell and Gifford, 2010). Spatial value was defined as the total value of a location to people (Lussault, 2003; Quesnot and Roche, 2015). In the same view, the terms *social elements* and *specificity* were defined as descriptions of social prominence of a place. They have considered a direct relationship between prominence and people's feeling to a place (Scannell and Gifford, 2010). While the term *social elements* is defined as prominence of a place at societal levels, *specificity* is denoted as a measure of how a place is differentiated from other places with respect to people's attachment to the place (Scannell and Gifford, 2010).

Social bonding and cultural attachment lead to *prominence* and *saliency* (Almuzaini, 2017) of places. One way to think of saliency is to view it as an answer to how successful a place is in initiating emotional bonding of people to the place, compared to other places (Almuzaini, 2017). While saliency is categorised as an emotive facet, the cause of a place's saliency can be rooted in other facets of the place, such as physical attractiveness (e.g., aspect of design for a church), functional aspects (e.g., the function of a hotel), and even spatial properties (e.g., height for a building) (Nothegger, Winter, and Raubal, 2004).

3.4.1.2 *Geographic Facets*

This category, the other top-level category in Figure 14, encompasses physical, and spatial properties of places. The facets in this group are related to *where* a place is, and what does a place *look like*. The category of geographic facets comprises characteristics of place described as *environment*, *physical setting*, *physical elements*, *structural properties*, and *formal dimension* in the conceptualisations of place proposed by Gustafson (2001), Relph (1976), Scannell and Gifford (2010), Vasardani and Winter (2016), and Couclelis (2010), respectively.

The group of geographic facets can be further divided based on the concepts underpinning its facets –i.e., whether the facet is related to form and material properties such as *physical features* and *form* (e.g., Jordan et al., 1998; Vasardani, Tomko, and Winter, 2016; Vasardani and Winter, 2016), or associated with spatial properties such as *location* and *boundary* (e.g., Montello et al., 2003; Vasardani and Winter, 2016; Winter and Freksa, 2012). Several facets, such as *positive space*, *local symmetries*, *the void*, *alternating repetition*, *good shape*, *strong center*, and *echoes* which are all parts of the place properties model (Vasardani, Tomko, and Winter, 2016; Vasardani

and Winter, 2016), cannot be assigned to one of the sub-categories due to their mixed spatial and physical meaning. For example, *local symmetries* can be evaluated based on form, from an architectural perspective, and at the same time *local symmetries* may relate to particular spatial properties in spatial information sciences. The aforementioned facets therefore belong to this level of categorisation and can currently not be sub-divided into finer levels.

GEOGRAPHIC PHYSICAL FACETS Physical facets have been referred with different terminologies such as *physical dimension* (Ballatore, 2016), *physical features* (Jordan et al., 1998), *material form* (Gieryn, 2000), and *aspect of design* (Canter, 1997). Physical facets can be defined as any type of information related to structure, material, form and architectural style of a place (Ballatore, 2016; Vasardani, Tomko, and Winter, 2016). Physical facets thus capture *what a place looks like based on the common senses*. While the question can be answered in different ways, two noticeable types are observed in the literature: referring to style and form, and structure and parts. For example, the possible answers to the specific question: "What does your home look like?" can be answered by relating to style and form: "It is an old Victorian building"; or in terms of parts and structure: "It is spacious, it has three bed rooms, a hall, and a big kitchen".

Form (Canter, 1997; Markus, 1987) is a facet describing how a place looks like with respect to architectural and morphological aspects. One of the properties of place is known as *not separateness* (Alexander, 2002; Vasardani and Winter, 2016). The *not separateness* assures the coherence of the whole (Alexander, 2002). While *not separateness* is not only about morphological aspects of a place, it can be applied to the notion of form. The previous example, "an old Victorian building", is grounded in 'not separateness' – i.e., in a way that the *home* is referred to as a single object (a whole) with shared properties that apply to the whole: the age of the building, and its Victorian architecture. *Roughness* (Alexander, 2002; Vasardani and Winter, 2016) is another facet relating to form. It is defined as imperfections in a way that one can notice similarity but not identicality of form within a place as a whole (Alexander, 2002) – e.g., a wooden fence around a suburban house constructed from woods with different colors and different wood types, but captured and perceived as a whole.

Parts (Tversky and Hemenway, 1983), *constitutive dimension* (Couclelis, 2010), and *composition* (Papadakis, Resch, and Blaschke, 2016) are facets with similar meanings related to the internal structure of a place. Thus, a place be considered as a container of other places (Scheider and Janowicz, 2010) (e.g., a home has rooms, and a country has states or provinces), and consequently what does a place look like can be answered through references to its parts.

Despite using *not separateness* for determining a place based on form, here *contrast* and *gradient* (Alexander, 2002; Vasardani and Winter, 2016) are used for describing a place based on its parts. *Contrast* is defined as variation of properties between elements or parts of a place, and *gradient* is referred to gradual changes which one can observe in the structure of a place from one part to another (Alexander, 2002). Analyzing a place based on its part can be evaluated through *simplicity and inner calm*, and *deep interlock and ambiguity* (Alexander, 2002; Vasardani and Winter, 2016). A place in which its parts can be simply linked and understood is a coherent whole supporting the criteria of *simplicity and inner calm*, and on the other hand, when the parts interpenetrate in linking together, it is described as *deep interlock and ambiguity* (Alexander, 2002).

GEOGRAPHIC SPATIAL FACETS Spatial facets anchor a place in space. *Spatial band* (Janowicz, Scheider, and Adams, 2013), *spatial properties* (Vasardani, Tomko, and Winter, 2016; Vasardani and Winter, 2016), and *space* (Markus, 1987) are generic terms proposed for any type of properties describing relationships between places and space. These spatial facets can be further divided into sub-groups that can be used to investigate 1- a place located in space, 2- the meaning of place in relation to space, and 3- the spatial relationships between places. While the first sub-group of facets describes places in isolation and captures mostly *at which locations the place is found* (e.g., Agnew, 2011; Goodchild, 2011), the second sub-group refers to facets that focus on generic relation of place and space such as *scale of interaction* (Canter, 1997). Finally, the third sub-group includes facets about spatial relationships between places, often conceptualised as objects, such as *containment* (Scheider and Janowicz, 2010; Scheider and Purves, 2013; Winter and Freksa, 2012), and *equipment* (Scheider and Janowicz, 2014).

Location is one of the facets mentioned in all publications. Locational information of a place is the answer to *where-questions* that are asked in a wide range of situations, from our everyday life to human-machine interactions. It is therefore no surprise that locational information is the primary part of different place models. *Localisation* (Scheider and Janowicz, 2014), *footprint* (Goodchild, 2011), and *geographic location* (Gieryn, 2000) are terms related to location. Localisation is closely associated with another facet of place, i.e., *boundary* (Vasardani, Tomko, and Winter, 2016; Vasardani and Winter, 2016; Winter and Freksa, 2012). Based on the Jordan curve theorem, a boundary partitions space into three segments, the boundary itself, an inside, and an outside (Hales, 2007). However, unlike mathematical geometries, places in the geographical world do not necessarily have well-defined, crisp boundaries (Montello et al., 2003; Winter and Freksa, 2012). Unlike administrative places which

are defined by crisp boundaries, people have vague and subjective perceptions of boundaries for the socially-constructed places such as downtown (Hollenstein and Purves, 2010; Montello et al., 2003; Smith and Varzi, 2000), for natural places such as a mountain (Jones, 1959), and in natural communication where boundaries are often irrelevant.

The second subcategory of spatial facets are related to the notion of scale. *Scale of interaction* (Canter, 1997), *spatial level* (Scannell and Gifford, 2010), and *level of scale* (Vasardani, Tomko, and Winter, 2016; Vasardani and Winter, 2016) are terms for capturing the meaning and even the existence, of the scale of places. Due to aggregation and abstraction, places do exist only in certain scales or levels of detail. For example, when investigating global trading behavior places such as shopping malls or grocery stores are irrelevant, even though they are perfectly matched with the functional purpose of the study. In addition, people's bonding to places is scale dependent: There seems to be a U-shaped relationship between the strength of attachment and scale of interaction (Lewicka, 2010). In other words, (Lewicka, 2010) has found that for five types of places (i.e., apartment, building, neighborhood, city district, and city) the attachment to apartment and city are higher than the building, neighborhood, and district.

Due to the importance of scale for the analysis of places, multiple scale schemas were developed to formalise qualitative ordinal progressions of scale through a finite set of levels grounded in human relationships with objects and environments (e.g., Montello, 1993; Richter et al., 2013). For example, the schema of scale presented by (Richter et al., 2013) includes seven levels – from the finest level of furniture, through room, building, street, district, city, and finally with the coarsest level country.

The last sub-category of spatial facets is related to spatial relationships between places. These facets are not properties of an isolated place but focus on the relationships of places with each other. *Containment* (Scheider and Janowicz, 2014; Winter and Freksa, 2012) and other mereological (part-of) and topological relationships, other qualitative relationships (e.g., nearness, cardinal directions), and metric relationships (e.g., distances) belong to this category (Kim, Vasardani, and Winter, 2016; Vasardani et al., 2013). In addition, *equipment* (Scheider and Janowicz, 2010, 2014) is categorised as a spatial facet here. Equipment of a place determines what are the things associated with the place, in terms of spatial relationships, often in form of things contained in/at a place with a purpose to facilitate specific activities (Scheider and Janowicz, 2010, 2014). As an example, a laboratory at the University is equipped with computers.

Spatial properties of a place with respect to other places project into a facet known as *accessibility* (Ingram, 1971; Winter and Freksa,

2012), or *spatial accessibility* (Carmona et al., 2012; Jonietz, 2016). In geography, accessibility has often been evaluated by mathematical methods based on mutual spatial relationships of places. Ingram (1971) distinguishes two types of accessibility, 1- the relative accessibility which is a measure related to accessibility of a location in space with respect to another location, and 2- the total, or integral accessibility of a location in space which is the integration of relative accessibilities to all other locations. Despite the sophisticated mathematical formalisations of accessibility, people's choices for movement from one place to another are also influenced by their experiential knowledge of places' accessibility (Mondschein, Blumenberg, and Taylor, 2010). Due to importance of accessibility in human experiences of space, Jonietz (2016) considered accessibility of places as one of the major factors in evaluating their qualities.

3.4.2 *Derived Facets*

Derived facets are the result of integration of primitive place facets. A facet such as *meaning* (e.g., Buttimer, 1976; Canter, 1997; Gustafson, 2001) can be interpreted in a multitude of ways, for example as its functional role for a person, or a society, or as the strength of attachment for the place with regard to a person, or a part of the society. The part of meaning that is specifically related to people-place relations has also been called *investment of meaning and value* (Gieryn, 2000; Vanclay, 2008). However, physical and spatial facets of a place all have impacts on meaning. Physical facets capture the visual attractiveness and distinctiveness of a place, and accessibility and geographic location determine how easy it is to be inside or nearby of the place. Hence, when place is defined as part of space with shared meaning (Cresswell, 2004), meaningfulness is derived from a combination of primitive facets related to objective and subjective nature of the place. For example, the meaning of a 'home', a 'restaurant', a 'mountain', and a 'path' are related mostly to their emotive, functional, physical and spatial facets, respectively. In addition, sometimes, specifically for landmarks, the meaning of a place may be related to its *symbolic representation* (Jordan et al., 1998), e.g., the Statue of Liberty as an icon of freedom.

Identity (e.g., Entrikin, 1994; Relph, 1976; Tuan, 1977) is another derived facet frequently mentioned in the literature. Shared identity of a part of space enables the emergence of a place in a process of place-making. *Spatial identity* (Quesnot and Roche, 2015), and *identification* (Scheider and Janowicz, 2014) are other terms with equivalent meanings to identity. Identity and meaning relate to each other like cause and effect. For a specific part of space, shared identity leads to a specific meaning, and same meaning creates a particular sense of identity (Entrikin, 1996). Using the

aforementioned examples provided for *meaning*, the relation between meaning and identity are discussed. The shared identity of a 'path' is mostly rooted in the spatial configuration of connecting two other places. Similarly, for a mountain the shared identity is primarily related to morphological properties and physical settings which lead to contrasts of shape and form with its surroundings. In the same way, a restaurant has a shared identity noticeably grounded in its function – providing food to customers – compared to a nearby school. Finally, a person's home has different identity compared to other buildings in terms of feeling safe, secure, and comfortable. As a result, place is also defined as a part of space with shared identity (Relph, 1976).

In the same way, *telic dimension* (Couclelis, 2010), *place objectives* (Canter, 1997), and *purpose* (Papadakis, Resch, and Blaschke, 2016) are categorised as derived facets. These facets have a similar meaning with respect to objectives that people have with a place. Consequently, they can be derived from a mixture of emotive, functional, physical, and spatial facets. For example, the Avenue des Champs-Élysées in Paris may have different purposes from the perspective of a French citizen, a tourist, a refugee, an urban planner, and a foreign worker, all at the same time.

Derived facets of place are also used in place formalisations and models. *Type* (e.g., Goodchild, 2011), *typology* (Jordan et al., 1998), *classification* (Scheider and Janowicz, 2014), and *thematic band* (Adams and Janowicz, 2015) are the facets proposed in the different place models for a similar purpose. These facets are defined to differentiate generic kinds of places from each other. Consequently, street, house, restaurant, and even country are valid examples of place types. There is a direct relationship between the type of place and its purposes and goals it affords to a person or a society. Hence, a mixture of emotive, functional, physical, and spatial facets ground the characterisation of a place as a particular *type of place*.

3.4.3 *Linguistic Facets*

Linguistic facets are the manifestations of place in language. The difference between linguistic facets of place and the primitive and derived facets is similar to the difference between a reference to a place with the place itself. Linguistic facets enable us to communicate and to share our knowledge of places, and perhaps this is the reason why in the literature they are often confounded with the characteristics of the places themselves. This group of facets are the lexicon to establish annotation schema to study place in verbal utterances.

Place is a special case of a broader notion of named entities in language. Named entities includes abstract and physical entities that

can be referred to by names (Nouvel, Ehrmann, and Rosset, 2016). In models and formalisations of place, names are called *place names* or *toponyms*. Toponyms are direct references to places. Even such direct references to places are, however, often ambiguous (Leidner, 2007). The ambiguity of toponyms stems from the facts that a place can have multiple names, including vernacular and historic names, and that a single place name can be used to refer to different places (Leidner, 2007).

Verbal reference (Montello et al., 2003; Purves and Derungs, 2015; Purves, Edwardes, and Sanderson, 2008; Winter and Freksa, 2012) is a linguistic facet of place that includes direct references by toponyms as well as complex place descriptions (Winter and Freksa, 2012). For example, *Mitchell Library* and *the library on Macquarie Street in Sydney* are both valid references to the same place, one by place name and one by place description. The primitive types of verbal references have been discussed in a number of publications (Edwardes and Purves, 2007a; Edwardes and Purves, 2007b; Scheider and Purves, 2013). These publications have defined *semantic relation to objects* (also known as *elements* (Edwardes and Purves, 2007a; Edwardes and Purves, 2007b)), *semantic relation to activities*, *semantic relation to qualities*, *spatial relation to other places* as fundamental ways a place can be described and even localised via language (Edwardes and Purves, 2007a; Edwardes and Purves, 2007b; Scheider and Purves, 2013). Such linguistic facets are related to physical, functional, emotive, and spatial categories of facets, respectively. These primitive types of verbal references can be combined in complex verbal references known as *narrative descriptions* (Entrikin, 1997; Jordan et al., 1998) including place descriptions (Vasardani et al., 2013), route descriptions (Winter et al., 2018), and destination descriptions (Tomko and Winter, 2009).

Spatial references (e.g., Vasardani et al., 2013; Winter and Freksa, 2012), also called *spatial relations to other places* (Scheider and Purves, 2013), are specific types of verbal references. In a spatial reference, the location of a place is described with respect to other places using qualitative spatial relationships, including topological, directional, and qualitative distance relationships (Egenhofer, 1991; Frank, 1992). A specific type of spatial references are *addresses* (Almuzaini, 2017). Western-style addresses describe the location of a place in space in terms of hierarchical (containment) relationship. Extracting and formalising spatial references from textual information have been studied by several scholars. The proposed models for capturing the information in spatial references are built-upon three elements which are known as a triplet –i.e., locatum, spatial relationship, and relatum (Kordjamshidi, Otterlo, and Moens, 2011; Vasardani et al., 2013).

3.5 DISCUSSION

The differences between people's knowledge of places, and the information stored in computers has been investigated based on the differences between the qualitative configurational knowledge of environments captured in human place-based knowledge, contrasting with quantitative measurements stored as geometrical coordinates in computers (e.g., Vasardani et al., 2013; Winter, 2009; Winter and Freksa, 2012). Importantly, these differences can be viewed as the contrast between the subjective nature of human experiences in terms of physical, functional, and emotive aspects and the purely spatial, reductionist views manifested by maps. The gap between place in information systems and people's experiential knowledge increasingly leads to unsatisfactory and troublesome interaction of human and machines, e.g., in geographic Web search (Ballatore, 2014). Hence, computational models for emotive, functional, and physical facets of place present an important direction to bridge the gap.

Modelling places, i.e., whether place can be modelled by information science or not, is a current topic of debate in GIScience (Goodchild, 2011; Merschdorf and Blaschke, 2018; Purves, Winter, and Kuhn, 2019). While some researchers have considered place as a concept with intrinsic complexity that cannot be captured and formalised unless in case of very narrowly defined circumstances (Goodchild, 2011), recently Purves, Winter, and Kuhn (2019) developed a grounding to formalise places in information science, based on core concepts of spatial information (Kuhn, 2012). The multi-faceted model of place presented in this study which can be used to extend that grounding (Purves, Winter, and Kuhn, 2019) to fit a specific purpose. In other words, the facets, extracted from the literature and categorised in this chapter, can be used as a list of terminologies to extend the current grounding to fit for different applications. In the next chapter, this model is used to classify questions using the core concept of spatial information.

Current interactions about place between humans and computers (e.g., geographic Web search) are still limited to gazetteer information lookup (*name*, *type*, *location*), however, research can progress by exploring the importance of multi-faceted perspective on place for place-related human computer interaction presented here. Rich corpora of verbal descriptions of places available from diverse sources such as social media and blogs can be used for improving the current state-of-the-art in place-related human-computer interaction. The extraction, formalisation, and processing of place information is therefore a promising direction that can benefit from the multi-faceted model of place. The extracted

platial information can be used in computers to imitate and support communication about places.

Another observation of this study is the notable capability of language to convey meaning of primitive and derived facets of place. In contrast to map visualisations which are designed to convey spatial information, language is a channel to express emotive meaning, or to talk about functional and physical characteristics of places (Derungs and Purves, 2014). A verbal place description (e.g., *a fantastic Gothic church in the South of France*) can include references to diverse facets of places, while maps are limited to purely spatial features. However, linguistic facets are yet underexplored. For example, linguistic facets describing tactile, sensual, and auditory experiences of place (e.g., *a quiet place*) have not been studied sufficiently in comparison to visual experiences (Chesnokova et al., 2018). Hence, the linguistic facets can be used and extended for developing a schema for extracting platial information from place descriptions. Evidently, further development in language processing techniques is the primary prerequisite for this purpose. In the next chapter, the linguistic facets are used as a lexicon to develop an encoding schema to study place-related questions and the corresponding human-generated answers.

3.6 CONCLUSION

The notion of place as a multidisciplinary topic of interest has been investigated in the literature through different disciplines and focal perspectives. Consequently, different terminologies, conceptualisations and models of place have been proposed so far. In this chapter, the notion of place in geography, social science, architecture and GIScience is reviewed, and the results are reported in form of a multi-faceted model place. Using the concept of facets, these terminologies, conceptualisations and models have been extracted from the literature. In addition, hierarchical categories of facets are derived using a card sorting game and hierarchical clustering.

The extraction of place facets from the literature leads to a set of more than one hundred facets. The extracted facets have been categorised into primitive, derived, and linguistic facets. The primitive facets have been systematised hierarchically. In the first level of subdivision, primitive facets were categorised into anthropocentric and geographic facets. The anthropocentric facets were then sub-divided into emotive and functional facets, and geographic facets were sub-divided into physical and spatial facets. Furthermore, the list of facets in emotive and functional categories was discussed based on their capture of the individual vs. social relationship with places. The physical facets have been divided into

facets related to style and form, and structure and parts. Spatial facets have been further discussed in terms of their capture of spatial properties, generic relations to space (scale), and spatial relations between places.

The relation between linguistic facets, derived facets, and primitive facets has been introduced and discussed, outlining the reasons for the distinction of linguistic facets in particular. The linguistic facets provide a lexicon to establish schema for analysing place-related verbal utterances.

The collected facets and proposed multi-faceted model place are used in the following chapters as a grounding to design models to investigate place-related questions and their answers. In Chapter 4, the multi-faceted model is used to design a top-down classification of place-related questions. In Chapters 4 and 5, the linguistic facets are used as a building lexicon to design an encoding schema to study a corpus place-related questions and answers. Finally, in Chapter 6, the collected generic facets such as scale, prominence and type are used to study human answering behaviour to where-questions.

4

UNDERSTANDING PLACE-RELATED QUESTIONS

Theory without data is groundless,
but data without theory is just uninterpretable.

— Immanuel Kant

4.1 INTRODUCTION

In their everyday communication people frequently ask questions about places (Winter and Freksa, 2012). This is reflected in the frequency of place-related queries in human-computer interaction, e.g., Web search (Sanderson and Kohler, 2004) and question answering systems (Li and Roth, 2006). The popularity of conversational bots and assistants (Radlinski and Craswell, 2017) requires a lifting from retrieving documents as response to keyword-based queries to *natural answers to* NL *questions*.

Making place-related questions digestible for computers is the prerequisite of enabling NL communication of humans and computers about geographic places. Place-related questions not only concern location (*where-questions*), but also encompass a wide range of informational needs about places from their types and affordances to their qualities. Referring to Chapter 3, any facet of place can be enquired in place-related questions.

Extracting useful information from the questions and answers (Srihari and Li, 2000), and classifying the questions and answers (Hermjakob, 2001) are two fundamental tasks to make place-related questions digestible for computers. The former determines *what kind of information* is provided in place-related questions and answers. The latter defines *what kind of questions is asked* and *what information is relevant to respond*.

The information extraction task often involves extracting named entities and finding their dependencies (Sun et al., 2015). While these

o This chapter is mostly derived from ***"Place Questions and Human-Generated Answers: A Data Analysis Approach"*** which is written by Hamzei et al. (2020). My role as the first author of this paper was to develop the idea, design the method, to implement some part of the method and also to write and present the paper – except one subsection that describes a part of implementation which was written by the second author, Haonan Li. This work was supervised by *Prof. Stephan Winter* (my primary supervisor), *Dr. Martin Tomko* (my co-supervisor), *Prof. Tim Baldwin* and *Dr. Maria Vasardani* who provided feedback and suggestions in every stage of this study.

methods are useful for open-domain QA, in domain-specific QA the questions and answers can be further investigated. In other words, domain knowledge can be used to extract interesting patterns beyond simple named entity categories. Yet, to derive interesting patterns an encoding schema derived from the domain knowledge and an NLP method to extract each type of information is required.

The classification of questions and answers are performed through *bottom-up* or *top-down* approaches. In bottom-up approaches, the classification derived from analysing a large corpora of questions and answers often using supervised methods (e.g., Henrich and Luedecke, 2007). The time consuming annotation process is the major drawback of these methods, specially for analysing large corpora. In top-down methods, the classification schema is designed by experts and often rule-based methods are used to classify the questions (e.g., Ferrés and Rodríguez, 2006; Punjani et al., 2018). These schemas are sometimes heterogeneous due to biases in expert opinions and lack of a theoretical grounding.

A thorough analysis of place-related questions and answer sets is still missing, yet it is an essential prerequisite for future improvement of question interpretation and answer generation mechanisms. This chapter focuses on analysing and classifying place-related questions and answers by defining a data-driven encoding approach and a top-down method grounded in the core concept of spatial information (Kuhn, 2012).

To extract useful information from questions and answers, the list of place facets (see Chapter 3) are used as a building lexicon to devise a schema of information types. Then a NLP method is devised to extract information from the questions and answers based on the schema. Using the encoding schema and the extracted information, enormous questions and answers are grouped into an smaller set of encoding representations. The encoding representations capture the structural patterns of the questions and answers. The structural patterns are used to study and analyse a large corpora of place-related questions and answers, MS MARCO v2.1 (Nguyen et al., 2016).

Using the encoding representations, a bottom-up approach is designed to derive a classification of questions and answers of MS MARCO. Here, the results of information extraction are used to classify the questions and answers without performing the time-consuming annotation task. First, the representations are clustered to derive unlabelled categories of questions/answers. Next, the clusters are labelled and the classes are derived using the interpretability of encoding representations.

In the last part of this chapter, a top-down approach is proposed to classify place-related questions. This approach is grounded in core concepts of spatial information (Kuhn, 2012) and the

multi-faceted model of place (see Chapter 3). Based on the core concepts, places are conceptualised as *objects* (Purves, Winter, and Kuhn, 2019), and they can also have relationships with other concepts such as *fields*, *networks* and *events*. Thus the core question of objects is used as the grounding for the theory-driven classification. Then the relationships between places and events, networks and fields are investigated to augment the classification with more complex and yet important question types. Finally, the multi-faceted model of place is used to enrich the core questions and derive the theory-driven classification.

In the following sections, first the data-driven experiment and its results are described. Then the theory-driven classification is discussed. Finally, the data-driven and theory-driven classifications are compared in the discussion and conclusions.

4.2 DATA-DRIVEN ANALYSIS

In this section, the data-driven approach is described, and the results of analysing MS MARCO (Nguyen et al., 2016) questions and answers are provided. First, the dataset is briefly introduced and later the encoding method and results are discussed.

4.2.1 *Dataset*

The data-driven analysis of this chapter is based on MS MARCO V2.1 (Nguyen et al., 2016). The dataset is a representative collection of questions (including place-related questions) and their human-generated answers. The current version of MS MARCO V2.1 contains more than one million records[1] of search question to Microsoft Bing, including human-generated answers, retrieved documents, and question types. The dataset identifies five types of questions – (1) LOCATION, (2) NUMERIC, (3) PERSON, (4) DESCRIPTION, and (5) ENTITY (Nguyen et al., 2016), but this chapter focuses exclusively on questions and answers which are labeled as LOCATION.

The original dataset has been divided into *train*, *dev*, and *test* sets. This chapter only focuses on the *train* and *dev* subsets, with in total *56721* place-related questions. In MS MARCO V2.1,[2] the relation between questions and their answers is not necessarily one-to-one (Nguyen et al., 2016). Thus, the questions with no answer, or more than one answer can be found the dataset. Table 7 shows the number of answers per question in the combined dataset, indicating that around 22% of questions have no answer and about 2% have

1 http://www.msmarco.org
2 https://github.com/dfcf93/MSMARCOV2/blob/master/README.md

Table 7: Number of answers per location question in the combined dataset

Number of answers (per question)	Count
0	12486
1	**42884**
2	1345
3	4
4	1
5	1

multiple answers (i.e., due to question ambiguity or insufficient information in the retrieved documents).

4.2.2 *Encoding Method*

To analyse place-related questions and their answers, a semantic encoding approach is devised. First, the questions and answers are translated into an encoding representation (see Section 4.2.2.1). Next, the questions and answers are clustered to identify common patterns in the data, based on semantic encodings and contrasted to a second approach based on word-embeddings. Word-embedding is the state-of-art technique to capture the meaning of words in a vector representation, and it has been widely used for text classification recently (Yang, Macdonald, and Ounis, 2018). Here, the embedding-based clustering is used to test the quality of proposed encoding-based clustering. Finally, the relations between the questions and their answers are explored by linking the clusters of the questions and answers. In the following sections, the method for generating the semantic encodings, categorizing the questions and answers, and analyzing the relation between questions and their answers are discussed.

4.2.2.1 *Semantic Encoding*

To develop a descriptive encoding schema, linguistic facets (Section 3.4.3) are considered as the building lexicon. First, a randomly selected 1000 questions and answers are manually investigated to select relevant linguistic facets and devise the encoding schema. The selected encoding classes and their relation to the part-of-speech tags are described below:

- PLACE NAMES (TOPONYMS): Toponyms are the names assigned to places and they are expressed as proper noun phrases to refer to geographic places – e.g., *MIT*.

- PLACE TYPES: Place types are generic references that describe a set of places with shared meaning. A type is a generic noun phrase that refers to a group of places – e.g., *university* (Edwardes and Purves, 2007b).
- ACTIVITIES: Activity is a functional facet of place. In NL texts, activities are expressed as action verbs that describe the activities supported or allowed in some places – e.g., *study* (Kuhn, 2001).
- SITUATIONS: Situation is an state, opinion or feeling associated with places. In NL texts, situations are often expressed as stative verbs that describe the situations related or associated to places – e.g., *feel* (Mourelatos, 1978).
- QUALITIES: Qualities are descriptions of places that often expressed as descriptive adjectives – e.g., *warm* (Edwardes and Purves, 2007b).
- SPATIAL RELATIONSHIPS: Spatial relationships describe the configuration of places in space. Spatial relationships often expressed as spatial prepositions in NL texts – e.g., *near* (Vasardani et al., 2012).
- NON-PLATIAL OBJECTS: Places are often equipped with non-platial objects. These objects are expressed as noun phrases in natural language, and they are often related to object of activities – e.g., *coffee*.

To differentiate types of questions, the QUESTION WORDS are also considered in the encoding schema.

Table 8 shows the resulting alphanumeric encoding schema. For example, after the removal of stop words, the question *what is the sunniest place in South Carolina* is translated into the encoding `2qtrn`. This semantic encoding enables to analyze and categorize a large dataset of questions and answers by their structural patterns. Using the alphanumeric encodings, a significant number of questions and answers are captured through a few set of encodings. Moreover, using simple string matching or regular expression techniques, a subset of the questions and answers can be derived for further investigations.

4.2.2.2 *Information Extraction*

To extract place-related semantics and linguistic information from the questions and their answers, a pipeline of NL preprocessing steps, using the Stanford CoreNLP toolkit (Manning et al., 2014), is devised:

1. **Tokenization, tagging and dependency parsing:** First, the text is tokenized, and abbreviations are expanded into their

Table 8: Semantic representation encoding

Semantic Type	Part-of-speech	Code	Semantic Type	Part-of-speech	Code
where	WH-word	1	Place name	noun	n
what	WH-word	2	Place type	noun	t
which	WH-word	3	Non-platial object	noun	o
when	WH-word	4	Quality	adjective	q
how	WH-word	5	Activity	verb	a
whom	WH-word	6	Situation and event	verb	s
whose	WH-word	7	Spatial relationship	preposition	r
why	WH-word	8			

canonical forms by using a common place name abbreviation table. Next, a part-of-speech tagger and dependency parser are applied to the text.

2. **Noun encoding:** Nouns are encoded in the order of place names (toponyms), place types, and non-platial objects. First, all subsequences of a sentence are considered candidate toponyms. These candidates are then matched with the GeoNames gazetteer[3] to extract toponyms. Compound place names are preferred as better matches than simplex place names. Thus, *North Melbourne* is a compound place name and not an adjective followed by a place name. Next, nouns that are not place names have been considered as candidates for place type and the non-platial objects. Place types are identified using dictionary lookup, and all nouns that are neither place names, nor place types are encoded as non-platial objects. The dictionary of generic place types is constructed by crawling the tag values of the OpenStreetMap (OSM) spatial database,[4] due to the rich diversity of place types that are sourced from a large, global community of active volunteers.

3. **Verb encoding:** In this step, verbs related to activities or situations (states) are identified and ignored other types of verbs, known as accomplishments and developments (Mourelatos, 1978). To differentiate between activities and situations, two sets of dynamic verbs (action and stative verbs) are collected and integrated from online resources [5]. Then, a pre-trained contextualized word embedding model ELMo (Peters et al., 2018), trained from large scale bidirectional language models, is used to extract activities and situations by considering contextual information. To compare the semantic similarity between the verbs in sentences with the

3 www.geonames.org
4 https://www.openstreetmap.org
5 https://www.gingersoftware.com/content/grammar-rules/verbs

verbs in the aforementioned verb sets, first an embedding vector for each verb in the verb set is generated. Then, for the extracted verbs from the questions and answers, their embedding vector representations are compared with the vectors in the two sets. Based on the computed Euclidean distances between the vector of the verb in a sentence with the vectors of the verbs in the two sets, the extracted verbs are classified as activities or situations.

4. **Preposition and adjective encoding:** Dependency parsing has been used to encode prepositions and adjectives. For the universal dependencies (De Marneffe et al., 2014) of each sentence, `case` and adjective modifier (`amod`) dependencies for prepositions and adjectives are extracted separately. Prepositions anchored to a place name are encoded as spatial relationships, and adjectives modifying a place type or a place name are considered as qualities of places.

4.2.2.3 *Question and Answer Analysis*

After extracting the desired place-related semantics and linguistic information, the questions and answers are translated into semantic encodings. To find the categories of place-related questions and answers, clusters of the semantic representations are computed. In this clustering, first a randomly select 1024 different encodings from the unique set of all semantic encodings are analysed and subsequently re-modelled into the semantic encodings of the questions/answers as 1024-dimensional vectors. The values in these vectors are calculated based on the *Jaro similarity* (Jaro, 1989) between the semantic encodings of the questions/answers and the selected encodings. The Jaro similarity sim_j of two strings s_1, s_2 is given in Eq.2:

$$sim_j = \begin{cases} 0 & \text{if } m = 0 \\ \frac{1}{3}\left(\frac{m}{|s_1|} + \frac{m}{|s_2|} + \frac{m-t}{m}\right) & \text{otherwise} \end{cases} \tag{2}$$

where $|s_i|$ is the length of string s_i, m is the number of matching characters, and t is the number of transpositions. Specifically, two characters from two different strings are considered to match if they are the same and not farther than $\left\lfloor \frac{max(|s_1|,|s_2|)}{2} \right\rfloor - 1$. The number of matching characters (but different sequence order) divided by 2, defines the number of transpositions.

Next, k-*means clustering* is applied to the questions and answers, separately. To measure whether the semantic representations retain the contextual similarity of the sentences, the ELMo representations for the questions and answers are used with the same similarity

measure and clustering technique to provide a second set of clusters. The similarity of the clustering based on the semantic encodings and ELMo-based clusters is then evaluated. The results of clustering are also human-interpreted using the most frequent encodings in each cluster, enabling to derive the categories of place-related questions and answers.

In addition to categories of the questions and answers, frequently extracted place-related semantics are further analysed. Moreover, the geographical distribution of the identified places using GeoNames are investigated to provide a deeper understanding of the dataset. In order to disambiguate between candidate toponyms, the place names from the question with those in the corresponding answers establish a geographic context and later this set is used to localise and disambiguate the place names. The set of geospatial groundings that is associated with the minimum total distance between the toponyms is selected. In the case of a single toponym, this method cannot be applied. Consequently, only the resolved toponyms are used for describing the geographical distributions of the dataset. Finally, the relation between the questions with unresolved toponyms and unanswered questions is investigated.

4.2.2.4 *Question/Answer Relationship*

The links between categories of questions to categories of answers are investigated the analyse relationship between the content of place-related questions and their corresponding answers. As a category of place-related questions can be answered using one or more categories of answers (and vice-versa), the cardinally of the links are considered *many-to-many*.

The same approach to categorization of questions and answers is also applied to concatenated question-answers. Finally, the result of linking categories of questions and categories of answers is compared with concatenated categories. For questions with multiple answers, multiple concatenated encodings are generated and investigated, while questions with no answers are concatenated with a unique pattern (oo).

4.3 DATA-DRIVEN RESULTS

In this section, the extracted place-related semantics, geographical distribution of places and the classification of questions and answers are presented. First, the preliminary results of analysing MS MARCO dataset presented, then the encodings and frequent patterns are described. Finally, the classification of question, answers, and question-answer pairs are provided.

Encoding class	Precision	Recall
Place names	92.04%	88.63%
Types	72.02%	87.84%
Qualities	70.23%	67.24%
Spatial Relationships	95.18%	53.14%
Activities	56.44%	75.37%
Situations	52.35%	79.17%

Table 9: Evaluation of encoding extraction

4.3.1 *Preliminary Analysis*

Place-related questions are constructed with a small number of tokens, and their answers are mostly short descriptions. The tokenization and sentence segmentation show that 95.71% of questions contain less than ten tokens, and 98.17% of their answers are formulated with one, two, or three sentences.

4.3.1.1 *Extracted Place-Related Semantics*

Table 9 shows the results of the evaluation for encoding extraction. The test dataset includes 500 question-answer pairs that are randomly selected from the MS MARCO dataset. The annotated test dataset is then compared to the results of the information extraction method. As shown in the table, the extraction of place names and types are highly precise. Spatial relationships are precisely extracted but the recall is low due to complex relationships that are not necessarily prepositions – e.g., cross or flow. Moreover, the result shows that even state-of-the-art word embedding models are not performing well in the extraction of activities and situations.

Table 10 shows the top-five most frequent values extracted from the questions and answers for each type of place-related semantics. The dataset manifests a geographical bias to places in *the USA*.

The frequent patterns of place types and place names shows that places in *coarse geographical scales* are dominant in the questions and their answers – i.e. mostly in country and state levels of scale. The frequency of activities and situations in the questions is notably higher than in the answers. In other words, people use activities and situations as criteria to describe the intended places in the questions rather than asking about activities and situations happening in a place –i.e., *where is the place with particular situation/affordance* is more often asked than *what is the affordance/situation of a specific place*. Thus, a set of detailed characteristics of a place may be specified in the question, leading to a simple answer with nominal references.

Table 10: Top-five frequent place-related semantics extracted from the dataset. Frequency in brackets.

Encoding type	**in Questions**	**in Answers**
Place name	California (1393)	United States (4845)
	Texas (1391)	California (1482)
	Florida (1148)	Texas (989)
	New York (895)	Florida (961)
	Illinois (692)	New York (894)
Type	County (11702)	City (1714)
	State (2291)	State (1653)
	City (1630)	County (1438)
	Zone (745)	Area (882)
	Region (653)	Region (758)
Activity	Buy (340)	Go (64)
	Go (296)	Run (62)
	Play (120)	Leave (55)
	Build (88)	Build (53)
	See (86)	Move (38)
Situation	Find (1412)	Find (695)
	Live (746)	Have (405)
	Have (662)	Live (305)
	Grow (321)	Include (231)
	Originate (237)	Originate (125)
Spatial relation	In (3916)	In (10851)
	Near (153)	On (379)
	At (142)	At (362)
	On (109)	Near (275)
	Between (38)	Between (251)
Quality	Largest (242)	Largest (121)
	Biggest (106)	Census-designated (68)
	Highest (97)	Metropolitan (54)
	Expensive (56)	Small (46)
	Beautiful	Coastal (36)

Table 10 shows large differences between the frequencies of spatial relationships extracted from questions, and those extracted from

Table 11: Statistics of extraction and disambiguation process

	Number of Places	Percentage
All	305868	100%
Possible to resolve	242375	79.2%
Ambiguous	63493	20.8%

answers. The reason is the use of spatial relationships for localization of places. Finally, the qualities included in questions and answers differ. In questions, qualities are used as criteria for identifying or describing the intended places, with most of the values being *superlative adjectives*. However, qualities in answers are mostly used to provide additional information about intended places or to describe particular places using combinations of quality and type – e.g., coastal region, or metropolitan area.

4.3.1.2 *Toponym Resolution and Geographical Distribution*

Tables 11 and 12 show that 79.2% and 68.1% of extracted places and question/answers records can be disambiguated, respectively. Figure 15 illustrates the geographical distribution of resolved places in the questions and answers. The spatial bias in the geographical distribution of places is visible. This spatial bias can be an artefact of the monolingual English dataset, population distribution of users of the source search engine (Microsoft Bing), or unknown question sampling from the overall Bing logs used to generate this dataset.

The comparison of resolved/unresolved and answered/unanswered records shows a strong relationship between ambiguous questions and unanswered questions. Table 12 shows the contingency table of this relation, revealing that 72.4% of unanswered questions are ambiguous (9042 records out of 12486 unanswered questions), and 50.0% of ambiguous questions are not answered (9042 records out of 18093 ambiguous question/answer records). Hence, while people can interpret ambiguous toponyms considering other factors, such as salience of places, and they can answer half of the ambiguous questions, spatial ambiguity is one of the major reasons for unanswered questions. These questions may still be interpretable in context, yet this context is lacking in the MS MARCO dataset.

4.3.2 *Frequent Patterns*

Tables 13 and 14 present the top-five frequently observed semantic representation patterns. More than 40% of questions and answers can be described using the top-five representations. The top-five

Table 12: Contingency table of ambiguous/resolved and unanswered/answered questions

	Answered	Unanswered	Sum (percentage)
Resolved	35184	3444	38628 (68.1%)
Unresolved	9051	9042	18093 (31.9%)
Sum (percentage)	44.235 (78%)	12486 (22%)	56721 (100%)

Figure 15: Geographical distribution of resolved toponyms in questions and answers

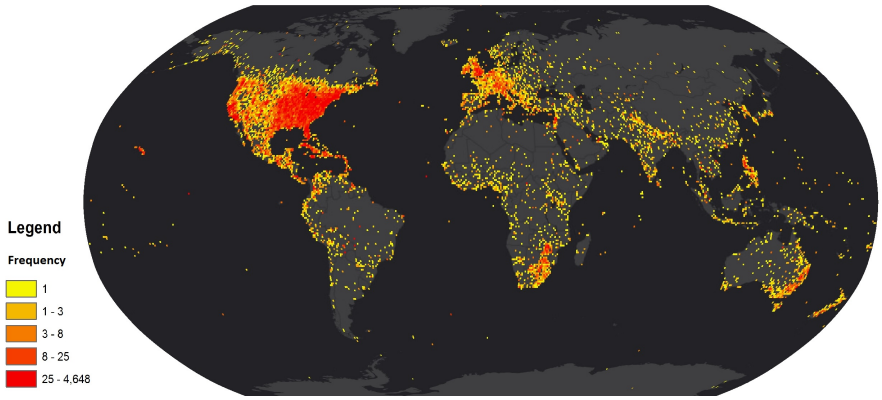

encodings for questions show that most of the spatial relationships are implicitly provided, e.g., *Barton County, Kansas* instead of *Barton County in Kansas*. In addition, the top-five questions are semantically different, first in terms of *What* and *Where* questions, and second in the manner the intended places are described. In questions with the semantic encoding `2tnn` and `2tn` the intended places are defined by type and spatial criteria, while questions with encoding `1n` and `1nn` use toponyms as verbal references to intended places with the aim to find information about their locations. Finally, pattern `1o`, which is related to implicit situations (e.g., *where are orangutans* instead of *where do orangutans live*), or to non-geographical places (e.g., *where is the key button*), is frequently also observed.

The most frequent patterns in answers provide evidence that people answer questions concisely. A single place name, which can be a compound or simple noun, is the answer for more than one quarter of the questions. Evidently, this may be also an artifact of the interface used by people to answer the questions (presumably, keyboard). In addition, implicit spatial relationships between places (e.g., `nn` as for *Tigard, Oregon*) are also more frequent in answers than explicit spatial relationships. Interestingly, one of the frequent patterns of answers is related to non-geographical places, such as *Hebrew* as an answer for a *where* question about languages.

Table 13: Frequent encoding patterns in questions

Pattern	Percentage	Example
2tnn	12.5%	*What is the county for Grand Forks North Dakota*
1n	8.6%	*Where are the Boise Mountains*
1nn	7.8%	*Where is Barton County Kansas*
1o	5.7%	*Where are ores located*
2tn	5.5%	*What is county for Seattle*

Table 14: Frequent encoding patterns in answers

Pattern	Percentage	Example
n	26.2%	*Germany* (question: *What country were gummy bears originally made in*)
o	6.2%	*Hebrew* (question: *Where did letter j originate*)
nn	4.2%	*Warren County, United States* (question: *Where is Lacona Iowa*)
rn	2.3%	*In Worcester County* (question: *What county Fitchburg Massachusetts*)
nrnn	2.3%	*Penhook is in Franklin County Virginia* (question: *What county is Penhook Virginia in*)

4.3.3 *Categorizations of Place-Related Questions/Answers*

To identify distinct types of place-related questions and answers, a clustering approach is used with manual interpretation based on the semantic encodings of the results. The same clustering technique has been used for finding clusters of questions and answers based on *word embeddings* and *semantic encodings*. The Calinski-Harabasz score (Caliński and Harabasz, 1974) is calculated to find the optimum number of clusters. The score has been computed for all clusterings in the range of 2-30 clusters. For both word embedding and semantic encoding approaches and for both questions and answers, the score is optimal for *three clusters*. Importantly, the results of clustering using semantic encoding and word embedding are highly similar, with a one-to-one relationship between the clusters with a 71.2% to 87.8% similarity using the Dice coefficient. To avoid presenting redundant information, the results of clustering based on semantic encoding which enables easier interpretability are reported in this chapter. The results of the clustering for questions and answers are interpreted manually using the most frequent semantic encodings of each cluster.

Table 15: Types of questions

Type	Percentage	Frequent Patterns
Non-spatial *Questions not aiming at localisation of places*	41.5%	`2tnn`, `2tn`, and `2tno`
Spatial *Qquestions about the location of places*	23.1%	`1n`, `1nn`, and `1nrn`
Non-geographical, and ambiguous *Questions about non-geographical places (e.g., fictional places) or ambiguous questions*	35.4%	`1o`, `1os`, and `1oo`

Tables 15 and 16 present the categories of questions/answers, and their overall percentage. The results show that questions differ in terms of *intention* and *formulation*. The first two types, non-spatial and spatial questions, relate to geographical places. The third type of questions is related to non-geographical entities such as virtual places (e.g., *Marvelous Bridge*, a place in the Pokemon World), and places in different types of space (e.g., *liver*, an organ in the space of the human body). In addition, non-spatial and spatial questions differ in the way the intended places are described in the questions. While in non-spatial questions places are described using place types, affordances, and situations, spatial questions mostly constructed using toponyms as a direct reference to the intended places. In the first and second types of questions, spatial relationships to other places are frequently observed implicitly (e.g., `2nn`, and `1nn`), as well as explicitly (e.g., `1nrn`, and `2trn`).

The answers are also classified into three categories (Table 16). There is a noticeable difference between answer patterns relating to geographical and non-geographical places (the first two vs. the third category of answers). The difference between the first two categories of answers derives from how the answer is formulated, using names and implicit relationships, or with spatial relationships which are explicitly mentioned in the answers. Moreover, the category of implicit localization includes notably fewer distinct semantic encodings — only 0.3% of all unique encodings, while the explicit localization category contains most of the patterns, 97% of all unique encodings. In other words, a broad range of simple to complex spatial descriptions are categorized as explicit localization. For example, `rnn` as for *in Nova Scotia Canada*, and `rntrnntrnnn` as for *in Monett township in Barry County and Pierce township in Lawrence*

Table 16: Types of answers

Type	Percentage	Frequent Patterns
Explicit localization and spatial descriptions	25.8%	`nrnn`, `rnnn`, and `rnn`
Implicit localization	42.5%	`n`, `nn`, and `nnn`
Non-geographical, and unanswered	31.7%	`oo`, `o`, and `ooo`

County Missouri United States both are categorised as explicit localizations.

4.3.4 *Question/Answer Relationships*

To investigate the relation between place-related questions and answers, two different scenarios are taken into account. First, the relation between types of questions and answers are explored. Next, the concatenated encodings of questions and answers are investigated using a clustering approach and manual interpretation of the frequently observed patterns in each cluster.

4.3.4.1 *From Questions to Answers*

A many-to-many relationship between categories of questions and answers is considered and presented as a contingency table (Table 17). Spatial questions are mostly answered with implicit localization answers and are less ambiguous than non-spatial questions. This is because they are formulated with direct references to the intended places through toponyms. Non-spatial questions are more ambiguous and are mostly answered through complex descriptions or remain even unanswered. For example, *what are the best airports in Southern Utah?* is unanswered in the dataset, and it cannot be answered without describing what *best* means here.

The relationships between non-geographical and ambiguous questions and implicit and explicit localization answers are observed. Several reasons contribute to the observed relationship, including (1) human interpretation of ambiguous questions, (2) issues in extraction of place-related semantics, which are propagated to the semantic encodings, and (3) similarity of patterns in formulating questions and answers for geographical and non-geographical places impacting on clusterings. There is a strong relationship between non-spatial question and answers with explicit localizations. Three primary reasons for this are:

1. **Answering style:** The content of question (how the intended places are defined) can also be repeated a part of answer (e.g.,

Table 17: Contingency table of linking the clusters

Q/A	Explicit localization	Implicit localization	Non-geographical
Non-spatial	9512	7966	6266
Spatial	2463	8400	2332
Non-geographical	2777	7970	9535

Question (`2trnn`) : *What beach is closest to Busch Gardens Tampa?* Answer (`nrnn`): *Clearwater beach is the closest to Busch gardens, Tampa.*).

2. **Ambiguous questions:** Ambiguous questions are answered in more detail (partially) related to localization information (e.g., Question (`2tn`): *What city is Olongapo City?* Answer (`nqtrnn`): *Olongapo City is a 1st class highly urbanized city in Central Luzon Philippines*).

3. **Misleading questions:** While the inquirer's question is semantically a *where-question*, it is formulated as a *what-question* (e.g., Question (`2nn`): *What is Bentonville Arkansas County?* Answer (`nrnn`): *Bentonville is in Benton County Arkansas*). Here, the encoding captures semantics different from the intent of the inquirer and this therefore affects the categorization of the question.

4.3.4.2 *Concatenated Clustering*

In the last experiment, the semantic representations of questions and their corresponding answers are concatenated and then clustered. The results of clustering show (Table 18) that there is a direct one-to-one relationship between the categories of questions and the concatenated clusters. Using the Dice similarity coefficient, the similarity of first, second, and third clusters of questions and the first, second, and third clusters of questions concatenated with answers are 88.86%, 75.83%, and 75.98%, respectively. This shows that the encoding of questions can determine the encoding representations of their answers.

4.4 THEORY-DRIVEN CLASSIFICATION OF PLACE-RELATED QUESTIONS

To derive a top-down classification of place-related questions, the core concepts of spatial information (Kuhn, 2012) is used as the grounding theory. In the following sections, first the core concepts are reviewed

Table 18: Frequent patterns in concatenated clusters

ID	Frequent Encoding	Example
1	`2tnn-n`, `2tn-n`, and `2tnn-nrnn`	*What county is Roselle NJ? Union County*
2	`1n-n`, `1nn-rnnn`, and `1nn-n`	*Where is Fairview? in Multnomah County Oregon United States*
3	`1o-oo`, `1oo-oo`, and `1no-oo`	*Where is the appendix to the liver? No Answer Present.*

and their relevance to place conceptualisation is described. Next, the top-down classification of place-related questions is derived using the *object-based conceptualisation of place* (Purves, Winter, and Kuhn, 2019) and the multi-faceted model of place.

4.4.1 *Core Concept of Spatial Information*

Kuhn (2012) presents a set of core concepts for spatial information that mitigates challenges in transdisciplinary research on geographic and spatial problems. The core concepts includes ten concepts that are hierarchically organized into *base* concept (i.e., location), *information content* concepts (i.e., object, field, network and event) and *information quality* concepts such as granularity and accuracy (Kuhn and Ballatore, 2015). This chapter only focuses on the base concept and information content concepts which are briefly described below[6]:

- LOCATION: Location is the answer to where-questions. It can be described as place descriptions using spatial relationships and anchor places, or captured using coordinates in a spatial reference system.
- OBJECT: Objects are identifiable and spatially bounded individuals. Objects have thematic properties that describe their meaning and they can have crisp or fuzzy boundaries.
- FIELD: Fields describe phenomena based on measured or calculated value(s) in each location. They answer the question of *what is here?* for anywhere in space.
- NETWORK: Network is the construct of relations among objects. Networks capture the relationships between objects, any relationships that connect or relate objects with each other.

6 more elaborate definitions can be found in the original paper, (Kuhn, 2012)

- EVENT: Events are the cause of changes in objects' properties, field values and the structure of networks.

Purves, Winter, and Kuhn (2019) conceptualised places in information science as spatial *objects* with the following properties:

- Places, as spatial objects, have shared and inter-subjective *identity* as they are socially-constructed phenomena;
- Places can emerge from *fields*. Bounded field values form regions in space, and such regions may become places with their identities derived from the field that they emerged from – e.g., warm temperate zones.
- Places are not isolated objects and they can form *networks* based on their relationships – e.g., a transportation network
- Places can participate in *events* and consequently their properties be changed.
- Places can emerge from events. The location of the event form a region, and from the region a place may emerge as 'the place where the event happened' – e.g., Gettysburg Battlefield.

These properties of place as spatial objects and the core question of objects (i.e., what are the properties of this object?) (Kuhn, 2012) are the theoretical basis of the top-down classification of place-related questions. The place facets (see Chapter 3) are used to enrich the core question to derive the classification of place-related questions. Using the relationship between places with fields, networks and events, the classification is extended to cover more sophisticated question types.

4.4.2 *Classification of Place-Related Questions*

Referring to Chapter 3, the multi-faceted model of place captures four sub-classes of facets:

- *anthropocentric emotive facets*, such as place attachment;
- *anthropocentric functional facets*, such as affordance;
- *geographic physical facets*, such as material form; and
- *geographic spatial facets*, such as location and accessibility.

Such classes of facets describe places. A fifth, distinct class of facets called *linguistic facets* captures referential properties of places such as verbal references and place names, which may be non-descriptive in nature but still identify the intended subject of a speaker.

These place facets can be used to derive place-related questions if they can be linked to the core question of objects, described in the

previous section. The core question of objects is the question template for places, and thus can be rewritten as: *What is the value of a place facet of this place?* Using this question template, the categories of place facets can be used to generate more specific question types.

Any place facet can become an intentional property of a particular place of interest. While emotive, functional and physical facets are attributed to individual places, spatial facets may characterize multiple places forming a network. If a place network is viewed as an identifiable whole with properties, then it has been conceptualized as an object – e.g., a public transportation system. Thus, the emotive, functional and physical facets can be applied in the same way to describe the whole network. For example, people's satisfaction with public transportation in Melbourne is an example of a question that inquires an emotive facet. To provide a list of more detailed place-related question types, the facet categories are used as answers to place-related questions.

Emotive and functional facets inform templates for questions about the relation between places and people. In both cases, these questions can be about a person (e.g., a subjective experience or opinion), a community (e.g., cultural values and traditions), or the whole world (e.g., historical information). The emotive facets (attachment, sense of place, commitment, and salience and familiarity) provide questions that are mostly of interest to social scientists, policy makers, and planners. Yet, these questions also cover everyday questions of residents and tourists about the experiences and opinions of others. The functional facets capture a different aspect of people's exploration and experience of places, focusing on what a person can do or is allowed to do in a place.

Physical facets capture properties of the material, form, and parts of places. Their kinds vary across different types of places. For man-made places, architectural properties may be of interest; for natural places (such as mountains), geological and morphological characteristics play that role. The question templates using emotive, functional and physical facets are shown in Table 22 with examples.

Spatial facets capture spatial properties of an individual place and spatial relations among places. In the first subclass, *location*, *footprint* and *boundary* (if any) are the major properties that describe where a place can be found in space. Additional spatial properties are *elevation*, *length*, and *area*. The second subclass is related to the network core concept because it captures relationships among places. Two types of relationships lead to two types of networks: (1) spatial relationships between places in a configurational network and (2) accessibility relationships among places in a transportation network. Considering the core question of networks (i.e., what connects the objects?), the question can be rewritten for places as *what connects*

Table 19: Question templates using emotive, functional and physical facets

Category	Subcategory	Template/Example
Emotive	Individual	*What is your feeling about place* p*?*
		How does it feel (What is your feeling) after visiting the orphanage?
Emotive	Social	*What are the social/cultural aspects of place* p*?*
		What is the cultural significance of Uluru?
Functional	Individual	*What can I do at place* p*?*
		What can I eat at this Subway restaurant?
Functional	Social	*What is (socially) allowed in place* p*?*
		What is allowed in this non-containment zone?
Physical	Material	*What are the material properties of place* p*?*
		What is the dominant soil type in East Texas?
Physical	Form	*What are the morphological properties of place* p*?*
		What is the architectural style of the Fallingwater?
Physical	Parts	*What places are part of place* p*?*
		How many rooms does this mansion have?

Table 20: Question templates using spatial facets

Concept	Facet	Template/Example
Object	Location	*Where is (What is the location of) place* p*?*
		Where is Melbourne?
Object	Boundary	*What is the boundary of place* p*?*
		What is the boundary of international waters around Australia?
Object	Spatial properties	*What is the value of this spatial property of place* p*?*
		What is the surface area of Australia?
Network	Spatial relationships	*What is the spatial relationship between place* p_1 *and place* p_2
		What is the distance between Melbourne and Sydney?
Network	Accessibility	*How to get to (What is the direction of) place* p_2 *from place* p_1*?*
		How to get to Sydney from Melbourne by car?

Table 21: Question templates using linguistic facets

Category	Subcategory	Template/Example
Linguistic	Specific	*What is this place* p *called?*
		What is the name of this mountain?
Linguistic	Generic	*What is the kind of this place* p*?*
		What is that silver thing?

these two places? Table 20 shows the question templates and examples derived from spatial facets.

Questions can be posed about the linguistic facets to find referential information about places. Linguistic facets include specific and generic references to places (Edwardes and Purves, 2007a). In the simplest form, questions can be asked about the names assigned to places (toponyms). Place types are an example of generic references and thus can be asked for in place-related questions. The question templates using linguistic facets are listed in (Table 21).

Table 22 shows the question types that are asked about the relationships of places with fields and events. Places as objects can be carved out of fields and events. In the former, places (e.g., the warm temperate zone) can emerge from bounded values of field(s) attribute (e.g., temperature within an interval). Similarly, when places are carved out of events, a new place emerges where an event has occurred (e.g., the Gettysburg Battlefield). These places can be

Table 22: Question templates for places in conjunction with events and fields

Concept	Relation	Concept	Template/Example
Place	Emerge from	Event	*Where did the event take place?*
			Where did the battle of Gettysburg take place?
Event	Change	Place	*What is changed in place* p *during the event?*
			What is the impact of global warming on Lake Urmia, Iran?
Place	Emerge from	Field	*Where are places with a specific (range of) field attribute(s)?*
			Where are warm temperate zones located?

subjected to the aforementioned question templates, and also they can be inquired based on their relation with fields and events. Moreover, places can participate in events, and their properties can be changed through these events. Table 22 provides question templates for places in conjunction with events and fields.

The theory-driven classification of place-related questions, as presented in this chapter, only captures *description questions* related to geographic places (Nielsen et al., 2008). Yet other classes of questions such as *comparison questions* (e.g., Is the Black Sea bigger than the Baltic Sea?) can be constructed with the same procedure – i.e., designing the core question of objects based on the question class and enrich the core question with place facets.

4.5 DISCUSSION AND CONCLUSIONS

In this chapter, the place-related questions were investigated using data-driven and theory-driven approaches. In the data-driven experiment, place-related questions and human-generated answers of MS MARCO were studied using the semantic encodings. The frequent patterns and the classifications were described to show what people expect from a search engine when they submit a place-related question. In the second part of the chapter, a theory-driven classification was presented that includes detailed and more diverse set of classes compared to the data-driven results. In the followings, the main points of the data-driven and theory-driven approaches are described and their differences are discussed.

In the data-driven experiment, the relations between place-related questions and answers are studied by linking categories of questions and answers, and concatenated clusters. The most frequent encoding patterns found to provide qualitative insights in place-querying behavior. Using geocoding and disambiguation techniques, the

geographical representation of places in questions and answers reveals that while the dataset has a global coverage, it is highly biased to North America, Western Europe, and Australia. While this bias may also be in part attributed to the biased coverage of the GeoNames gazetteer used for toponym resolution (Graham and De Sabbata, 2015), users of the MS MARCO dataset are advised to be mindful of this coverage bias and its potential impact.

Through extraction of place-related semantics, a semantic bias related to the scale of places is observed. Analyzing the frequent place types shows that while questions to fine-grained places such as *schools, libraries, and airports* occur, most of the frequently asked place types are related to coarse geographic scales. The top-ten most frequent types of places identified in the questions and answers are associated with the *city to country level* scales. The diversity of activities, situations, qualities, and spatial relationships in the dataset shows further potential for studies in geographic information retrieval.

Expressing the patterns in questions and answers through the proposed semantic encoding proved to be useful for categorization. It produces results consistent with those achieved by the state-of-the-art word embedding approach using the pre-trained ELMo model. Yet, the encoding approach allows for manual inspection of the patterns and their qualitative interpretation, thus providing a valuable insight into place-related querying and answer-giving behavior.

Using a data mining approach and interpretation of the frequent patterns in the results, three main types of questions in MS MARCO – i.e., non-spatial, spatial, and non-geographical and ambiguous questions. Applying the same strategy to the answers similarly reveals three main categories of answers, explicit localization (spatial descriptions), implicit localization (names and addresses), and non-geographical and unanswered questions.

In the theory-driven classification, fourteen classes of place-related questions are identified using the places facets and core questions of objects and networks. Moreover, three classes of questions are provided to classify the questions about places in relation to fields and events. In comparison to the results of the classification derived from data-driven experiment, the theory-driven classification is more detailed and covers more diverse types of questions.

The reason behind the more details in the top-down classification is derived from possible biases to question corpora, in this case, MS MARCO. Place-related questions can be asked in different situations, from everyday communication to asking a smart assistance bot. Yet, a dataset collected from a specific situation may not present the complete view to the all possible question types. For example, a place-related question about emotive facets is not

popular in MS MARCO dataset because of the situation with which the data is collected from.

In this chapter, the encodings are used to investigate the place-related questions and their answers. While encodings are useful for classification purposes or to find patterns in the structure of the data, their potential is limited for other tasks such as translating questions to queries. In the next chapter, the encodings are extended to capture more information types and also parsing is introduced to capture the relationship between the encoded information units. In Chapter 5, the parsing method is used to translate place-related questions to formal queries.

5 FROM PLACE-RELATED QUESTIONS TO QUERIES

It is the duty of machines and those who design them to understand people. It is not our duty to understand the arbitrary, meaningless dictates of machines.

— Norman (2013)

5.1 INTRODUCTION

Consider the following place-related question, taken from the Geospatial Gold Standard dataset (Punjani et al., 2018):

Question: *How many pharmacies are in 200 meter radius of High Street in Oxford?*

Structured data stored in databases and knowledge bases are more suitable than unstructured texts to answer such questions. Using structured data, GeoQA systems can perform spatial (i.e., 200 meter radius) and non-spatial operations (i.e., how many), and retrieve information based on specific criteria (e.g., pharmacies in High Street). The complementary role of structured data for QA systems has been noted in several studies (e.g., Chen, 2014; Dimitrakis, Sgontzos, and Tzitzikas, 2019), yet answering NL questions using structured data is challenging. Using structured data in QA systems requires formal-language queries to be constructed from NL questions.

Translating place-related questions to formal query languages such as *SPATIAL SQL* or *GeoSPARQL* is a necessary step to answer the questions using structured data (Chen, 2014; Punjani et al., 2018). Such translations enable place-related QA systems to benefit from available knowledge sources and to perform spatial analysis to answer complex place-related questions – e.g., a 200 meter buffer zone is needed to answer the introductory question. Yet, the translation task demands a detailed analysis of place-related questions. For example, the encoding approach, proposed in

o Section 5.2.4 of this chapter is an extended version of my contributions to a conference paper, ***"Extracting Interrogative Intents and Concepts from Geo-Analytic Questions"*** (Xu, Hamzei, Nyamsuren, Kruiger, Winter, Tomko, and Scheider, 2020). My role as the *second author* of this paper was to develop the idea for parsing geographic questions, implement the method and to write the method section of the paper.

Chapter 4, is not expressive enough to enable automatic translation – i.e., the encoding does not capture the relationships between extracted information units.

The translation task often involves two major steps: (1) parsing the questions into an intermediate structure, and (2) generating queries using the parsing results (Punjani et al., 2018). In the former, information units are extracted from the content of questions, and their dependencies are often stored in graph or tree data structures – e.g., <*pharmacies, in 200 meter radius, High Street*>. In the latter, the extracted information is used to define variables and their relationships using the parse results and the predefined syntax of the targeted formal language(s).

Two types of information construct the questions, *intent* and *criteria*. The former determine clues about the answer(s) and sometimes involve spatial and non-spatial operations – e.g., "how many pharmacies". The latter describes how to find relevant information in the information sources – e.g., pharmacies in 200 meter radius of High Street in Oxford. Intent and criteria are necessary to be extracted, and sometimes inferred, from the NL questions, and without an accurate extraction the translation fails.

The aim of parsing is to extract information and information dependencies to a structured and machine-digestible representation. Parsing results should be informative enough to be translated automatically into formal query language(s). A suitable parsing method in a domain-specific QA should be linked to the domain concepts in order to ease the query generation step. Domain concepts and their relationships are available as grounding, and the approach to define or identify them is already captured in domain ontologies. Hence, the concepts can be used to design a domain-specific parsing that benefits from the available domain knowledge and avoids reinventing redundant terms.

A parsing based on domain concepts ease the process of query generation. In query generation, the task is simplified into defining concepts and their relationships based on the extracted information from the NL questions. The extraction of intent and criteria determines what should be the intent of the generated query, and what should be stated as the criteria to narrow down the search domain.

In this chapter, the object-based conceptualization of place (Purves, Winter, and Kuhn, 2019) is used as the grounding to design a parsing method for place-related questions. Moreover, the encoding approach (see Chapter 4) is extended and their links to domain concepts such as *location*, *object (place)* and *event* are studied[1]. Grammatical parsing methods are used to identify these concepts

1 The object-based conceptualisation of place and core concepts of spatial information are reviewed and discussed in Section 4.4.1

inside the questions (e.g., High Street) and extract their relationships (e.g., High Street, in, Oxford). The results are formulated as First-Order Logic (FOL) statements that are machine-digestible and human-readable. FOL statements can capture the intent and criteria explicitly which is necessary to generate formal queries. Finally, a dynamic approach[2] is proposed to generate GeoSPARQL queries from the FOL statements to demonstrate the applicability of FOL statements in translating NL questions to formal queries.

The novelty of the proposed method is twofold. First, the method is not biased to specific types of place-related questions by supporting the concepts and their relationships in the object-based conceptualisation of place. Second, the method is reusable to translate place-related questions to other formal languages(s) with minimal efforts using FOL statements as the intermediate representation. As will be discussed in Section 5.3, the method also improves existing state-of-the-art approaches to translate place-related questions to GeoSPARQL queries.

5.2 METHODOLOGY

5.2.1 *Dataset*

In this chapter, the Geospatial Gold Standard dataset (Punjani et al., 2018) is used to test the method for the translation task. This dataset contains 200 place-related questions collected for translating questions into GeoSPARQL queries. The dataset is generated by students of the Artificial Intelligence course at the National and Kapodistrian University of Athens. The participants were instructed to formulate questions about geographic places that can be answered using information available in knowledge bases or through spatial analysis afforded by GeoSPARQL.

5.2.2 *Method*

Figure 16 shows the proposed method to translate place-related questions to formal queries. To perform the translation, a workflow is designed based on encodings, their relationships and the core concepts of spatial information (Kuhn, 2012). In short, the workflow includes:

- **Extracting encodings**: This step extracts relevant information units from the content of place-related questions;

2 The dynamic approach uses templates for GeoSPARQL *constructs* rather than a template for the *whole query*. Consequently, the query was generated dynamically from the constructs, and the method is less biased to specific types of queries.

Figure 16: The proposed method to translate place-related question to GeoSPARQL queries.

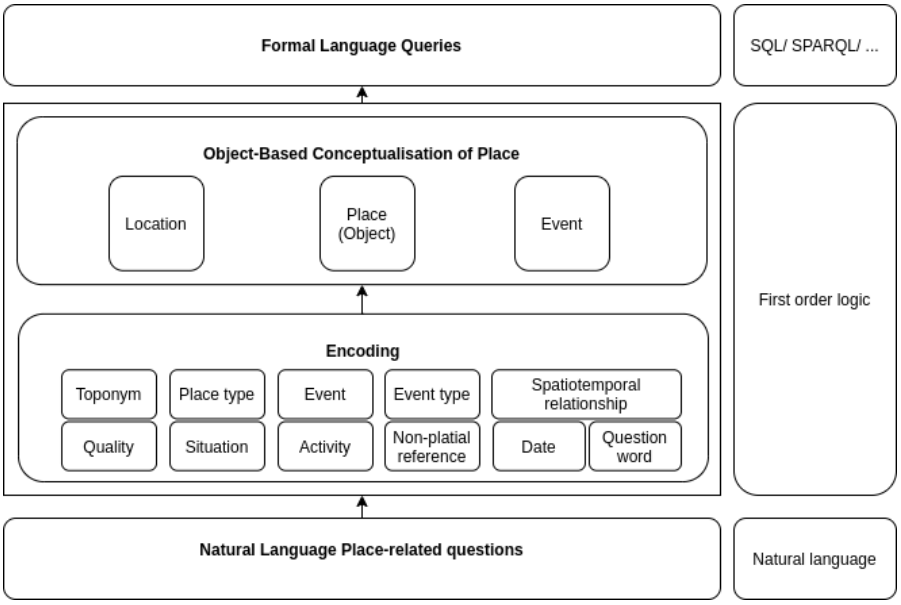

- **Parsing**: In parsing, the relationships between encodings are identified and captured;
- **Identifying concepts**: Using the encodings and their relationships, high-level concepts of location, place (object) and event are identified;
- **Generating logical statements**: The concepts and their relationships are expressed in logical statements that define what is expected in answers (intent) and what are the criteria, in a machine digestible manner;
- **Generating GeoSPARQL queries**: Using the logical statements and the predefined syntax and specification of GeoSPARQL, formal queries are generated for the place-related questions.

In the following sections, the details of each step and its implementation are described.

5.2.3 *Encoding Extraction*

For the translation, encoding classes are extended to capture events and event types, as well as dates and numbers. Logical operations such as *and*, *or*, *negation* and *comparison* are then added to the encoding schema. Table 23 shows the schema of encoding classes and their codes.

The Geospatial Gold Standard dataset contains questions which are grammatically correct, and consequently the encodings can be

Table 23: Semantic representation encoding

Semantic Type	Code	Semantic Type	Code
where	1	*place name*	P
what	2	*place type*	p
which	3	*event*	E
when	4	*event type*	e
how	5	*non-platial object*	o
how many	6	*activity*	a
why	7	*situation*	s
is/are	8	*spatial relationship*	R
date	d	*temporal relationship*	r
place quality	Q	*object quality*	q
comparison	<, >, =	*and*	&
or	|	*negation*	!

derived from Named Entity Recognition (NER) and part-of-speech tags without a need for using gazetteers[3]. In the following, the method to extract the encodings using pretrained NLP models is described:

- **Noun encoding:** Noun phrases can be place names, place types, event names, event types or other objects. To detect place names and event names, first proper nouns are selected using part-of-speech tagging, next a pretrained model for fine-grained NER is used to detect the place names and event names (Lample et al., 2016). Generic noun phrases can be place types or event types. Here, a look-up approach is used to test whether a generic noun refers to a place type or an event type. The place type list is collected from OpenStreetMap tags with the same method described in Chapter 4. The list of event types are collected from the DBPedia ontology of events which includes social and natural events. The remaining unlabelled noun phrases are labelled as non-platial objects.

- **Verb encoding:** The verb encoding is identical to the method in Chapter 4 – i.e., word embedding vectors are used to label verbs as activity or situation.

- **Preposition encoding:** Preposition phrases can be related to spatial relationships, temporal relationships and comparisons. If a preposition refers to place(s) (either place names or place

3 This is different from MS MARCO questions which are not necessarily grammatically correct, See Chapter 4

types), the preposition is captured as a spatial relationship. If the preposition phrase refers to a date then the preposition is captured as a temporal relationship. Otherwise, if the preposition phrase includes *comparative adjectives*, it is captured as comparison – e.g., greater than.

- **Adjective encoding:** In this study, both qualities of places and qualities of other objects are captured through two encodings, Q and q respectively. After resolving noun phrases into proper encoding classes, adjective phrases that are referring to place names and place types are captured as place qualities and adjectives referring to events, event types and other objects are captured as object qualities. Only the superlative (e.g., oldest) and descriptive (e.g., old) adjectives are captured as qualities. Comparative adjectives (e.g., bigger) are a constituent of preposition phrases (e.g., bigger than) which are encoded as comparison due to their different role in constructing FOL statements and GeoSPARQL queries.

- **Conjunction encoding:** Conjunctions are identified through part-of-speech tagging, and their encodings are identified through a look-up approach. The conjunctions are encoded either as *or*, *and*, or *negation*.

5.2.4 *Constituency and Dependency Parsing*

Constituency parsing and dependency parsing are two approaches to study the structure of NL sentences. The constituency parsing captures how tokens can be combined to construct phrases and how phrases create more complex phrases or sentences (Figure 17). On the other hand, dependency parsing identifies the relationships between tokens and is able to capture long distance relationships between them (Figure 18).

In this study, both constituency parsing and dependency parsing are used to extract relationships between encodings. If a relationship can be extracted from phrases and their constituents, constituency parsing is used; otherwise dependency parsing is used. Moreover, the extraction of intent and criteria from the content of questions is achieved by a heuristic method derived from constituency parsing results.

The parsing method starts with constituency parsing and identification of phrase-level information such as *conjunctions phrases*, *quality phrases* and *location phrases*. The following steps are performed in analysing constituency parsing results:

1. **Preprocessing:** The result of constituency parsing are trimmed and enriched by the extracted encoding. In the trimming task,

Figure 17: Constituency parsing

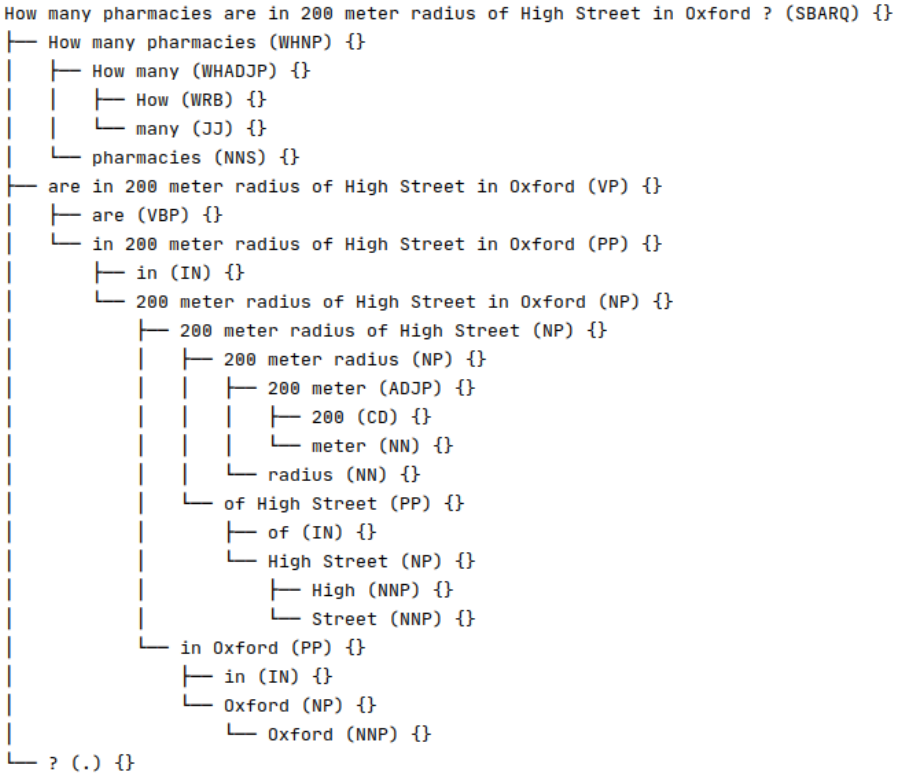

the extracted compound phrases (e.g., High Street) are captured as leaf nodes in the tree representation. Hence, the constituency tree expands to a meaningful phrase level that is not necessarily consisting of individual tokens. Then, the extracted encodings are labelled in the tree representation. For example, in Figure 17, *High Street* is captured as two leaf nodes while in Figure 19 (after analysing the constituency tree), the whole phrase is captured as a leaf node of the tree.

2. **Conjunction phrases:** First, leaf nodes labelled with conjunction encodings (i.e., `&`, `|`, and `!`) are selected from the constituency tree. For *and/or* conjunctions, if the parent node includes constituents with the same encoding role (e.g., multiple place names), the parent node is labelled as a *conjunction phrase* – e.g., towns or cities in *what are the towns or cities in UK?* For negations, if the parent nodes contain place names or event names the parent nodes are labelled as *negation phrase* – e.g., 'except London' in *What is the largest city in UK except London?*

3. **Quality phrases:** The leaf nodes encoded as qualities (i.e., `q` and `Q`) are retrieved from the constituency tree, and their parent nodes are selected for further investigation. If the parent

Figure 18: Dependency parsing

```
are (root) {['AUX']} []
├── How (advmod) {['ADV']} []
├── pharmacies (nsubj) {['NOUN']} []
│   └── many (dep) {['ADJ']} []
├── in (prep) {['ADP']} []
│   └── radius (dep) {['NOUN']} []
│       ├── meter (nn) {['NOUN']} []
│       │   └── 200 (dep) {['NUM']} []
│       └── of (prep) {['ADP']} []
│           └── Street (pobj) {['PROPN']} []
│               └── High (amod) {['PROPN']} []
├── in (prep) {['ADP']} []
│   └── Oxford (pobj) {['PROPN']} []
└── ? (dep) {['PUNCT']} []
```

node includes nodes with relevant encoding classes (i.e., `p`, `P`, `e`, `E` and `o`), the parent node is captured as a quality phrase.

4. **Location phrases:** Location phrases are constructed by spatiotemporal relationships and their corresponding anchor places/dates – e.g., in 200 meter radius of High Street. Here, location is considered in space and time. Location phrases are captured by finding an ancestor node of spatiotemporal relationships which includes the anchor places/dates. Figure 19 shows two labelled locations, *in Oxford* and *in 200 meter radius of High Street in Oxford*.

5. **Measure Phrases:** Phrases that are only constructed with cardinal numbers and non-platial objects are detected measures phrases. In the introductory example, *200 meter* is a phrase with numeric (200) and non-platial object (meter) constituents (Figure 19).

6. **Comparison phrases:** Comparison is a binary relationships that has a source (the left-side of comparison) and target (the right-side of comparison). While the target and comparison tokens (e.g., less than) are often form a phrase, the source can have a long distance relationship with comparison tokens, depending on the structure of the questions. Hence, the relation between comparison tokens and the target are captured through constituency parsing and dependency parsing is used to detect the source of the comparison. If the parent node of an identified comparison encoding includes a valid target phrase (i.e., `p`, `P`, `e`, `E`, `o`, and `n`), the parent node is labelled as a comparison phrase – e.g., more than 10 districts.

Figure 19: Labelled constituency parse tree

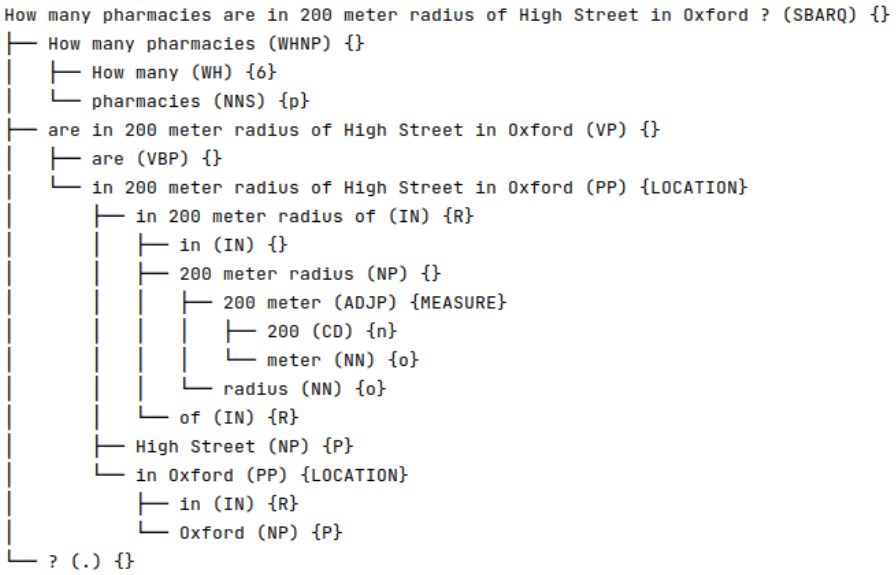

The results of detecting phrase-level information is shown in the labelled constituency tree in Figure 19.

After identification of phrase-level information, intent extraction finds the goal of the place-related questions. The following heuristics are used to detect the intent of questions:

1. **Question word rule:** The question word gives an initial idea about the intent of the questions – e.g., *where*, *what* and *Is/are*. For yes/no questions (*Is/are*) not only intent is clear (i.e., to find whether criteria is met or not) but also the answer domain is evident, i.e., a boolean value. For *where* questions, the intent is the location of the selected variable (either a place name or a place type) which can be identified using the following rules. For *what* and *which* question, the following rules are used to determine the intent. The question words *how many* is also special case which refers to the count of the selected variable.

2. **Less specificity rule:** The more specific the concepts are, the less likely they are to be the intent of the question — e.g., a *non-platial object* or a *place type* (if mentioned in the question) is more likely to be the intent of the question in comparison to a *place name*. For example, in the question *Which cities in England have at least two castles?*, cities and castles are more likely to be the intent of the question than England. This rule determines the specificity of concepts using their encoding classes, i.e., *objects*, *place types* and *place names*. The objects (place properties) are the least specific concept. Next more specific concept is *place types*, and finally the most specific concepts are *place names*.

3. **Phrase rule:** If the intent is ambigous and the candidates belong to place types or place names (e.g., multiple place types are identified), the location phrases are used to reduce the ambiguity. If a place type or place name is the intent, it must not belong to a *location phrase*. For example, *Where in California is UCSB?*, California is part of a location phrase (in California) and must be removed from the list of intent candidates. Such location phrases are spatial criteria and they should be avoided in the intent recognition process. In the same manner, non-platial objects which belong to activity/situation phrases or complex spatial relationships phrases are removed from the candidate list. Referring to the introductory example, *radius* is part of a complex spatial relationship, *in 200 meter radius of*, and removed from the candidate list.

4. **Earlier position rule:** If the intent is still ambiguous, the earlier position rule is used. This rule is based on the structure of the questions. The subject of the questions is the earliest concept extracted from the questions. For example, in question: *Which cities in England have at least two castles?*, the earliest concept (i.e., cities) is determined as the intent of the question.

Referring to the introductory question, the intent is identified as *how many pharmacies*.

After finding the intent and phrase-level information, dependency parsing is used to detect the relationship between (1) places and events with location phrases, (2) situation and activities with non-platial objects, (3) places with situations and activities, and (4) comparison phrases and their source.

1. **Preprocessing:** The result of dependency parsing are trimmed and enriched by the extracted encoding and identified phrases from constituency parsing. In the trimming task, each of the extracted compound phrases (e.g., High Street) or phrase-level information from constituency parsing is captured as a single node in the dependency tree representation. Hence, the dependency tree is summarised to extract relationships among phrases rather than individual tokens. Finally, the extracted encodings and phrases are labelled in the tree representation. Figure 20 shows the dependency tree of the introductory example after the preprocessing step.

2. **Places/event references and location phrases:** Using the terminology introduced by Vasardani et al. (2013), a spatial description includes three elements, – i.e., *locatum*, *spatial relationship*, and *relatum* (see Section 2.2.1.2). A location phrase is a combination of the spatial relationship and relatum that describe the location of a place or an event in space – e.g., in

Figure 20: Dependency parse tree after preprocessing

```
are (root) {['AUX']} []
├── How many (advmod) {['ADV']} [6]
├── pharmacies (nsubj) {['NOUN']} [p]
├── in 200 meter radius of (prep) {['ADP']} [R]
│   └── High Street (pobj) {['NOUN']} [P]
├── in (prep) {['ADP']} [R]
│   └── Oxford (pobj) {['NOUN']} [P]
└── ? (dep) {['PUNCT']} []
```

200 meter radius of High Street. The relationships between the locatum (a place or an event) and the location phrase captured through dependency parsing. If the location phrase is a dependent (*prep* relation[4]) of a place or an event, their relationships are captured as a locatum and location phrase relationship – e.g., pharmacies (locatum), in 200 meter radius of High Street (location phrase).

3. **Situations/activities with non-platial objects:** Situations and activities often include references to non-platial objects. Hence, in these cases, the verbs do not completely describe the situation and activities – e.g., to buy coffee or to see birds. If a non-platial object is a dependent (object phrases of the verb, *dobj* relation) of the situation/activity verb, then the relationship is identified between the verbs and noun phrases.

4. **Places with activities/situations:** To identify which place in the question is described through situation/activity phrases, dependency relationships are used. If a situation/activity phrase is a dependant (subject phrase of the verb, *nsubject* relation) of a place reference, their dependency is captured as a place relationship with situations/activities.

5. **Source of comparison phrases:** The relationship between the source of comparison to comparison phrase is captured through dependencies. The source of a comparison belong to a valid encoding type (i.e., `p`, `P`, `e`, `E` and `o`) and the comparison phrase is a direct dependent of it in the dependency parse tree. For example in *does England have more counties than Ireland*, England has a valid encoding type (i.e., `P`) and the comparison phrase (i.e., more counties than Ireland) is its dependent in the parse tree.

4 Details of part-of-speech tags and dependency types can be found in Appendix B

5.2.5 *Concept Identification*

In concept identification, the references to places, events and locations are identified. A reference to a place or an event, could be a proper noun or a generic reference with additional spatiotemporal conditions – e.g., hurricanes in the US since 1980, or pharmacies in 200 meter radius of High Street in Oxford.

In this step, references to places and events are first detected using the extracted encodings and their relationships. Place names and event names are labelled as place concepts and event concepts – e.g., *Oxford* in the introductory example. Then, generic references to places and event with their relationships (e.g., spatial relationships) are captured together as place concepts and event concepts – e.g., *pharmacies in 200 meter radius of High Street* in the example.

Finally, the anchor places with their corresponding spatial relationships are captured as location phrases. Here, temporal relationships with anchor dates are also considered as location concept. Consider the following example, *how many cities are affected by hurricanes in the US since 1980?* In this example, a complex location phrase identified by merging spatial and temporal references – i.e., *in the US since 1980*.

Referring to the introductory example, three place concepts are identified as (1) Oxford, (2) High Street in Oxford, and (3) pharmacies in 200 meter radius of High Street. Moreover, two locations are detected, (1) in 200 meter radius of High Street and (2) in Oxford. The locations describe the spatial frame to disambiguate and identify the intended place(s) – i.e., pharmacies. Later, the detected concepts are declared as constants or variables through FOL statements and their relationships captured through functions with the constant or variable as arguments.

5.2.6 *Generating First-Order Logic Statements*

In the FOL syntax, three concepts are available, (1) terms, (2) functions and (3) statements. A FOL term is either a constant (e.g., High Street) or a variable (e.g., x). The functions are symbols that either declare terms or describe their relationships – e.g., declaration: Place(High Street) and relation: IN(High Street, Oxford). Finally, the statements are either query statements that return values of variable terms or logical statements which can be either true or false.

In the context of place-related QA, the terms are either place or event concepts. Specific references (proper nouns) are the constants and generic references are the variables. Referring to the introductory example, *High Street* and *Oxford* are the constants, and for the *pharmacies*, variable x0 is assigned to represent the generic reference.

To declare terms, special functions are considered. Two specific functions are PLACE() and EVENT() that describe place and event concepts – e.g., PLACE(HighStreet). Generic references are declared using the assigned variables and the generic place type – e.g., PHARMACY(x0). Predefined rules are considered to link the generic declarations to PLACE and EVENT. In other words, every generic declaration is either a place or an event – e.g., $\forall x\, \text{PHARMACY}(x) \rightarrow \text{PLACE}(x)$

The other set of functions in FOL statements captures the relationships between the terms or describes a single term. This set includes, spatiotemporal relationships, situation/activity relationships, comparisons and qualities. Qualities describe a single term, and are represented by a function with single arguments. For example, *Scottish counties* are represented as $\text{COUNTY}(x0) \wedge \text{SCOTTISH}(x0)$. The other functions have more than one arguments and define relationships between the extracted terms.

In the implementation[5], spatial relationships are either binary or ternary. The three essential constructs of such relationships are *locatum*, *spatial relationship*, and *relatum* (Vasardani et al., 2013). In FOL representation, spatial relationship is the name of the function, and locatum and relatum are the two argument. In case of complex spatial relationships, the additional information is presented as the third argument. In the example, the complex relationship is represented as IN_RADIUS_OF(x0, High Street, 200 meter).

Similarly, situations and activities are represented by functions with the place in relation as the first argument and the object of activity or situation (if present) as the second argument. For example in *where can I buy coffee in Melbourne?*, the activity is represented as BUY(x0, coffee).

Comparisons are also defined as binary relationships with the comparison source as the first argument and target as the second argument. Considering the following example, *Is Melbourne bigger than Sydney?*, the comparison is represented as a FOL function with two arguments: BIGGER_THAN(Melbourne, Sydney).

The extracted conjunction relationships are treated differently when generating FOL statements. If a constituent of a conjunction participates in relationships to other terms, then automatically a similar relationship is applied to the other constituent of the conjunction. For example in *which libraries are near to CBD or Parkville?*, the spatial relationships between libraries and *CBD* is replicated for *Parkville* because of their conjunction relationships – i.e., $\text{NEAR}(x0, \text{CBD}) \vee \text{NEAR}(x0, \text{Parkville})$.

After term declarations and function definitions, the FOL statement is formulated using the identified intent (Section 5.2.4). If

5 available in https://github.com/hamzeiehsan/place_related_question_parser

the question is a yes/no question, then the statement is simply derived by appending term declarations and function definitions. Otherwise, the FOL statement is a query statement that returns the intended terms (e.g., x0) or functions (e.g., COUNT(x0)). The FOL statement for the introductory question is presented in Equation 3.

$$\begin{aligned}&\text{COUNT}(x0) : \text{PLACE}(\text{High Street}) \wedge \text{PLACE}(\text{Oxford})\\&\wedge \text{PHARMACY}(x0) \wedge \text{IN_RADIUS_OF}(x0, \text{High Street}, 200\,\text{meter})\\&\wedge \text{IN}(\text{High Street}, \text{Oxford})\end{aligned} \tag{3}$$

5.2.7 *Generating GeoSPARQL Queries*

To translate logical statements into GeoSPARQL queries, a template-based approach is designed. FOL and GeoSPARQL are declarative formal languages which eases the process of translation from one language to another. However, the level of formalism is different between FOL and GeoSPARQL, and consequently the translation task is not trivial. GeoSPARQL includes prefixes to access ontologies and knowledge bases, ontology-based declarations, and predefined spatial and non-spatial functions that need a careful translation from FOL terms and functions.

GeoSPARQL queries are either *ASK* (yes/no) or *SELECT* (return stored values) queries. This defines the type of the generated query, and completely depends on the intent of the questions. GeoSPARQL queries have several constituents, including *PREFIXES*, *ASK/SELECT* statements, and *WHERE* clauses. PREFIXES define the namespaces from which to access knowledge bases, ontologies, and implemented functions. The ASK/SELECT statements determine the output of the query, and the WHERE clauses capture the criteria mentioned in the question. These constituents are constructed based on predefined templates, described below.

Three sequential steps are proposed to translate FOL statements to GeoSPARQL queries: First, the overall structure of a given query (i.e., ASK vs. SELECT query) is determined from the extracted intent. Second, the WHERE-clause is dynamically generated by concatenating individual concept and relation definition statements. Finally, in case of SELECT queries, the intent of questions determines what variables the inquirer wants to know their value.

A predefined set of prefixes are used to access SPARQL/GeoSPARQL functions and to generate shorter and easier-to-read queries (Query 1).

```
PREFIX geosparql: <http://www.opengis.net/ont/geosparql#>
PREFIX geof: <http://www.opengis.net/def/function/geosparql/>
PREFIX db: <http://spatial.au/ontology#>
PREFIX spatialF: <http://jena.apache.org/function/spatial#>
PREFIX units: <http://www.opengis.net/def/uom/OGC/1.0/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl:<http://www.w3.org/2002/07/owl#>
```

Query 1: Predefined set of prefixes

The WHERE-clause is part of the general structure for both ASK and SELECT queries. To generate the WHERE-clause, for each extracted geographic concept (e.g., *Place* and *Event*) a unique variable is assigned. Each concept is defined using predefined templates and the corresponding variable name. The concept definition statements define a place/event based on the *name* or *type* using an *is-a* relationship. The extracted *non-platial objects* are defined as properties of their corresponding concept variables. The extracted relations among the concepts (i.e., *activities*, *situations*, *comparison* and *spatiotemporal relationships*) are translated to the formal query language using the participating variables, their FOL definitions and the predefined relation templates. Definition statements for place/event concepts using their *name* and *type* are presented in Queries 2 and 3, respectively.

```
?<VARIABLE> a db:<CONCEPT> .
?<VARIABLE> db:name ?<VARIABLE>NAME;
    geosparql:hasGeometry ?<VARIABLE>GEOM.
FILTER(regex(?<VARIABLE>NAME, <NAME>, "i" )).
```

Query 2: Place/event definition template using name

```
?<VARIABLE> a db:<CONCEPT> .
?<VARIABLE> db:type ?<VARIABLE>TYPE;
    ?<VARIABLE> db:name ?<VARIABLE>NAME;
    geosparql:hasGeometry ?<VARIABLE>GEOM.
FILTER(regex(?<VARIABLE>TYPE, <TYPE>, "i" )).
```

Query 3: Place/event definition template using type

The place attributes are captured through *has-a* relationships (see Query 4). In the case of spatial relationships, a lookup table is used to map the spatial prepositions to their corresponding spatial functions implemented in GeoSPARQL. The relationships are constructed by the *predicate* and *argument(s)* which are captured in the corresponding FOL function definition. Similarly, the situation/activities and comparisons are defined using their corresponding templates from their FOL definitions. The template for distance relationships, as examples of the relation templates, is shown in Query 5.

```
?<VARIABLE> db:has_<ATTRIBUTE> ?<OBJECT>.
```

Query 4: Attribute relation template

```
FILTER(geof:distance(<VARIABLE1>, <VARIABLE2>, units:<UNIT>) < <
    DISTANCE>).
```

Query 5: Distance relation template

In the case of SELECT queries, the SELECT-statement can: (1) select some variables, or (2) apply a function (e.g., distance) over the values of selected variables. The selected variables are related to the intent of the questions which is clearly stated in the FOL statement.

Finally, when aggregation (e.g., counting) or sorting (e.g., superlative qualities) is needed their corresponding templates are used. In the special case of aggregation and sorting, GROUP-BY and HAVING statements, and ORDER-BY and LIMIT-statements are concatenated at the end of the generated query, respectively. Sorting is needed when superlative qualities (e.g., longest river) are found and aggregation is necessary when comparison-based situations (e.g., cities that have more than 10 suburbs) are extracted. The template for aggregation and sorting are shown in Queries 6 and 7, respectively.

```
GROUP BY <VARIABLE>
HAVING (COUNT(*) > <LIMIT>)
```

Query 6: Aggregation template

```
ORDER BY <ASC/DESC> (<VARIABLES>) LIMIT <LIMIT>.
```

Query 7: Sorting template

The automatically generated GeoSPARQL query for the introductory example is shown in Query 8.

```
PREFIX geosparql: <http://www.opengis.net/ont/geosparql#>
PREFIX geof: <http://www.opengis.net/def/function/geosparql/>
PREFIX db: <http://spatial.au/ontology#>
PREFIX spatialF: <http://jena.apache.org/function/spatial#>
PREFIX units: <http://www.opengis.net/def/uom/OGC/1.0/>
PREFIX xsd: <http://www.w3.org/2001/XMLSchema#>
PREFIX owl:<http://www.w3.org/2002/07/owl#>


SELECT (COUNT(distinct ?x0) as ?countx0)
WHERE {
        ?c0 a db:PLACE .
        ?c0 db:name ?c0NAME;
                geosparql:hasGeometry ?c0GEOM.
        FILTER(regex(?c0NAME, "High Street", "i" )) .

        ?c1 a db:PLACE .
        ?c1 db:name ?c1NAME;
                geosparql:hasGeometry ?c1GEOM.
        FILTER(regex(?c1NAME, "Oxford", "i" )) .

        ?x0 a db:Place .
        ?x0 db:type ?x0TYPE;
                geosparql:hasGeometry ?x0GEOM;
                db:name ?x0NAME.
        FILTER(regex(?x0TYPE, "pharmacy", "i" )) .

        FILTER(geof:distance(x0GEOM, c0GEOM, units:meter) < 200).

        c0GEOM geosparql:ehCoveredBy c1GEOM.
}
```

Query 8: GeoSPARQL query of the introductory example

5.3 RESULTS AND DISCUSSION

In this section, first the results of encoding and FOL statements are evaluated and next the automatically generated GeoSPARQL queries are evaluated.

5.3.1 *Extraction Results and Logical Statements*

Table 24 shows the results of evaluating encoding extraction. The averaging strategy is called macro averaging that shows the average of precision, recall and f-score of encoding extraction for each question. The count shows the frequency of each information unit in the dataset. For instance, if the count is equal to one that shows there is only one reference to activities in the whole dataset.

As shown in Table 24, the state-of-the-art pretrained models perform well in toponym identification from grammatically correct NL questions. Moreover, part-of-speech rules are successful in identifying numbers, comparisons and conjunctions. The dependency parsing and constituency parsing are also proven to be

Table 24: Extraction performance

Semantic Type	Average precision	Average recall	Average f-score	Count
Place name	98.9	99.5	99.2	263
Place type	100.0	99.4	99.7	178
Event name	–	–	–	0
Event type	–	–	–	0
Non-platial objects	95.9	97.9	96.9	55
Number	100.0	100.0	100.0	26
Spatiotemporal relationship	100.0	96.8	98.4	191
Situation	94.2	98.0	96.1	51
Activity	50.0	100.0	66.7	1
Place quality	100.0	91.7	95.7	25
Object quality	85.7	100.0	92.3	6
Comparison	100.0	100.0	100.0	24
Conjunction	100.0	100.0	100.0	6

successful in identifying more complex encoding classes such as spatiotemporal relationships and qualities.

Table 24 shows that event-based place-related questions are completely missing in the dataset. However, the extraction of events from the questions is also expected to perform well considering the performance achieved by the fine-grained NER models (Lample et al., 2016). Moreover, activities are also rarely observed in the dataset, which is another limitation of Geospatial Gold Standard dataset (Punjani et al., 2018, 2021). However, such limitations have minimal impacts on the proposed method due to its reliance on place conceptualization as the grounding theory.

FOL statements are constructed by declaration of places, events and non-platial objects, and functions that describe or relate these terms based on spatiotemporal relationships, situations, conjunctions and comparisons. Table 25 shows the precision, recall and f-score of FOL term declarations and function definitions. The results show that declarations, spatiotemporal relationships and comparisons can be formulated in FOL statements with high precision and recall. Hence, the combination of grammatical rules from dependency parsing and constituency parsing are proven to be successful.

Table 25 shows that qualities are defined as FOL functions with high precision. Yet the recall is lower (i.e., 82.8%), and the reason is that using constituency parsing alone may not be sufficient for all cases. For example in *Which site of Manchester is the most popular?*, the long distance relation between the quality (*most popular*) and the place type it describes (*site*) cannot be captured through constituency parsing.

Situations are also formulated as FOL functions with high precision and lower recall (i.e., 64.7%), because of errors in dependency parsing which fail to link situation verbs to place

Table 25: FOL statements evaluation

Logical statement	Average precision	Average recall	Average f-score
Declaration	98.2	99.7	98.9
Spatiotemporal relationships	97.6	91.9	94.7
Quality	100.0	82.8	90.6
Comparison	91.7	91.7	91.7
Conjunction	80.0	100.0	88.9
Situation	100.0	64.7	78.6

attributes (i.e., non-platial objects) in some cases. Hence, when the link is not captured, the situation is missing and the recall is lower. On the other hand, the high precision shows that whenever the dependency parsing identifies the relationships between situations and place attributes, the corresponding FOL functions are generated correctly.

In contrast to qualities and situations, conjunctions are reflected in the FOL statements with high recall and lower precision (i.e., 80%). The conjunctions mentioned in Section 5.2.6 are reflected by applying the same FOL functions for both sides of the conjunction which are bounded by either *and* or *or* logical operators. Hence, errors in defining other functions (e.g., spatial relationships) can be propagated through conjunctions. Consequently, errors in other functions have an impact on the precision of conjunctions as well.

5.3.2 *Query Generation Results*

Table 26 shows the results of analysing automatically generated GeoSPARQL queries. In overall, more than 80% of the GeoSPARQL queries are completely correct. The evaluation is performed manually to check whether the generated query is syntactically and semantically correct. That is a considerable improvement on the current benchmark results of the same question dataset (Punjani et al., 2018, 2021) (noting that due to the different evaluation metrics and query generation approaches, their results are not directly comparable with the proposed method). The results of this chapter also overcome the performance of an under-review method by Li et al. (2021, Under Review) that achieved 71% overall accuracy for query the generation task.

In terms of the number of questions that can be handled, the benchmark covers only 43% questions in the dataset, while the proposed approach can deal with all questions. Of all the questions that can be handled by Punjani et al. (2018) (43% of questions), the GeoQA method successfully generates answers for 51% of the questions. While in the proposed approach more than 80% of

Table 26: GeoSPARQL queries evaluation

GeoSPARQL query	Average precision	Average recall	Average f-score
Overall (Query)	82.8	–	–
Intent (SELECT/ASK)	97.5	–	–
Criteria (WHERE)	84.3	–	–
Aggregation (Group By)	100.0	83.3	91.5
Sorting (Order By)	92.9	89.7	91.2

GeoSPARQL queries are correct for all questions in the dataset. However, benchmark results also include ontology mapping and fetching the query results from an integrated knowledge base that is not considered in this chapter. The ontology mapping aims at matching extracted information to predefined ontologies in existing knowledge bases. Fetching the results is also important because even with correctly generated queries, the data may not be found in existing knowledge bases.

The major difference between the proposed approach and the benchmark method is the use of place-based conceptualization in translating questions to queries. Using the conceptualization, the proposed method covers more diverse types of place-related questions and consequently is able to handle all questions in the dataset. Using the place conceptualization, the proposed method is not biased specific types of questions and can handle missing question types in the dataset such as event-based questions – e.g., *What cities are affected by hurricanes in the US since 1980?*

Table 26 shows that integrating dependency parsing and constituency parsing is useful to extract the relationships between encodings. The benchmark method only use dependency parsing to extract information from the content of the questions. The improvement of the results shows that using phrase level information from constituency parsing can improve the quality of information extraction.

The other advantage of the proposed approach is the translation of questions to FOL statements that can capture both intent and criteria of NL questions in machine-digestible manner. Using FOL statements as the intermediate representation, the NL questions can be translated to other formal query language(s) with minimal effort – e.g., translating questions to *SPATIAL SQL*.

Table 26 shows the evaluation results for the whole query or specific parts such as WHERE clause. All GeoSPARQL queries include intent (SELECT/ASK statement) and criteria (WHERE clause). Hence, the recall and f-score is meaningless in evaluating these parts – i.e., 100% recall, because these parts are never missing. On the other hand, sorting (29 questions) and aggregation (6

questions) are only applicable for some specific questions. Consequently, the recall and f-score are reported for sorting and aggregation (Table 26).

Table 26 shows that the intent heuristic performs well and consequently the mistakes in formulating the criteria are the main reason for incorrect GeoSPARQL queries. Criteria are more complex to compose than intent and may include an arbitrary number of criterion which all should be formulated correctly. Hence, even a minor mistake in formulating one criterion in a WHERE clause produces an incorrect query.

5.4 CONCLUSIONS

To enable spatial analysis using structured spatial data in GeoQA, NL questions must be translated to formal queries. In this chapter, a method was devised using the object-based conceptualization of place in information science (Purves, Winter, and Kuhn, 2019). First, the essential information was extracted from the questions and encoded based on their types. Next, the relationships between the extracted encodings were extracted using both dependency parsing and constituency parsing. Then, the concepts were identified based on the extracted encodings and their relationships and expressed as FOL statements. Finally, the FOL statements were translated to GeoSPARQL queries using a template-based approach.

The novelty of the proposed method was mainly to (1) utilize domain knowledge (place conceptualization) in question analysis, (2) use both dependency parsing and constituency parsing in relation identification and (3) propose FOL as the intermediate representation. Using domain knowledge, the proposed method is proven to be less biased to the dataset and can analyze diverse types of place-related questions. Moreover, the method is grounded theoretically in previous works and avoids re-investing new terms and concepts.

Using part-of-speech tagging, named entity recognition, and dependency and constituency parsing methods, the quality of information extraction and relation identification was improved in comparison to the benchmark method. The results show that the state-of-the-art neural network models perform well in named entity recognition and part-of-speech tagging. Moreover, the grammatical rules based on the dependency and constituency parsing proven to be accurate in digesting place-related questions.

Finally, using FOL statements as the intermediate representation increases the flexibility and applicability of the method. In this chapter the translation to GeoSPARQL using a template-based approached was investigated and demonstrated. Yet, several steps of the method (i.e., encoding, parsing and generating FOL statements) can be used in translating the questions to other formal query

languages such SPATIAL SQL. The FOL statements have two main advantages which ease such translations. First, the statements are machine digestible and can also be empowered with logical reasoning. Second, these statements can capture both intent and criteria of the NL questions which is also required in translating NL questions to other formal query language(s).

The proposed method is only able to generate GeoSPARQL queries. To be used in real-world application, the ontology mapping step should be developed and tested. Several challenges including entity and relation disambiguation, and ontology matching can be resolved using available QA toolkits such as Frankenstein platform (Singh et al., 2018).

In this chapter, the translation of NL questions into formal query languages was investigated. Since this is not the final stage of question answering, the answer formulation is still a necessary step ahead. In the next chapter, human-generated answers to a special kind of place-related questions, where-questions, are investigated in detail to complete the workflow of question answering.

6 ANSWERING WHERE QUESTIONS

The answering of questions is not a simple retrieval and response of stored information; rather the process is embedded in a general structural framework containing knowledge of the questioner, the question, and the world around it.

— Norman (1973)

6.1 INTRODUCTION

Consider the following question and its corresponding answer, taken from the MS MARCO dataset v2.1 (Nguyen et al., 2016):

Question: *Where is Putney Bridge?*
Answer: *Putney Bridge is a bridge crossing of the River Thames in west London.*

This *where*-question is answered using a *place description* – a description that characterizes the location of interest (Putney Bridge) based on a few anchor places (River Thames and London). Place descriptions, however, are not the only way to answer where-questions. Where-questions can also be answered via other representations such as maps or sketches (Church et al., 2010). Invariant to the chosen representation, the answers localize the place in question based on its spatial relationships with chosen anchor places (Couclelis et al., 1987). Hence, answering where-questions poses the following challenges no matter what representation is used:

- Generating informative answers – i.e., the answer should complete the inquirers' gap of knowledge in a way that obvious or already-known responses should be avoided and useful and necessary information are included (Shanon, 1983). In the example, obvious, inadequate or undetailed answers

o This chapter is based on a journal paper (Hamzei, Winter, and Tomko, 2021) and a conference paper (Hamzei, Winter, and Tomko, 2019). My contribution as a first author of these publications includes proposing the hypothesis and research problem, designing the method, analysing the results and presenting the discussions and conclusion. These works were supervised by *Prof. Stephan Winter* (my primary supervisor) and *Dr. Martin Tomko* (my co-supervisor) who provided feedback and suggestions.

such as *on Earth* or *in the UK* or *over a river* are avoided by the responder.

- Answering the question in a cognitively efficient manner (Wilson and Sperber, 2002) – e.g., producing short and straightforward place descriptions (Hamzei et al., 2020) and personalized map labelling strategies in map visualizations (Wilson, 2018). In the example, the responder excludes unnecessary information such as the nearby theaters and restaurants to keep the answer as simple and relevant as possible.

- Determining the level of granularity of answers – e.g., a suitable zoom level for maps (Ballatore, 2019) and referring to places of suitable granularity in place descriptions (Hamzei, Winter, and Tomko, 2019). In the example, the name of the roads and streets that are connected to the bridge are neglected in the answer based on the judgement of the responder for the relevant scale level.

- Selecting places that can be assumed to be known by the inquirer – e.g., labelling the places known to inquirers in maps (Suomela, Lakkala, and Salminen, 2009) and referring to them in place descriptions as anchors. In the example, the location of *River Thames* and *London* are assumed to be known to the inquirer.

Where these challenges are met by an answer, the communication succeeds.

Addressing these challenges is a necessary step towards answering where-questions. Corpora of human-generated answers can be used as an information source that reflects the selectivity of choosing anchor places in answering where-questions. The analysis of human-generated answers can be later used to imitate such selectivity in formulating automatic responses. The results can be applied for generating answers to where-questions as NL responses (place descriptions). Selecting relevant anchor places is an essential part of generating place descriptions that succeed in answering where-questions. Moreover, information about anchor places can be used in static maps to be visualised in a proper context frame (Ballatore, 2019).

In recent years, several GeoQA studies are conducted for answering geographic questions. Previous studies provide solutions to retrieve responses from knowledge bases (Stadler et al., 2012) and documents (Buscaldi et al., 2006; Luque et al., 2006). GeoQA studies are mostly focused on what/which questions about geographic places (e.g., Scheider et al., 2020b; Vahedi, Kuhn, and Ballatore, 2016). In answering where-questions, the task is either simplified

into retrieving stored coordinates (Luque et al., 2006; Stadler et al., 2012), or selecting a part of text without explicit adaptation to the question (Buscaldi et al., 2006).

Current geographic question answering systems are often focused on coordinate retrieval as answers to where-questions (e.g., Luque et al., 2006; Stadler et al., 2012). While coordinates are useful for communication between location-based services to perform spatial analysis or visualisation (Jiang and Yao, 2006), it is not necessarily a relevant response to inquirers without a proper map visualization. Yet, a characterization of relevant anchor places to localize a place in question is still missing.

To enable computers to provide responses with similar qualities to human-generated answers, the responses need to be relevant. An answer is relevant if its positive cognitive effects to inquirers are large and the processing effort to achieve the effect is small (Wilson and Sperber, 2002). In other words, answers should be informative enough and as straightforward as possible. Assuming human-generated answers are relevant responses, machine-generated responses should imitate the selectivity in human-generated answer to provide useful pieces of information and avoid unnecessary ones. Generating such relevant responses is the major prerequisite of intelligent GeoQA as defined by Winter, 2009.

In this chapter, *human*-generated answers to where-questions are investigated to inform the properties of such answers and to devise and test a method to imitate their structures in machine-generated responses. To achieve these goals, the information in a where-question and its answer is modelled as *an ordered set of places* that are mentioned in their content. Then the properties of places in questions and corresponding answers are derived and further investigated. This model forms a template (i.e., an ordered set of place properties) that enables computers to learn and imitate human answering behaviour. In other words, place properties are utilized to understand why a set of places are chosen as anchors to localize the place in question and how this selectivity can be imitated by computers.

The properties that are used in the templates are generic geographic information that describe the shared meaning of places in form of generic types from a finite set of categories. Referring to the example above, the place in question is a *bridge* which is localized by referring to the river it goes over and the city it belongs to. Here, the template captures the structure of the answer as relationships between *bridges and rivers*, and *bridges and cities*. The generic information captures shared meaning of geographic places. While generic geographic information is not used in QA, it has been used to investigate and characterize place descriptions (Edwardes

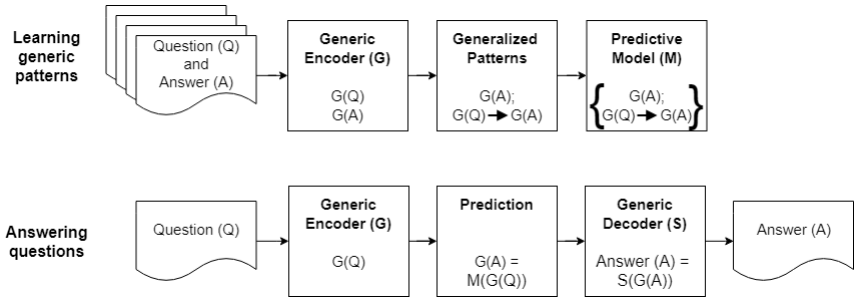

Figure 21: Specific-generic translation approach

and Purves, 2007a; Richter et al., 2013), route descriptions (Raubal and Winter, 2002), and regions (Tomko and Purves, 2008).

6.2 METHODOLOGY

In this chapter, a generalized approach of *specific-generic translation* is devised to study human answering behaviour using generic properties. Next, a method using specific-generic translation is devised to investigate the QA in interaction of people with a general-purpose search engine. Other QA scenarios (e.g., human-human dialogue, human-social bot interaction) may require different design of specific-generic translations.

6.2.1 *Specific-Generic Translation*

Figure 21 shows the proposed approach to derive characterizing patterns in human-generated answers and generating template to answer where-questions. The approach includes two stages: (1) *learning generic patterns* where the objective is to investigate and to characterize human answering behavior into a machine learning model, and (2) *answering questions* where the model is used to generate answers. The novelty of the approach is in the encoding of questions and answers into their generic meaning and to model the relation between questions and answers in their generic forms. Later, the model is used to generate generic forms of answers to where-questions. Finally, the answer is constructed by decoding generic form (e.g., type) of the answer into its specific representation (i.e., toponym).

The specific-generic translation approach involves the following steps:

1. Selecting a set of generic information classes (e.g., place type, scale, prominence, functions and access) based on the context of QA and availability of data;

2. Defining a schema for each selected generic information class;

3. Designing an information extraction approach to encode the questions and answers into generic forms (*Generic Encoder* in Figure 22);

4. Evaluating how effective each generic class is in capturing the relation between the questions and their human-generated answers (*Generalized Patterns* in Figure 22). The results of evaluation also provide insights about human answering behavior in the context of the QA problem.

5. Training a predictive model that can learn generalized patterns of human-generated answers (*Predictive Model* in Figure 22);

6. Defining a decoding approach to map generic forms of answers into specific (toponym) representation (*Generic Decoder* in Figure 22). This step can be followed by additional steps such as natural language generation to be used in real-world applications.

In this chapter, the results of the first five steps are discussed in details and the last step is only demonstrated using examples.

6.2.2 *Type, Scale and Prominence (TSP) Encoding*

Based on the specific-generic translation, TSP encoding is proposed to investigate where-questions constructed only with toponyms. The generic forms which are used to investigate these questions and their answers are *type*, *scale* and *prominence* of the toponyms. In the followings, first the terminology is presented and next the proposed TSP encoding method is discussed in detail.

6.2.2.1 *Definitions*

The investigated types of where-questions are defined as:

- **Toponym-Based Where-Question (TWQ)**: A toponym-based where-question is a geographical where-question that is generated completely using toponyms. For example, *Where is Beverly Hills?* is a toponym-based where-question, while *Where is Clueless filmed?* (without toponym) and *Where can I buy furniture?* (affordance, buying furniture) do not belong to this type.

- **Simple Where-Question (SWQ)**: Simple where-questions are a sub-class of TWQs that contains only one toponym in their body (e.g., *Where is Beverly Hills?*).

- **Detailed Where-Question (DWQ)**: Detailed where-questions are a sub-class of TWQs with more than one toponym in their

content (e.g., *Where is Beverly Hills, California?*). In DWQs, contextual details are provided in the content of the questions that shows what the inquirer already know – e.g., Beverly Hills is located in California.

Type, *scale* and *prominence* are defined as:

- **Type**: A type (e.g., restaurant, mountain) is a reference to a group of places with similar characteristics (e.g., affordance, or physical properties). Type defines similar places and differentiates dissimilar ones, sometimes in a hierarchical or taxonomical manner. In this chapter, the relation between a specific reference to a place (unambiguous toponym) and its type is considered as a one-to-one relation.
- **Scale**: Scale is defined as a finite hierarchically-organized ordinal set of levels grounded in the human understanding and differentiation of places based on their size and their relationships (i.e., containment). The relation between scale and an unambiguous toponym is considered as one-to-one. Due to the specific context of the QA scenario, very fine levels of scale of geographic entities, such as room-level or object-level, can be neglected here, while in everyday human-human communication these levels of scale may have a more important role.
- **Prominence**: Prominence is a measure of how well-known a place is. In this research, prominence is characterized by a finite ordinal set of levels. While prominence of places is subjective and differs from person to person based on their experience, here prominence is considered as an absolute and objective measure to rank places, established through a proxy measure defined later. This approach enables to avoid individual experiential biases and is supported by the evidence of success in day to day communication in which the absolute prominence evaluation is adapted between hearers and speakers.

Type, scale and prominence are used to characterize place descriptions (Edwardes and Purves, 2007a; Richter et al., 2013). These geographic concepts can be used to capture different kinds of relationships among places. These relationships can be used to understand the relation between where-questions and their answers. For example, such relationships between rivers and seas (*flows to*), and cities and countries (*part of*) can be captured using place type. Considering the relation between where-questions and their answers, *containment* (different levels) and *nearness* (a same level) can be captured through differences among scale levels. Finally, prominence is a measure to check whether the answer is

interpretable by the inquirers – i.e., more prominent places are expected to be better known by inquires.

The aspects of human-generated answers which are investigated in this chapter are defined below:

- **Content**: The content of an answer is a collection of distinct information units that are presented to satisfy the inquirer's information need. Content can be generic (e.g., type) or specific (e.g., toponym). Content is the most important aspect of the answers, in a way that the difference between correct and incorrect responses are completely based on their content.
- **Style**: The style of an answer is the way that the content is presented. Style directly influences the perception of naturalness of the response. Referring to the introductory example, *... Putney Bridge is a bridge crossing of the River Thames in west London* and *... Putney Bridge is a bridge in west London which goes over River Thames* are two different styles of answers (with same content) to the question. Here, the former is preferred by the responder.

6.2.2.2 *TSP Sequences*

In TSP encoding, sequence representations of generic/specific information are selected to model the questions and answers. A sequence is defined as an ordered set of items (here, references to generic/specific geographic information). First, questions and their corresponding answers are modelled as sequences of toponyms (specific representation). Then, these toponym sequences are encoded into type, scale and prominence sequences by translating each specific toponym into its corresponding generic type, scale and prominence reference. Referring to the introductory example, the specific representations (toponym sequences) and the encoded type sequences (an example of a generic sequence) of question and answer are presented below:

- **Toponym sequences**: [Putney Bridge] [River Thames, London]
- **Type sequences**: [bridge] [river, city]

Here, the *content* refers to the information items in the sequences, and their order defines the *style* in which the information is presented.

6.3 IMPLEMENTATION

Figure 22 shows the proposed workflow[1] to investigate TWQs and their answers. In the following sections, the details of the dataset,

1 Additional details of implementation are presented in Appendix C, Section C.1

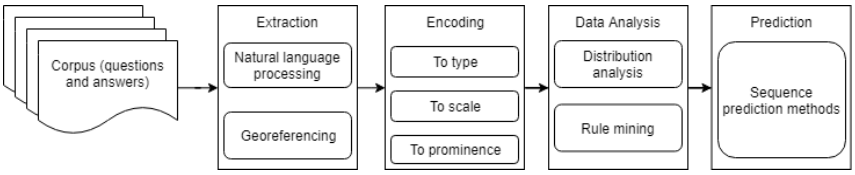

Figure 22: The proposed implementation approach

and extraction, encoding, pattern mining and prediction methods are discussed. A complete implementation of the proposed TSP encoding approach also includes decoding from generic to specific information. Here, the decoding step is demonstrated through examples, and a fully automated implementation remains out of scope of this chapter.

6.3.1 *Data*

The questions in MS MARCO v2.1 (Nguyen et al., 2016) are categorized into five categories using tags: (1) *numeric*, (2) *entity*, (3) *location*, (4) *person*, and (5) *description* (Nguyen et al., 2016). Geographic questions can thus be easily extracted using the predefined *location* tag. The dataset contains over one million records divided into *training*, *development* and *testing* subsets.

Each record in the dataset includes a *question*, *human-generated answer(s)* (except for records in the *test* dataset, where the answers are deliberately excluded), a *tag*, and a *set of documents* retrieved by the Microsoft Bing search engine[2].

The 'location questions' in MS MARCO (56,721 question-answer pairs) include 36,939 geographic questions, and the remainder are questions about fictional, mystic and other non-geographic places (Hamzei et al., 2020). Among the geographic questions, 13,195 pairs of questions and answers are geographic where-questions (Hamzei et al., 2020).

There are several reasons to choose MS MARCO for this study considering other available datasets such as SQuAD (Rajpurkar et al., 2016):

- MS MARCO is the largest available QA dataset;
- The questions are labelled and geographic questions can be easily extracted;
- All questions are asked in a specific real-world scenario (i.e., Web search);
- Inquirers pose questions to resolve their information needs while in some datasets such as SQuAD, questions are made

2 More information about the dataset can be found in: https://github.com/dfcf93/MSMARCO

from documents. In other words, questions in SQuAD is more about what a document can answer rather than what actual inquirers want to know.

- The answers are provided using an open form strategy. The answerers can utilize suggested documents (one or more) and their own knowledge to answer a question. Hence, the answers are not directly extracted as a single span of a document.

6.3.2 *Extraction*

First, the questions are extracted using the *location* label when they start with a *where*-token. Next, the text is the toponyms inside the questions and answers are identified using Named Entity Recognition (NER) and gazetteer lookups using both OSM Nominatim and Geonames gazetteers. Here, the Stanford NLP toolkit is used to extract named entities (Finkel, Grenager, and Manning, 2005). In this step, if a compound noun phrase is tagged as location, first the compound noun is checked by gazetteer lookup; if it is not identified, then its constituent simple nouns are checked. If a compound or simple noun phrase is found in both gazetteers, it is stored as a toponym. For the extracted toponyms, the gazetteers records are retained if (1) the OSM and Geonames records have same name, and (2) the Geonames' point-based coordinates are inside the region-based representation of their corresponding OSM records.

The toponym disambiguation is then undertaken based on the *minimum spatial context* heuristic (Leidner, Sinclair, and Webber, 2003). Here, the bounding boxes are used to determine which combination of the geographic locations satisfy the minimum spatial extent condition. In cases of duplicate places in GeoNames which lead to the same bounding boxes, the combination with more specific place types is selected. For example, populate place (PPL) is a place type in GeoNames which could be a village, city, state and even country. Hence, administrative divisions (e.g., state) are chosen over the populated places. Finally, if the toponym ambiguity still exists, the importance value is used to select the more prominent combination.

More sophisticated heuristics in toponym disambiguation (e.g., Lieberman and Samet, 2012; Wang et al., 2010) are not used due to reliance on significant assumptions – e.g., the relation between place types in the toponyms, city-country relation. These heuristics constrain the relationships between type, scale and prominence of resolved places in the text. This may impact the results of this study and lead to stronger associations based on type, scale and prominence between toponyms in questions and answers. Here, to present fair results, these disambiguation methods are avoided.

6.3.3 *Encoding*

The gazetteers' attributes for the extracted toponyms have been used as proxies to capture type, scale and prominence of the toponyms in questions and answers. Using these proxies, the sequence representations for each question-answer pair are encoded into TSP sequences.

Type: The Geonames type schema[3] has been used without modification to encode generic place types. This schema contains 667 unique types of geographic features, covering both natural and man-made geographic types.

Scale: To capture scale, the schema from Richter et al., 2013 are extended in this study. This schema contains seven levels of granularity: (1) furniture, (2) room, (3) building, (4) street, (5) district, (6) city, and (7) country. The coarse levels of granularity are extended by adding *county*, *state*, *country*, and *continent*, and removing the *furniture* and *room* levels from the schema. OSM records include an attribute related to the OSM definition of scale (i.e., place_rank[4]), a number between 0 to 30. The extracted gazetteers' records are translated into the appropriate scale level based on a look-up table that maps OSM scale levels into the proposed scale schema (see Appendix C, Section C.1).

Prominence: To capture prominence, the *importance* attribute in the extracted OSM Nominatim record is used. The OSM importance value is estimated using Wikipedia importance score (Thalhammer and Rettinger, 2016) with some minor tweaks[5]. The value is defined between 0 and 1, and it is designed to be used for ranking search results. These values are translated into seven finite levels of prominence, derived by *natural breaks* classification (Jenks, 1967) of the frequency spectrum of the values.

6.3.4 *Distribution Analysis and Rule Mining*

Distribution analysis and rule mining techniques are used to extract and investigate patterns in the human-generated answers and the relation between the questions and their answers. Distributions of type, scale and prominence sequences are used to compare the questions and answers, and SWQs versus DWQs. To derive patterns in the questions and their answers, association rule mining, a-priori algorithm (Agrawal and Srikant, 1994), is used.

The strength of the extracted rules are evaluated using the standard measures – i.e., *support*, *confidence*, and *lift*. Support defines

3 https://www.geonames.org/export/codes.html
4 https://wiki.openstreetmap.org/wiki/Nominatim/Development_overview
5 https://lists.openstreetmap.org/pipermail/geocoding/2013-August/000916.html

how frequently an association rule is observed in the whole dataset, and confidence determines how often the rule is true. Lift is a measure to evaluate the importance of the rules – i.e., lift greater than one shows positive and strong dependency among the elements of the extracted rule.

6.3.5 *Prediction*

The input for the prediction is an encoded sequence of TWQs, and the output is the generic sequence of their corresponding answers. The problem can then be formulated as a sequence prediction from concatenated generic sequences for the questions and their answers, where a part of a sequence is known, and the rest is predicted. Table 27 shows the sequence prediction methods which are used in this study. In the implementation, an open-source toolkit for sequence analysis (Fournier-Viger et al., 2016) is used and extended to implement the prediction methods.

These classic methods are divided into probabilistic (Cleary and Witten, 1984; Padmanabhan and Mogul, 1996; Pitkow and Pirolli, 1999) and non-probabilistic categories (Gueniche et al., 2015; Gueniche, Fournier-Viger, and Tseng, 2013; Laird and Saul, 1994; Ziv and Lempel, 1978). The probabilistic methods are based a graph representation of conditional probabilities (Cleary and Witten, 1984) or Markov chain's transition probability matrix (Padmanabhan and Mogul, 1996; Pitkow and Pirolli, 1999) of the sequence elements. The non-probabilistic methods compress the sequences in a lossy (Laird and Saul, 1994; Ziv and Lempel, 1978) or lossless approaches (Gueniche et al., 2015; Gueniche, Fournier-Viger, and Tseng, 2013) into tree-based (Gueniche et al., 2015; Gueniche, Fournier-Viger, and Tseng, 2013) or graph-based (Laird and Saul, 1994) data structures (for a review of sequence prediction methods see Tax, Teinemaa, and Zelst, 2020).

The structure of sequence and the relation between prior elements in the sequence to their succeeding elements are trained into a model. The model is then tested on an unseen part of data using K-fold cross validation (K=10). Two baseline methods are considered to evaluate the performance of the sequence prediction methods: (1) random sequence generation and (2) most frequent pattern.

The random generation baseline only utilizes the schema of type, scale and prominence without any information about the distributions of values in the answers. The most frequent patterns baseline predicts templates of answers using the schema and the distribution of generic references in the answers. The difference between the prediction performances of random generation and the most frequent patterns shows the impacts of using the distribution of generic values in generating templates of answers. The sequence

Table 27: Sequence prediction methods

Method	Publication	**Year**
Lampel-Ziv 1978 (LZ78)	(Ziv and Lempel, 1978)	1978
First order Markov Chains (Mark1)	(Cleary and Witten, 1984)	1984
Transition Directed Acyclic Graph (TDAG)	(Laird and Saul, 1994)	1994
Dependency Graph (DG)	(Padmanabhan and Mogul, 1996)	1996
All-k-Order Markov Chains (AKOM)	(Pitkow and Pirolli, 1999)	1999
Compact Prediction Tree (CPT)	(Gueniche, Fournier-Viger, and Tseng, 2013)	2013
Compact Prediction Tree Plus (CPT+)	(Gueniche et al., 2015)	2015

prediction methods also consider the relation between generic values in the questions to their answers. Consequently, the improvement in generating the templates compared to the most frequent patterns baseline is related to the association between generic values of questions and their answers.

In prediction, each generic form of questions is used to predict the same generic form of their answers. In addition, an approach is devised to predict one of the generic forms of an answer using all generic forms (i.e., type, scale and prominence) of its corresponding question. Algorithm 1 shows the process to use all three type/scale/prominence sequences to predict a generic form of the answers in each generic class. Here, each combination of type, scale and prominence values are mapped to a unique code. Using these codes, a new sequence is generated for each question/answer to capture type, scale and prominence together. Next, these sequences are used to predict the generic form of answers. Finally, a reverse mapping is used to decode these sequences into type, scale and prominence sequences.

Algorithm 1 Training and prediction based on type-scale-prominence together

1: **procedure TSP_Prediction**(type, scale, prominence)
2: generate a code for each unique combination of type-scale-prominence (TSP)
3: create encoded sequences based on generated TSP codes
4: train a model to predict TSP in answers based on TSP in the questions
5: **for** every question **do**
6: given a question (TSP); predict the answer (TSP)
7: decode the predicted answer (TSP) to answer (type/scale/prominence)
8: **if** multiple predictions are allowed **then**
9: avoid counting duplicate decoded values for type/scale/prominence
10: **end if**
11: **end for**
12: **end procedure**

Table 28: Extraction evaluation

Extraction	#Extracted	#Investigated	Precision	Recall
TWQs	6274	314 (5%)	91.7% (288 out of 314)	–
SWQs	3285	121 out of 288	89.4%	90.2%
DWQs	2989	167 out of 288	92.7%	92.1%
Toponyms	22307	1133[6]	88.6%	90.8%

Table 29: Encoding results

Encoding	#TWQs	#SWQs	#DWQs
Type sequences	6,274	3,285	2,989
Scale sequences	3,936	1,985	1,951
Prominence sequences	6,051	3,098	2,953

6.4 RESULTS

6.4.1 *Extraction and Encoding*

The assessment of toponym extraction, finding TWQs, and categorizing the questions into SWQs and DWQs are presented in Table 28. Here, average precision and recall of the extraction results are calculated using manually annotated data (5% of TBWQs and their answers). For the task of finding TWQs in the dataset, the *false negatives* (TWQs that have not been extracted) are not investigated, hence the recall is unknown.

As shown in Table 28, 6,274 TWQs and their answers are found in the dataset. The TWQs are approximately 11.1% of the *location questions* of the dataset. For evaluation, 5% of extracted TWQs (314 questions) are investigated and the precision of extraction is 91.7% – i.e., 288 of 314 extracted questions are TWQs. Using the 288 TWQs, the precision and recall of extracting toponyms and classifying the questions to SWQs and DWQs are presented in Table 28.

Table 29 shows the number of records that are completely encoded for question-answer pairs in type, scale and prominence sequences. Here, if even the information for one place (which is mentioned either in the question or its answer) is missing, the question and its answer are not used to extract patterns or test the predictability of generating generic form the answer. As shown in the table, the encoding into scale and prominence is not always possible due to incompleteness of attribute information (i.e., *place_rank* and *importance*) in OSM Nominatim.

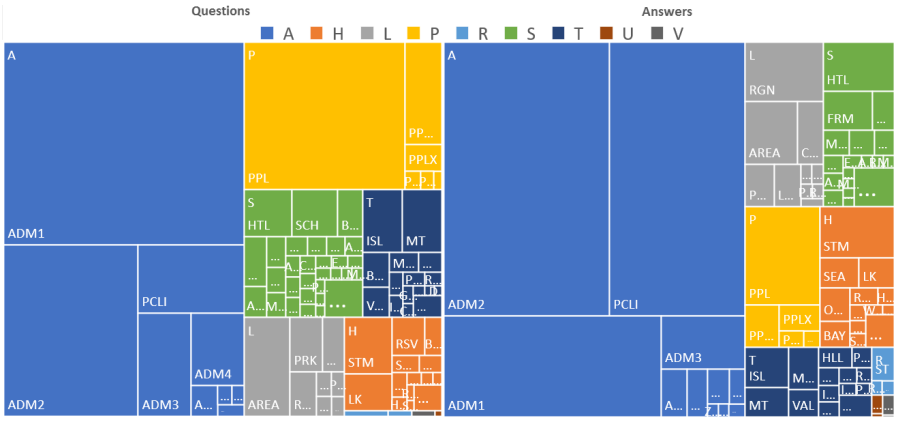

Figure 23: Distribution of place types in the questions and in the answers.

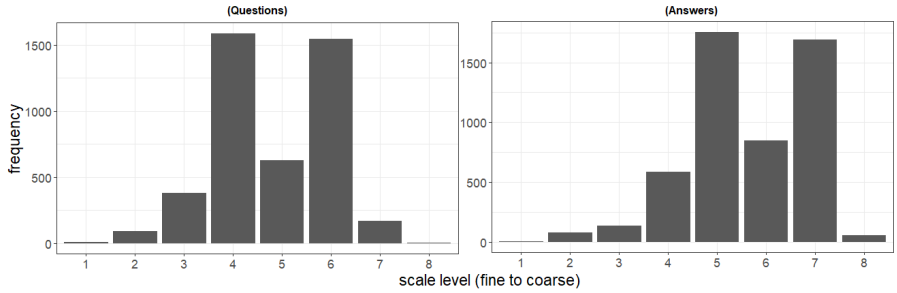

Figure 24: Distribution of levels of scale in all toponym-based where questions and answers.

6.4.2 *Distributions*

The distribution of TWQs and their answers based on type, scale and prominence are shown in Figures 23[7], 24 and 25. Figure 23 shows that the diversity of types in the questions is higher than in the answers. While administrative divisions are more frequent than other generic types in both questions and answers, they are more dominant in the answers.

Figure 24 shows the scale in the answers is systematically one level coarser than in the questions. In addition, the distribution shows that city-level and state-level scales are frequently observed in the questions, while the answers mostly contain references of county and country levels of scale. The results further show that the coarsest level of scale (i.e., continent level) is rarely observed in the answers. This observation shows an answer at the continent level would be under-specified in most cases, and therefore uninformative.

The distributions of prominence levels in questions and answers are similar to the distributions by scale (Figure 25). In the questions, a bi-modal distribution of levels of prominence in the content of questions cab be observed. The distribution of prominence in the

7 The codes for place types are described in AppendixC

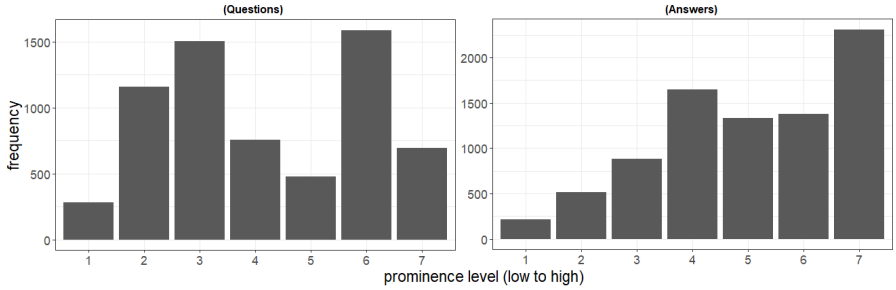

Figure 25: Distribution of prominence levels in the questions versus answers

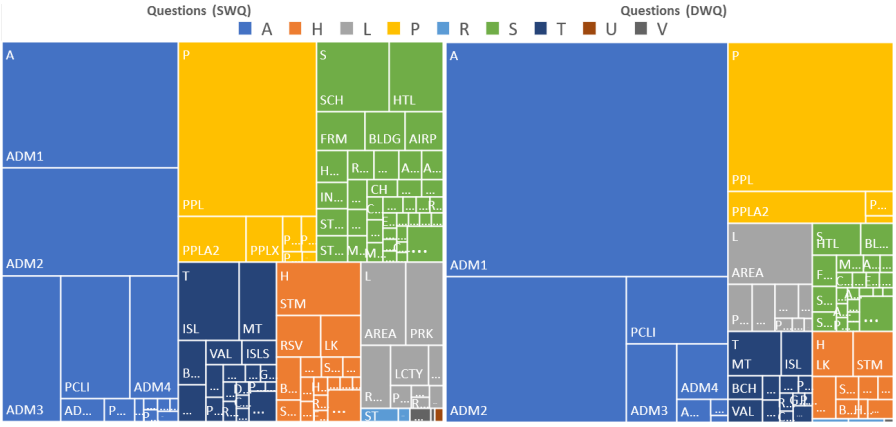

Figure 26: Distribution of types in the SWQs versus DWQs.

answers, however, shows that higher levels are dominant. In contrast to the distributions by scale, the most prominent level is dominant in the answers. Hence, people tend to refer to well-known places in their answers. Unlike with scale, the highest levels of prominence do not necessarily lead to obvious or irrelevant answers.

Figure 30 shows the distributions of place types in SWQs and DWQs. In DWQs, administrative divisions are almost twice as frequent as in SWQs. The extra detail (additional toponyms) included in the content of DWQs dominantly refers to administrative divisions. For example, *Where in California is Disneyland?* can be viewed as a SWQ, *Where is Disneyland?*, with extra details that shows what the inquirer already knows about its coarse location (*in California*). Hence, the differences between the two distributions illustrate where the additional detail is included.

Answers to SWQs (Figure 31) have a linearly incremental distribution (except for the coarsest level of scale). The figure shows the answers are generated with toponyms that belong to *coarser levels* of scale compared to places mentioned in the SWQs. The distribution of scale in the answers to DWQs shows a dominance of references to places of scale *between* the two mean values of the questions' bi-modal distribution. Thus, people tend to generate responses at a level of scale greater than that of the place that is

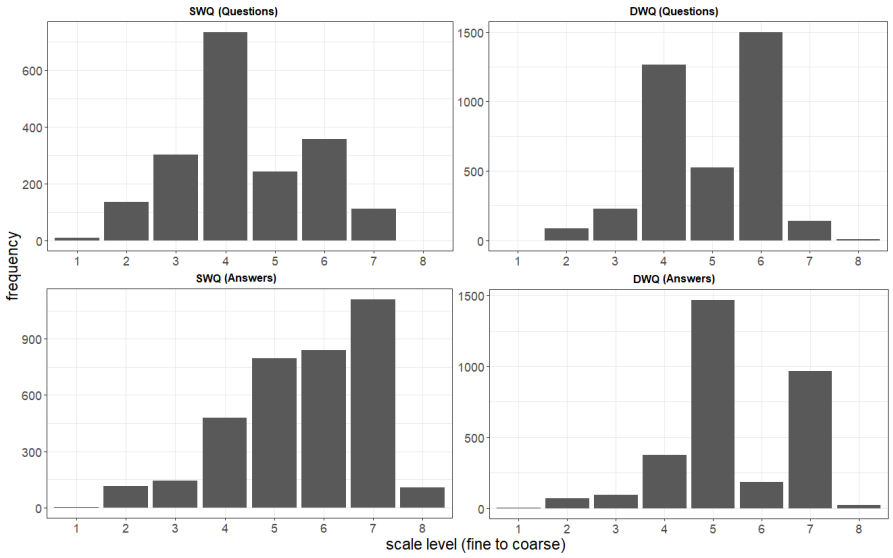

Figure 27: Distribution of the levels of scale in the SWQs versus DWQs.

asked for, and lower than the scale of the places mentioned as additional details. This is because places at a same or coarser level of scale than the places in questions would lead to obvious, and hence irrelevant, answers. The clear distinction of the difference in answering SWQs and DWQs are shown in Table 46, enabling the comparison of the levels of scale of places identified in the questions and their answers.

Figure 32 shows the distribution of prominence in SWQs, DWQs, and their answers. The results show that people ask about less-known places and answer referring to well-known ones. In DWQs, both less-known and well-known places are frequently observed. The well-known places in DWQs are the details presented in the content of the questions. In DWQs, inquirer's state of knowledge can be estimated better, because the question contains what they already know and what they want to find out. Similar to the pattern observed based on scale, prominence levels of places in the answers to DWQs are mostly between the levels of prominence of the toponyms mentioned in their corresponding questions (Table 47).

6.4.3 *Extracted Rules*

To study the structure of human-generated answers and their association with their corresponding where-questions, association rule mining is used to identify strong rules. The association rules extracted from the answers can be used to describe how answers are constructed in detail. The relationship between the content of the questions and their answers can thus also be further investigated using the extracted rules.

Table 30: Comparison of levels of scale in the answers compared to their questions

Level of scale in answers	Lower than	Equal (Between)	Greater than
SWQ			
each value	17.0%	13.1%	**69.9%**
min value	25.8%	16.0%	**58.2%**
median value	14.5%	14.9%	**70.6%**
max value	10.6%	5.9%	**83.5%**
DWQ			
each value	7.2%	**59.6%**	33.2%
min value	10.5%	**77.7%**	11.8%
median value	6.7%	**77.3%**	16.0%
max value	3.9%	40.8%	**55.3%**

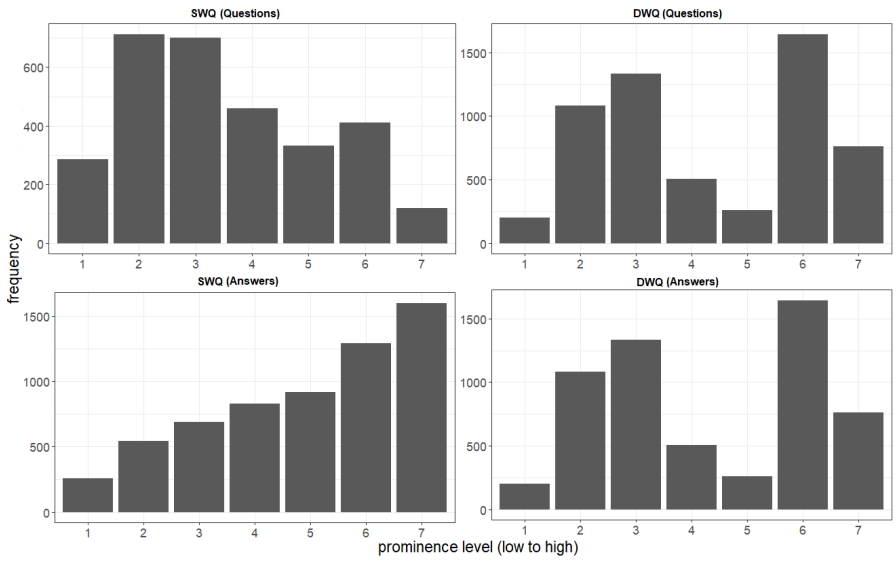

Figure 28: Distribution of prominence levels in the SWQs versus DWQs

Table 31: Comparison of prominence levels in the answers compared to their questions

Prominence level in answers	Lower than	Equal (Between)	Greater than
SWQ			
each value	19.8%	12.5%	**67.7%**
min value	30.7%	14.7%	**54.6%**
median value	19.4%	9.5%	**71.1%**
max value	9.7%	7.9%	**82.4%**
DWQ			
each value	8.7%	**68.4%**	22.9%
min value	13.0%	**77.9%**	9.1%
median value	6.8%	**79.7%**	13.5%
max value	3.5%	**55.6%**	40.9%

Table 32: Extracted rules from type sequences

Rank	rule	support	confidence	lift	frequency
Simple where-questions					
1	{A-ADM2} =>{A-ADM1}	0.15	0.52	1.28	478
2	{Q-ADM1} =>{A-PCLI}	0.08	0.74	1.69	259
3	{A-ADM1,A-ADM2} =>{A-PCLI}	0.08	0.54	1.24	259
4	{Q-ADM2} =>{A-PCLI}	0.06	0.52	1.21	188
5	{Q-PPL,A-ADM2} =>{A-PCLI}	0.04	0.54	1.23	112
Detailed where-questions					
1	{Q-ADM1} =>{A-ADM2}	0.57	0.76	1.12	1701
2	{Q-ADM1} =>{A-PCLI}	0.38	0.50	1.13	1126
3	{A-PCLI} =>{A-ADM2}	0.35	0.79	1.17	1053
4	{A-ADM2,Q-ADM1} =>{A-PCLI}	0.31	0.54	1.21	916
5	{Q-PPL} =>{A-ADM2}	0.22	0.78	1.15	656

Tables 32-34 show the top five extracted rules (based on *frequency/support*) for type, scale and prominence, respectively. In the tables, the values starting with *Q-* relate to the contents of the questions and the values starting with *A-* to the content of the corresponding answers. As shown in the tables, some rules describe the structure of answers (e.g., {A-ADM1, A-ADM2} =>{A-PCLI}) while the others describe the relationships between questions and answers (e.g., {Q-ADM2} =>{A-PCLI}).

Table 32 shows the dominant role of administrative divisions in the human-generated answers. Association rules extracted based on the scale (Table 33) show the *greater-than* and *between* levels of the answers to SWQs and DWQs. The top five patterns of answers are mostly constructed with references to the highest level of prominence (*A-*

Table 33: Extracted rules from scale sequences

Rank	Rule	support	confidence	lift	frequency
Simple where-questions					
1	{Q-6} =>{A-9}	0.21	0.55	1.01	417
2	{Q-6} =>{A-8}	0.20	0.54	1.24	404
3	{A-7} =>{A-9}	0.16	0.56	1.73	307
4	{Q-6} =>{A-7}	0.15	0.54	1.37	295
5	{A-7} =>{A-8}	0.15	0.54	1.25	293
Detailed where-questions					
1	{Q-8} =>{A-7}	0.65	0.80	1.08	1277
2	{Q-6} =>{Q-8}	0.49	0.81	0.99	952
3	{Q-6} =>{A-7}	0.48	0.80	1.07	940
4	{A-9} =>{Q-8}	0.45	0.87	1.07	887
5	{A-7,Q-6} =>{Q-8}	0.42	0.88	1.07	823

Table 34: Extracted rules from prominence sequences

Rank	Rule	support	confidence	lift	frequency
Simple where-questions					
1	{A-4} =>{A-7}	0.14	0.54	1.09	425
2	{A-5} =>{A-7}	0.13	0.50	1.02	417
3	{Q-3} =>{A-7}	0.12	0.52	1.05	382
4	{Q-6} =>{A-7}	0.08	0.58	1.18	260
5	{Q-4} =>{A-7}	0.08	0.53	1.07	250
Detailed where-questions					
1	{Q-6} =>{A-4}	0.32	0.56	1.19	957
2	{Q-6} =>{A-7}	0.30	0.51	1.06	884
3	{A-4} =>{A-7}	0.25	0.54	1.11	742
4	{Q-3} =>{Q-6}	0.24	0.54	0.93	695
5	{Q-3} =>{A-7}	0.22	0.51	1.04	650

7). This shows the major impact of prominence in human answering behavior to where-questions – i.e., people refer to prominent places in answering where-questions.

Tables 32, 33 and 34 show that stronger association rules with higher support are extracted from DWQs in comparison to SWQs. The rules show strong associations between antecedent and consequent parts of the extracted rules with lift value greater than one. The results show that stronger rules with higher confidence and support are extracted using *scale* in comparison to *type* and *prominence*.

The tables only present the extracted rules with highest frequency and support. These tables show how a small set of generic rules describes a large proportion of data in the MS MARCO dataset. Sorting rules by confidence or lift will change the order of the rules. For example, the maximum lift (equal to *8.93*) in the extracted rules belongs to {Q-6, Q-9} =>{A-8} for detailed-where questions using scale. The frequency of this rule is *43*, and it describes the relevant scale level (between minimum and maximum levels of the question) for detailed where-questions. The maximum confidence is 0.93 for detailed where-questions encoded by type. This association rule is {Q-PPLA2, Q-ADM1} =>{A-ADM2} with a frequency *109*. This rule describes that *populated places* in detailed where-questions are mostly localized by referring to *counties* they belong to.

6.4.4 *Predicting the Generic Form of Answers*

The predictability of the generic sequence of an answer given the generic sequence of the corresponding question are assessed using the accuracy measure. Here, different prediction scenarios are considered including (1) the same generic class prediction (e.g., predicting type sequence of answers using type sequence of questions), and (2) prediction of one generic class using all generic classes (e.g., predicting type sequence of answers using type/scale/prominence sequences of questions, see Algorithm 1).

The prediction accuracy is used to evaluate prediction of *content* and *content-and-style* of the answers. Referring to the introductory example, if the type sequence of the answer is predicted as [river, city] then it is captured as a correct prediction for both content and content-and-style. The other permutation of this sequence (i.e., [city, river]) is considered as a correct prediction of content and incorrect prediction for content-and-style. Evidently, any other type sequence is an incorrect prediction for both content and content-and-style scenarios.

Each prediction scenario is applied over all questions, SWQs and DWQs to investigate the impacts of *question types* on the prediction accuracy. Each scenario is tested using all six sequence prediction

methods and is compared with the two baseline approaches (i.e., random generation and most frequent patterns). Only the best prediction performances among the six sequence prediction methods are presented. The *best performance* is the maximum prediction accuracy achieved by one of the methods for a prediction scenario. Prediction performance when multiple predictions are allowed is also evaluated using prediction accuracy – i.e., *top-k predictions* for k from one to five. In top-k predictions, k unique sequences are predicted for each answer and if one of the sequences matches with the generic form of the answer then the prediction is successful.

Table 35 shows the best performances in predicting type sequences of answers. The prediction accuracy based of TSP sequences is noticeably higher than that of predictions using only type sequences. This shows a complementary role of scale and prominence in predicting type sequence of the answers.

Contrasting DWQs and SWQs shows that extra details in DWQs are useful for prediction of the generic form of answers. As shown in the results, the subjectivity in style of answers and flexibility of language to convey information lead to noticeable less accuracy in prediction of content-and-style of answers in comparison to prediction of content. This observation is related to the flexibility of natural language, in which the same meaning can be presented in different ways. Finally, the number of predictions (k in the table) shows that the accuracy dramatically increases in the case of multiple predictions.

Tables 36 and 37 show that compared to type sequence prediction, the TSP sequences contribute less effectively in predicting the prominence and scale sequences – i.e., only slightly improve the prediction accuracy. When considering multiple predictions, TSP sequences lead to worse results than prominence sequences or scale sequences alone. This can be explained by overfitting to specific patterns in the training dataset. Here, overfitting is observed because the schema of types is more than 20 times larger than the scale and prominence schemas. Hence, using type in prediction of scale or prominence leads to very detailed patterns that are not generalizable enough and decrease the prediction accuracy on unseen data. Finally, scale is the most predictable, and prominence is the least predictable generic class. Similar to the observations based on type prediction performances, DWQs are more predictable than SWQs based on scale and prominence.

Table 38 shows the improvement of accuracy in best prediction performances compared to two baselines – i.e., random generator, and most frequent pattern(s)[8]. The minimum improvement is +18.3% in prediction of type sequences of answers using type sequences of questions in comparison to the most frequent pattern(s). This

8 Details are presented in Appendix C, Section 1

Table 35: Prediction accuracy for type sequences

#Predictions (k)	Content		Content and Style	
	Type → Type	TSP → Type	Type → Type	TSP → Type
All questions				
1	45.2	55.7	29.0	40.7
2	68.9	77.1	44.6	60.5
3	80.2	83.3	57.8	73.3
4	83.6	84.7	64.0	76.1
5	84.4	85.5	68.3	77.4
Simple where-questions				
1	39.5	47.5	14.2	27.4
2	60.8	69.4	32.7	48.5
3	73.2	75.8	48.2	63.4
4	77.2	77.5	58.1	66.2
5	78.5	78.2	63.1	67.0
Detailed where-questions				
1	59.1	67.3	47.1	59.6
2	80.4	88.7	61.3	76.3
3	84.0	91.2	65.9	84.4
4	88.0	91.3	73.6	86.4
5	88.5	92.1	75.6	87.1

Table 36: Prediction accuracy for scale sequences

#Predictions (k)	Content		Content and Style	
	Scale → Scale	TSP → Scale	Scale → Scale	TSP → Scale
All questions				
1	55.0	56.7	38.2	42.2
2	79.4	79.2	61.0	62.8
3	91.6	86.1	79.0	76.0
4	96.3	88.7	92.0	81.9
5	98.0	89.3	96.0	83.5
Simple where-questions				
1	48.5	49.5	20.4	28.6
2	79.6	71.8	49.1	49.8
3	89.9	78.3	71.9	67.0
4	95.6	81.8	90.3	74.0
5	97.5	82.6	94.9	75.5
Detailed where-questions				
1	69.6	68.2	59.8	60.6
2	88.4	89.6	78.4	77.3
3	95.8	93.3	88.6	87.1
4	97.5	95.2	94.8	92.1
5	98.6	95.2	97.0	92.7

Table 37: Prediction accuracy of prominence sequences

k	Content		Content and Style	
	Prominence → Prominence	TSP → Prominence	Prominence → Prominence	TSP → Prominence
All questions				
1	50.8	53.0	19.9	30.7
2	74.1	73.4	39.2	49.1
3	85.0	81.9	61.6	66.4
4	92.1	86.7	79.2	77.1
5	96.1	88.6	89.2	81.8
Simple where-questions				
1	45.4	45.6	14.3	19.5
2	75.4	69.4	34.9	39.1
3	84.7	77.0	54.5	56.9
4	91.3	80.5	73.7	68.2
5	95.6	81.9	87.9	72.7
Detailed where-questions				
1	53.3	58.2	26.9	43.8
2	75.1	80.0	50.9	60.4
3	86.0	88.9	70.9	78.4
4	93.1	93.0	82.5	87.0
5	96.8	95.4	91.9	92.6

Table 38: Accuracy improvement using sequence prediction compared to the baselines

Prediction Scenario	Random	Most Frequent Pattern(s)
Type → Type	+48.9%	+18.3%
Scale → Scale	+58.1%	+27.6%
Prominence → Prominence	+39.2%	+30.4%
TSP → Type	+61.6%	+31.0%
TSP → Scale	+54.1%	+23.6%
TSP → Prominence	+42.3%	+33.5%
Overall	+50.7%	+27.4%

observation shows that strong patterns exist in the distributions of answers and consequently, the baseline method performs well in prediction of type sequences of answers. The strongest improvement is +61.6% when comparing the best predictive performance of type sequences using type/scale/prominence sequences together, compared to the random baseline. This is because of the large number of distinct types in type schema that lead to false predictions for the random baseline. The accuracy improvements illustrate the strong relationship between the generic content of questions and generic content of their answers.

To compare the sequence prediction methods, the difference between the prediction accuracy of each method to the best performance achieved by all methods is used to define the best performing method for each prediction scenario. Table 39 shows the root mean square error (RMSE) for each sequence prediction method. The RMSE shows how well-performed a method is in comparison to other methods. If the RMSE of a method is lower than others, the prediction accuracy of the method is higher than the others. The prediction scenarios in Table 39 are simplified groups of actual predictions. For example, prediction scenario of scale is related to predicting scale sequences of answers using (1) scale sequence of questions or (2) type/scale/prominence sequences of questions.

As shown in Table 39, in all scenarios the **CPT** method is the *best* performing method and **TDAG** performs *worst* based on the RMSE values. The results suggest that **CPT** is the best method to construct predictive models to predict the generic form of answers.

6.5 DEMONSTRATION: FROM GENERIC TO SPECIFIC

Translating generic encoding of answer to specific form (e.g., type sequence to toponym sequence) is the last phase in the proposed

Table 39: RMSE of sequence prediction methods

Prediction Scenario	LZ78	Mark1	TDAG	DG	AKOM	CPT	CPT+
Type	7.4%	15.2%	**21.8%**	13.4%	17.3%	**7.1%**	12.9%
Scale	9.8%	12.5%	**17.9%**	10.6%	14.3%	**5.7%**	11.8%
Prominence	8.7%	13.3%	**19.2%**	9.9%	15.2%	**4.9%**	11.5%
Content	8.9%	15.5%	**22.7%**	9.1%	17.4%	**1.9%**	10.2%
Content and Style	8.6%	11.7%	**16.1%**	13.2%	13.7%	**8.2%**	13.8%

approach. The approach to the generic-to-specific translation problem is grounded in the following assumption: *places mentioned in the questions have relationships to places referred to in their answers, and these relations can be found in a knowledge base*. In addition, the specific form of questions and generic form of answers are available through encoding and prediction, respectively. Based on this assumption and the available information, the specific form of answer can be derived using a SPARQL query template (Query 9). While the *structure* of a suitable knowledge base for this purpose has been studied before by Chen et al., 2018, no such knowledge base is yet available with the definitions of type, scale and prominence as used in this study. Hence, the translation is only demonstrated here using the introductory example[9].

DBPedia (Auer et al., 2007) and Geonames are used as sources to demonstrate how SPARQL queries can be used to find specific forms of answers. Considering the information stored in DBPedia and Geonames, this demonstration is limited to type sequences of the answers because the prominence and scale are not available in the place ontology of these knowledge bases. Even the type schema used in DBPedia is different from the Geonames' type schema, and consequently in the following example, mapping to the DBPedia type schema is done manually.

```
PREFIX [KNOWLEDGE BASE]

SELECT distinct ?question ?answer
WHERE {
    VALUES ?question [SPECIFIC] .
    ?answer a [GENERIC] .
    {?question ?r ?answer} UNION {?answer ?r ?question} .
}
```

Query 9: SPARQL template

Referring to the introductory example, the where-question and its answer is modelled as follows:

- specific representation (question): [Putney Bridge];

9 More examples are provided in Appendix C, Section C.4

- TSP encoding (question): type sequence [BDG], scale sequence [4], prominence sequence [3];
- TSP encoding (answer): type sequence [STM, ADM2], scale sequence [6, 6], prominence sequence [6, 7];
- specific representation (answer): [River Thames, London]

The SPARQL queries for finding the specific forms of answers are presented in Queries 16 and 17 using DBPedia and Geonames ontologies. The results of these queries are shown in Table 50. Using DBPedia, the generic forms are correctly translated into River Thames and London. However, the generic to specific translation using Geonames is partially successful. In Geonames, places are conceptualized as points and it supports only containment. This example shows that point-based conceptualization of places is not sufficient for generic to specific translation and more diverse support of spatial relationships can be useful to find the correct specific forms.

```
PREFIX dbo: <http://dbpedia.org/ontology/>

SELECT distinct ?q1 ?a1 ?a2 WHERE {
  VALUES ?q1 {<http://dbpedia.org/resource/Putney_Bridge>}

  ?a1 a dbo:PopulatedPlace .
  {?a1 ?r1 ?q1} UNION {?q1 ?r1 ?a1} .

  ?a2 a dbo:River .
  {?a2 ?r2 ?q1} UNION {?q1 ?r2 ?a2} .
}
```

Query 10: SPARQL query of the example (DBPedia)

```
PREFIX gn: <http://www.geonames.org/ontology#>

SELECT distinct ?q1 ?a1 ?a2 WHERE {
  VALUES ?q1 {<http://sws.geonames.org/6619925/>}

  ?a1 gn:featureCode gn:A.ADM2 .
  {?a1 ?r1 ?q1} UNION {?q1 ?r1 ?a1} .

  ?a2 gn:featureCode gn:H.STM .
  {?a2 ?r2 ?q1} UNION {?q1 ?r2 ?a2} .
}
```

Query 11: SPARQL query of the example (Geonames)

6.6 DISCUSSION

The results of the proposed method show how generic information can be used to characterize and imitate human answering behavior to generate templates for answering the questions. While the results are limited to the human-search engine interaction, the proposed

Table 40: SPARQL results to find specific form of the answer

Knowledge Base	Q1	A1	A2
DBPedia	Putney Bridge	**London**	**River Thames**
Geonames	Putney Bridger	**London**	**–**

methodology (specific-generic translation) is flexibly defined to be applicable to other QA scenarios as well.

Type, scale and prominence are considered as generic classes to investigate the MS MARCO dataset. Their potentials are compared in describing human answering behavior and their performances in predicting the generic forms of the answers. As a result, two major observations are reported.

First, while strong patterns for each generic class have been observed, *scale* is found as the most predictive class. This is because where-questions are a specific subset of spatial questions and scale directly captures inherent spatial aspect of places. Meanwhile, in the notions of type and prominence, other aspects of places contribute as well – e.g., functional and physical aspects. In addition, scale is a generic class that captures hierarchical relationships between places, and previous studies show that these relationships are the basis for answering where-questions (Shanon, 1983). Moreover, type is a better measure than prominence in both characterizing and predicting human-generated answers. This observation is highly influenced by the proxies used to capture type and prominence.

Second, when comparing SWQs and DWQs, the result shows that the generic templates to answer to DWQs, compared to SWQs, can be generated more accurately. Stronger rules and patterns are extracted in the answers to DWQs than in answers to SWQs. This is because DWQs contain richer details which helps narrowing down the list of possible relevant answers. To illustrate this point, two examples are provided (1) *Where in California is Beverly Hills?* and (2) *Where is Beverly Hills?* In the first question, the list of possible relevant answers is narrowed down to *Los Angeles County* because the inquirer already knows it is in *California*. For the latter, respondents are free to subjectively guess the state of the inquirer's geographic knowledge and provide answers such as *Los Angeles County*, *California*, and *United States*.

6.7 CONCLUSIONS

Generating responses with a similar quality to human-generated answers is a challenge to current search engines and QA systems. In particular, where-questions are hard to answer because the

responses can sometimes be either vague or obvious to the inquirers. To avoid generating ambiguous or obvious responses or retrieving unnecessary information as a part of the answers, a proper set of anchor places must be identified to localize the place in question. The assumption that answers to where-questions can be found completely, without any further modification, inside a textual document or as a node or its properties in a knowledge base may not hold in general. Consequently, an approach is introduced to generate templates to answer where-questions based on relevant pieces of information.

The approach is based on the automatic extraction of patterns of generic geographic forms from human-generated QA. These captured in predictive models and are used to generate templates of answers similar human-generated responses. Three generic classes (i.e., type, scale and prominence) are used to investigate the properties of the anchor places in human-generated answers. In this chapter, the questions and answers from MS MARCO v2.1, an extensive dataset constructed from questions submitted to a general-purpose search engine, are studied.

Using distribution analysis and rule mining techniques, recurrent patterns in the questions and their answers are identified. With sequence prediction methods, the generic forms for answers are successfully predicted using the generic forms of the corresponding questions. Moreover, how such generic forms may help with the generation of the appropriate answers are demonstrated with examples, based on the information available in the knowledge bases.

The results show that the prediction of answer structures based on *scale* is more precise, compared to predictions relying on *type* and *prominence*. The rules extracted based on scale have higher support and confidence than the rules extracted from type or prominence. Moreover, the results show how the type of questions (i.e., SWQs vs. DWQs) influence the strength of the extracted rules and lead to noticeable differences in prediction performances. Finally, different sequence prediction methods are assessed and CPT method (Gueniche, Fournier-Viger, and Tseng, 2013) is found as the best performing approach in all scenarios. However, the results of this study are limited to human interaction with a general purpose search engine. Consequently, an important future direction of this research is to investigate other corpora of QA related to different scenarios – e.g., human-human dialogue.

7 DISCUSSION AND CONCLUSIONS

Asking for information about places has been frequently observed and reported in studies of Web search query logs and question answering datasets (e.g., Alonso et al., 2020; Aloteibi and Sanderson, 2014; Sanderson and Kohler, 2004). These studies show that resolving place-related information needs is vital to IR and QA systems. Yet place-related information needs are not separately investigated, but mostly studied only as a part of GIR and GeoQA. GIR and GeoQA also focus on other geographic concepts, such as events and spatial processes. The wider scope of GIR and GeoQA limits the depth of analysis.

The indeterminacy of boundaries (Montello et al., 2003) is not the only challenge in answering place-related questions. Strong context dependency, metaphorical uses, and complex and vernacular references are also major challenges for resolving place-related information needs (e.g., Agnew, 2011; Ballatore, 2016; Hollenstein and Purves, 2010; Winter and Freksa, 2012). In short, the notion of place as experienced or learned by humans is complex to be modelled in IR and QA systems.

Geographic places and their meanings are of the interest of several disciplines including Human Geography, GIScience, Social Sciences, Environmental Psychology, and Architecture (Turner and Turner, 2006). This thesis asserts that the numerous theoretical studies on the notion of place in the aforementioned disciplines provide a rich domain knowledge that can be used to resolve challenges in modelling places and resolving place-related information needs.

The major contributions of this thesis are twofold: First, this thesis contributes theoretically to place-related research by (1) extracting place facets using a systematic and interdisciplinary literature review, and (2) proposing a multi-faceted model of place by synthesising the extracted facets into a unified model (see Chapter 3). The extracted facets and their definitions provide a lexicon for building place-related models for diverse applications, including place-related QA. The multi-faceted model of place also provides a grounding for future studies in place-related research, whenever the properties of places are under investigation.

The second major contribution includes a set of methods and models in place-related QA that utilise domain knowledge (i.e., extracted place facets). In summary, the contribution includes (1) developing a novel encoding approach to study large corpora of place-related questions (Chapter 4), (2) proposing top-down and

bottom-up classifications of place-related questions (Chapter 4), (3) developing a novel method to translate place-related questions to formal queries (Chapter 5), and (4) generic templates to discover patterns of human-generate answers and imitate human answering behaviours to where-questions (Chapter 6). The ultimate goal of these contributions is to show how domain knowledge can be used in several stages of place-related QA including question classifications and analysis, query generation for NL questions, and answer formulations.

In this chapter, I first summarise the results, major findings and limitations of this PhD study, according to the content of Chapters 3, 4, 5 and 6, respectively (Section 7.1). Next, the core contributions of this thesis are summarised and evaluated against the hypotheses in Section 7.2. Finally, I discuss the future directions of place-related research with a focus on QA are discussed in Section 7.3.

7.1 DISCUSSIONS ON MAJOR FINDINGS AND LIMITATIONS

In this section, first the summary, major findings and limitations of place-facets and multi-faceted model of place are discussed. Next, the results of analysing place-related questions from MS MARCO dataset is presented, and the major outcomes and limitations of the question classifications are discussed. Then the findings of translating questions to GeoSPARQL queries and their limitations are provided. Finally, the patterns from human-generated answers and the results of templates to study human answering behaviour are discussed.

7.1.1 *Place Facets*

Place is a central concept in geography and a topic of interest in the social sciences, urban planning, architecture, and most recently in information science. The notion of place has therefore been studied with different foci of interest. Consequently, heterogeneous terminologies, conceptualizations, models, and ontologies have been proposed to capture this elusive concept. Yet these studies complement each other.

In Chapter 3, the notion of place in Geography, Social Science, Architecture and GIScience has been reviewed. Utilizing the concept of place facet as a particular type of information about places, the terminologies, conceptualizations and models of place have been extracted from the literature. Moreover, a unified multi-faceted place model is derived using open card sorting and hierarchical clustering of the extracted facets. The multi-faceted model of place, first distinguishes three classes of facet – i.e., primitive, derived and

linguistic facets. Later the primitive facets are divided into spatial, physical, functional and emotive categories.

The analysis of the categories of *primitive facets* reveals a diversity related to conceptualizations and subsequent formalizations. While in the list of *spatial facets* a balance between conceptualizations and formalizations can be noted, in the *emotive facets*, the dominant theme is conceptualization with no formalism. In other words, the current inadequacy of state-of-the-art computational models of place to handle emotive facets is evident. Consequently, there is currently no noticeable difference in how geospatial databases handle places such as *the Sistine Chapel*, as an example of a sacred and iconic place for a large number of people, with the representation of *a small shopping mall*, which has at most a local influence to a limited number of individuals. While for different groups of individuals these places have different meanings in terms of the emotive, functional, and physical facets, in geospatial databases these places are modelled as geometries with crisp boundaries and a set of attributes.

Based on the systematically selected publications, patterns of research can be observed in place-related research. The patterns show that publications contributing to the anthropocentric facets mostly belong to Human Geography, Social Sciences and Environmental Psychology, while the facets in geographic facets are mostly extracted from publications belonging to GIScience and Architecture.

The multi-faceted model of place has potential limitations. First, the meaning of place as an interdisciplinary concept is reduced into geographic places and consequently non-geographic places (e.g., virtual places) and their facets are not included. Secondly, the review is biased to the literature in Geography, GIScience, Social Sciences and Environmental Psychology. Consequently, place as studied in Anthropology (e.g., Coleman and Collins, 2006), Digital Humanities (e.g., Dunn, 2017), Public Health (e.g., Cummins et al., 2007) or Philosophy (e.g., Casey, 1997) is not covered.

7.1.2 *Question Analysis*

Chapter 4 investigates place questions submitted to search systems and their human-generated answers. Place-based search is motivated by the need to identify places matching some criteria, identify them in space or relative to other places, or to characterize the qualities of such places. The diversity of place questions challenges today's search engines that provide only rudimentary geographic information retrieval support. In Chapter 4, an analysis of the patterns in place-based queries is undertaken using the largest dataset of English language questions/answers currently available,

MS MARCO V2.1. Further, Chapter 4 studies the patterns in the questions, corresponding answers, and their combinations. Clustering approaches and interpretation of the frequent patterns in each cluster are used to categorize the place questions and answers and investigate the relations between these categories. Finally, a theory-driven classification of place-related questions is presented using object-based conceptualization place and the multi-faceted model of place, and it is compared to the data-driven question classification. The theory-driven classification covers the categories of place facets in the multi-faceted model of place, including spatial, physical, functional, emotive and linguistic facets. The result is a detailed classification of place-related question which includes 14 classes.

The encoding approach is designed to be performed as an unsupervised technique to characterise large datasets. The unsupervised design enhances the usability and applicability of the method without any time consuming data annotation. Yet, the results are human interpretable and the encoding results can be used to characterise large corpora of place-related questions with minimal supervisions to interpret the results.

The encoding results reveal a high potential of the machine comprehension datasets such as MS MARCO for future place-related studies. It is, however, highly recommended to consider appropriate sampling strategies in order to deal with the geographical and semantic biases observed and discussed throughout Chapter 4. Using the proposed encoding and georeferencing technique, the geographic bias of the MS MARCO dataset has been identified. Most of the questions in this dataset are about places in the North America, Western Europe and Australia, and only a small portion of data are the questions of places in Africa, Asia, South America and Eastern Europe. Moreover, the bias to specific place types such as countries, states and counties is evident based on the encoding results.

The encoding method has two major limitations: First, when the questions are complex and contain multiple criteria, the encodings become less expressive and hard to be interpreted. Second, some specific question types such as comparison questions are out of the scope of the encoding method (this limitation later resolved in the extended encoding in Chapter 5). Semantically richer encoding representations that can differentiate more complex types of questions and answers may lead to a more nuanced characterization of the QA dataset.

In Chapter 4, the data-driven and theory-driven classifications are compared. The data-driven classification is highly sensitive to biases in the data collection and thus only reflect question types represented in, or supported by the MS MARCO dataset. As a result,

questions about spatial facets are overrepresented in the data-driven classification, while emotive, functional and physical facets are marginalized. In contrast, the theory-driven classification is a generic, cross-disciplinary approach based on the spatial core concepts (Kuhn, 2012), and with less expert bias using the multi-faceted model of place.

7.1.3 *From Question to Queries*

Translating NL questions to formal queries (e.g., SQL) is an essential task of database-oriented QA. To use structured data in databases or knowledge bases in QA systems, formal queries must be generated automatically from NL questions. Using structured data in place-related QA is especially useful whenever spatial analysis is needed to answer the questions. Hence, such translation is a vital step toward answering place-related questions in practice.

In Chapter 5, a novel approach using domain knowledge and object-based conceptualization of place is designed to translate NL questions to GeoSPARQL queries. This approach includes five major steps, (1) information extraction and encoding, (2) parsing and relation identification, (3) concept identification, (4) generating FOL statements and (5) translating FOL statements to GeoSPARQL queries. The novelty of the approach is mainly to (1) use domain knowledge and avoid reinventing new terms, and (2) utilise FOL statements as an intermediate representation which can be later translated not only to GeoSPARQL, but any other formal query language with minimal effort. The method is tested using the Geospatial Gold Standard dataset (Punjani et al., 2018).

The method utilizes state-of-the-art NLP models for part-of-speech tagging, named entity recognition, constituency parsing and dependency parsing. Hence, the method is able to translate NL questions to GeoSPARQL queries without time-consuming data annotation. The detailed evaluation of the method shows that it can overcome existing state-of-the-art approaches in the translation task. Moreover, the results shows that integrating phrase-level information (from constituency parsing) and the extracted dependencies (using dependency parsing) improves the information extraction and relation identification significantly.

Using the object-based conceptualisation of place as the grounding theory, the method is less sensitive to a biased dataset. Hence, the method is able to translate diverse types of place-related questions. The results also show that the dataset does not contain enough diversity to test the performance of the whole capability of the proposed method. The results of this study shows that the Geospatial Gold Standard dataset is biased to specific question types and does not cover important types of place-related questions such

as activity-based or event-based questions. Hence, the need for a larger dataset that is less biased to spatial relationships and cover sophisticated place-related questions is evident based on the results of Chapter 5.

The limitations of the proposed translation can be summarised as follows: First, the dataset is small and more reliable testing is needed to evaluate the performance of the unsupervised method. Next, some steps are not considered in the translation task including ontology mapping and fetching the results from knowledge bases. While the queries are syntactically and semantically evaluated, the ontology mapping may introduce new errors in the process. In ontology mapping, the identified concepts should match to actual design of the knowledge bases and the information stored in it. Fetching the results is also an error prone task due to incompleteness of data in current knowledge bases that is not evaluated in Chapter 5. Yet these two challenges should be resolved before applying the method to practical usage.

7.1.4 *Answering Questions*

Where-questions are answered by place descriptions in everyday communication. To answer where-questions automatically, computers should be able to generate relevant place descriptions that satisfy inquirers' information needs. Human-generated answers to where-questions are constructed based on a few anchor places that characterize the location of inquired places. The challenge for automatically generating such relevant responses stems from *selecting relevant anchor places*. In Chapter 6, *templates* are introduced that allow to characterize the human-generated answers and to imitate their structure.

The templates are patterns of generic geographic information derived and encoded from MS MARCO dataset. First, the toponyms are identified and encoded into sequences of generic information. Next, sequence prediction methods are used to model the relation between the generic information in the questions and their answers. The results show strong patterns in the questions and answers, and associations among generic information that describe the relevance of answers to the corresponding questions.

To study human-generated answers to where-questions, three generic classes are selected – i.e., *scale*, *type* and *prominence*. The results of analysing templates show that *scale* is most effective generic class to study and imitate human answering behaviour due to stronger associations and more accurate results. Moreover, the patterns show that human-generated answers to different type of where-questions (i.e., simple where-questions vs. detailed where-questions) are semantically different. For example, in simple

where-questions the scale level of places in answers are often *coarser* in comparison to the scale level of the place in question. However, in detailed where-questions, the scale level of places in answers are often *in between* the scale levels of of places in the questions. Finally, patterns of templates using place types show that meaningful insights can be extracted from the dataset to describe why the anchor places are relevant to localise the place in question – e.g., bridges are often localised by referring to rivers, and cities are often localised by states.

The limitations of templates and the approach to imitate the structure of human-generated answers in machine responses are twofold:

Theoretical limitations:

1. The generic-to-specific translation approach is only focused on where-questions, and other types of geographic questions are neglected;
2. The proposed approach is focused only on the questions and their relationship with the answers, when no other contextual information about inquirers is available;
3. The approach is designed with an exclusive focus on toponyms while qualitative spatial relationships have an important role in answering where questions;
4. The additional impacts of qualitative spatial relationships (e.g., *in southern part of*) as modifiers of scale are neglected in the TSP encoding.

Results limitations:

1. The biases of the English language MS MARCO dataset directly influence the results. The data are extracted from the Microsoft Bing search engine, and hence the results are necessarily biased to the questions asked by users of this search engine. In addition, the sampling approach used when extracting MS MARCO questions from the Microsoft Bing query logs may have a direct and unquantifiable impact on the generality of the results;
2. The results are influenced by the geographic biases and incompleteness of data in Geonames and OSM Nominatim. The bias and incompleteness of gazetteers are well-documented by Acheson, Sabbata, and Purves (2017);
3. The biases in the proxies that have been used to capture the TSP encoding also have an impact on the results.

Despite these limitations, the identified patterns of Chapter 6 align well with everyday experience and provide a grounding for answering where-questions.

7.2 CONTRIBUTIONS OF THIS THESIS

This PhD thesis has addressed the identified major challenges of place modelling, question analysis, translation of questions to queries, and answer generation in the scope of place-related QA. The major contributions of this PhD study are summarized and listed below:

1. Designing a systematic approach to extract place facets from interdisciplinary review of the literature.
2. Proposing a multi-faceted model of place that provides a unified view to the notion of place based on the interdisciplinary literature review.
3. Developing an encoding approach to study large corpora of place-related questions and their answers.
4. Proposing data-driven and theory-driven classifications of place-related questions using the largest available data (MS MARCO) and object-based conceptualization of place, respectively.
5. Proposing a novel method to translate NL questions to formal queries based on object-based conceptualization of place that improves the existing state-of-the-art approaches.
6. Developing a generalized approach to investigate human answering behavior to where-questions using generic geographic information;
7. Analysing the human-generated answers to where-questions asked in Web search, using patterns of type, scale and prominence of places.

The contributions are evaluated against the four hypotheses of this study (Chapter 1). In the following, first the hypotheses are reviewed and then further discussions are provided.

- **Hypothesis 1**: A multi-faceted model of place can be derived by categorisation of extracted place properties from a systematic literature review.
- **Hypothesis 2**: Place-related questions can be classified using structural patterns derived through NLP and information extraction in a bottom-up approach, or can be derived from a theoretical conceptualisation of place in information system in a top-down approach.

- **Hypothesis 3**: Formal query representations of place-related questions can be derived from a pipeline process of information extraction, textual parsing and template-based query generation.
- **Hypothesis 4**: Generic information can be used to characterize human answering behavior and ultimately to generate templates for answering where-questions.

The first hypothesis is addressed in Chapter 3. The unified multi-faceted model of place is derived using open card sorting and hierarchical clustering. The open card sorting experiment allows participants to propose different categorizations based on their opinions. Yet the results of individual categorizations show a *high agreement between the expert opinions*. The applicability of the model and the extracted facets are demonstrated in Chapters 4, 5 and 6 to classify place-related questions, extract useful information from large QA corpora, find patterns that describe human answering behaviour, and also to study the structure and relevance of answers to where-questions. Therefore, the results show the multifaceted model of place can be derived from individual categorisation with high agreement among the experts, and the model is supported to be useful as a grounding to involve domain knowledge in place-related question answering.

For the second hypothesis, the largest available QA data (Nguyen et al., 2016) and state-of-the-art conceptualisation of place (Purves, Winter, and Kuhn, 2019) are used to derive data-driven and theory-driven classifications of place-related questions. To classify the questions of MS MARCO dataset, a novel encoding approach is introduced in Chapter 4. The results show that even using the largest QA dataset, data-driven classifications are biased to specific question types such as *spatial* questions. Yet, the theory-driven approach is more detailed and covers the identified categories of facets including *spatial*, *physical*, *functional* and *emotive* facets. The theory-driven classification also compared with current classifications. The results show that the previous classifications are a subset of the theory-driven classification, and also the proposed classification includes more diverse types of questions. Consequently, the second hypothesis is supported, and the results show the superiority of theory-driven classification over the data-driven classification.

The third hypothesis is investigated and addressed in Chapter 5 by introducing a novel approach to translate NL questions to GeoSPARQL queries. The results show that using both dependency and constituency parsing improve the relation identification in comparison to the state-of-the-art method (Punjani et al., 2018). Using FOL statements, several parts of the method can be used for translating questions to other formal query languages such as

SPATIAL SQL. Due to theoretical grounding in the object-based conceptualization of place, the method also can translate the missing question types in the Geospatial Gold Standard dataset. Therefore, the third hypothesis is supported, and the results show the impact of domain knowledge to improve the existing translation methods.

For the fourth hypothesis, a new method using templates of generic information is designed to study the association between where-questions and their human-generated answers. Next, the method utilizes sequence prediction approaches to imitate human answering behaviour automatically. The method is proven to be useful in extracting patterns from questions and answers, and finding insightful associations between places mentioned in the questions and answers. Moreover, the properties of relevant anchor places to answer where-questions are captured in a predictive model. This predictive model can be used in answering where-questions or testing the relevance and quality of the responses generated by other approaches. The results show that generic information can describe the relevance of anchor places to the place in question, and predictive models can be constructed to capture this knowledge in a machine digestible manner. Therefore, the applicability and usability of domain knowledge in analysing and answering where-question is supported.

7.3 FUTURE DIRECTIONS

In this section, I conclude the thesis by listing suggestions for future works, identified during my study. The five future directions presented below can be used in QA systems to resolve place-related information needs.

7.3.1 *A Computational Model of Place*

The current focus of place formalisms on pragmatic applications, such as finding similar places and differentiating contrasting places (from a small set of well defined, utilitarian types of places) is evident based on the results of Chapter 3. The spatial facets of place are dominant. Consequently, further research is required to investigate place similarity measures which take the emotive, functional and physical facets of place into account. This direction is of increasing importance as people's interaction with each other is increasingly mediated by computers – and often the natural language interaction with intelligent spatial assistant systems may be the only interaction desired. In short, the emotive facets are much under-investigated in terms of formal models compared to the facets in the functional group of anthropocentric facets. Moreover, the facets in the spatial group are mostly formalized into computational

models while the facets in the physical group are less developed in terms of formalization.

A computational model of place that includes the four categories of primitive facets and utilise them in finding the similarity and differences of geographic places can be used in diverse applications including place-related QA, GIR and place recommendation systems. Yet to develop such computational models, first the emotive and physical facets should be formalised. In addition, further research is needed to derive useful similarity measures that can combine different facets into consideration.

7.3.2 *Questions to SPATIAL SQL: A Step Towards Natural Language Interfaces to GIS*

Spatial databases can perform a wide range of spatial analysis on structured data. Especially, the relational databases support most of the GIS functionalities and analysis toolboxes. Moreover, the best practice of storing GIS data in databases for better data management, increases the importance of relational databases in both research and practice. Yet to enable the user-centered design and reduce technical barriers between users and GIS, a natural language interface is required.

The major challenges in translating NL questions to SPATIAL SQL compared to the translation achieved in Chapter 5 stem from the difference between SPATIAL SQL and GeoSPARQL capabilities. While GeoSPARQL only provides basic functionalities such as metric and topological relationships, SPATIAL SQL has been developed into a full-fledged query language that supports most of the functionalities afforded by current GIS systems. Hence, the questions to be answered by databases may include more diverse types of spatial analysis and such difference needs further research to be addressed. Similar to challenges in translating questions to GeoSPARQL queries, ontology mapping is vital to translate questions to SPATIAL SQL. Especially, ontology mapping is needed to semantically match extracted information from questions to find tables and columns relevant for answering the questions. The results of enabling natural language interface to spatial databases not only improve the user experience of current GIS systems but also has an impact on future developments of geoportals, and IR and QA systems.

7.3.3 *Generic Geographic Information to Specific References*

In a geographic document with a relatively coherent context (e.g., a question-answer pair about place, a paragraph describing a place in a Web document), several references to geographic places can be

found. These references can be presented in specific form (e.g., Melbourne) or generic form (e.g., the beautiful city). While extracting specific references, toponyms, have been studied with a focus on identifying and resolving their ambiguity (e.g., Leidner, 2007; Santos et al., 2018); a solution to extract the generic information and to resolve their ambiguity is still missing.

Developing an automatic approach to decode generic references to specific representations (i.e., toponyms) is a necessary step to complete the process of the specific-generic translation approach (see Chapter 6). Available information in documents or knowledge bases can be used to derive the specific representations from the identified generic references. The results of generic-to-specific study can be used in rewriting generic references inside the documents and place-related questions to their specific representations. After rewriting generic references, the text will be ready for geoparsing, and more information can be extracted using the existing geoparsing techniques. Consequently, the results will influence a wide range of applications from GIR to place-related QA.

7.3.4 *Benchmark Dataset for Place-Related Question Answering*

While several machine comprehension datasets are available now, a dataset that fits for the community of geographic information science is still missing. First, the existing datasets – at best – provide a 'location' label which is not specifically defined for the purpose of our community. Second, the essential information which is needed for future research in the geographic information community is not provided – e.g, georeferenced information. Consequently, every research goes through reinventing an approach for these tasks or at least putting time and efforts to extract this information which may not directly relate to their contributions. Finally, the quality of these extractions differs in different studies and as a result, methods cannot be compared reasonably. An enriched and integrated dataset of geographic question can be used for future geographic and place-related QA research and enable the further developments of geographic and place-related QA systems.

7.3.5 *Place Recommendation: Moving from Content to Context*

This PhD study is mainly focused on the content of place-related questions and their answers. The role of context in place-related QA is also important in place-related QA and the lack of support for contextual information is a major limitation of this thesis. The study of inquirer's information and preferences is necessary for answering questions in practice. A model to capture user preferences is vital for resolving several place-related information needs, specifically in

place recommendation systems. Such systems recommend places to visit or places that afford specific activities to tourists or local people.

To present useful recommendations, contextual information should be captured alongside with information about places and their properties. The multi-faceted model of place (Chapter 3) can be used to define a model for the latter, yet a model that goes beyond location information of inquirers is still missing. Consequently, defining such a model for user preferences in place recommendation systems, and integrating contextual information with the place model are suggested as a promising research direction to resolve place-related information needs.

Part I

APPENDIX

APPENDIX A

A.1 EXTRACTED FACETS

The extracted facets from the literature are presented in Table 41.

Table 41: Extracted facets from the literature

Facet name	References
Accessibility	(Winter and Freksa, 2012; Ingram, 1971)
Action	(Das and Winter, 2016; Entrikin, 1997; Hockenberry, 2006; Jonietz, 2016; Jordan et al., 1998; Kuhn, 2001)
Activity	(Almuzaini, 2017; Calafiore, 2016; Couclelis, 1992; Das and Winter, 2016; Edwardes, 2007; Edwardes and Purves, 2007b; ElGindy and Abdelmoty, 2014b; ElGindy and Abdelmoty, 2014a; Jonietz, 2016; Kuhn, 2001; Relph, 1976; Tuan, 1977; Tversky and Hemenway, 1983)
Address	(Almuzaini, 2017)
Affordance	(Alazzawi et al., 2012; Almuzaini, 2017; Ballatore, 2016; Das and Winter, 2016; Edwardes, 2007; ElGindy and Abdelmoty, 2014b; ElGindy and Abdelmoty, 2014a; Galton, 2010; Gibson, 1979; Harrison and Dourish, 1996; Hockenberry, 2006; Jonietz and Timpf, 2015; Jordan et al., 1998; Kuhn, 2001; Ortmann and Kuhn, 2010; Sabbata et al., 2015; Scheider and Janowicz, 2014; Scheider and Krzysztof, 2010; Scheider and Purves, 2013; Winter and Freksa, 2012)
Agentive Dimension	(Couclelis, 2010)
Alternating Repetition	(Vasardani et al., 2016; Vasardani and Winter, 2016)
Aspect of Design	(Canter, 1997)
Attributes (Features)	(Tversky and Hemenway, 1983)
Behaviour	(Buttimer, 1976; Edwardes, 2007; Jonietz, 2016)
Boundaries	(Couclelis, 1992; Montello et al., 2003; Vasardani et al., 2016; Vasardani and Winter, 2016; Winter and Freksa, 2012)

Classification	(Scheider and Janowicz, 2014; Scheider and Purves, 2013; Tversky and Hemenway, 1983)
Composition	(Papadakis et al., 2016)
Constitutive Dimension	(Couclelis, 2010)
Containment	(Scheider and Janowicz, 2014; Scheider and Krzysztof, 2010; Scheider and Purves, 2013; Winter and Freksa, 2012)
Contrast	(Vasardani et al., 2016; Vasardani and Winter, 2016)
Deep Interlock and Ambiguity	(Vasardani et al., 2016; Vasardani and Winter, 2016)
Echoes	(Vasardani et al., 2016; Vasardani and Winter, 2016)
Elements	(Edwardes and Purves, 2007a; Edwardes and Purves, 2007b)
Emotional Attachments	(Mennis and Mason, 2016)
Environment	(Gustafson, 2001)
Environment-Self	(Gustafson, 2001)
Equipment	(Scheider and Janowicz, 2014; Scheider and Purves, 2013)
Event	(Roche, 2015)
Feasibility (of an activity)	(Jonietz, 2016)
Footprint	(Derungs and Purves, 2014; Goodchild, 2011)
Form	(Canter, 1997)
Formal Dimension	(Couclelis, 2010)
Function	(Canter, 1977; Papadakis et al., 2016)
Functional Differentiation	(Canter, 1997)
Functional Properties	(Vasardani et al., 2016)
Generalized Properties (Place Properties)	(Adams, 2015)
Geographic Location	(Adams, 2015)
Good shape	(Vasardani et al., 2016; Vasardani and Winter, 2016)

Gradients	(Vasardani et al., 2016; Vasardani and Winter, 2016)
Identification	(Scheider and Janowicz, 2014; Scheider and Purves, 2013)
Investment with Meaning and Value	(Gieryn, 2000)
Level of Scale	(Richter et al., 2013a; Richter et al., 2013b; Vasardani et al., 2013; Vasardani et al., 2016)
Local Symmetries	(Vasardani et al., 2016; Vasardani and Winter, 2016)
Locale	(Agnew, 2011; Capineri, 2016; Cresswell, 2004; Hollenstein and Purves, 2010; Jonietz, 2016; Massey, 1994; Purves and Derungs, 2015)
Localization	(Scheider and Janowicz, 2014; Scheider and Purves, 2013)
Location	All
Material Form (Physicality)	(Gieryn, 2000)
Meaning	(Buttimer, 1976; Canter, 1977; Canter, 1997; Couclelis, 1992; Edwardes, 2007; Entrikin, 1991; Entrikin, 1994; Entrikin, 1996; Harrison and Dourish, 1996; Jonietz, 2016; Massey, 1994; Relph, 1976; Tuan, 1977; Turner and Turner, 2006; Vanclay, 2008; Capineri, 2016)
Narrative Descriptions	(Jordan et al., 1998)
Not Separateness	(Vasardani et al., 2016; Vasardani and Winter, 2016)
Object of Action	(Hockenberry, 2006)
Observed Properties (Environment Properties)	(Adams, 2015)
Others	(Gustafson, 2001)
Others-Environment	(Gustafson, 2001)
Parts	(Tversky and Hemenway, 1983)
Physical Dimension	(Ballatore, 2016)
Physical Elements	(Scannell and Gifford, 2010)
Physical Features	(Jordan et al., 1998; Turner and Turner, 2006)
Physical Landscape	(Cresswell, 2004)
Physical Setting	(Relph, 1976)
Place Affect	(Almuzaini, 2017; Cresswell, 2004)

Place Attachment (1)	(Scannell and Gifford, 2010; Tuan, 1977)
Place Attachment (2)	(Vanclay, 2008)
Place Commitment	(Vanclay, 2008)
Place Dependence (1)	(Almuzaini, 2017)
Place Dependence (2)	(Vanclay, 2008)
Place Familiarity	(Vanclay, 2008)
Place Identity (1)	(Canter, 1977; Cresswell, 2004; Entrikin, 1991; Entrikin, 1996; Entrikin, 1997; Gustafson, 2001; Jonietz, 2016; Massey, 1994; Relph, 1976; Scannell and Gifford, 2010; Tuan, 1977; Turner and Turner, 2006)
Place Identity (2)	(Vanclay, 2008)
Place Name (Toponym)	All
Place Objectives	(Canter, 1997)
Place Reference (Verbal Reference)	(Alazzawi et al., 2012; Almuzaini, 2017; Bennett and Agarwal, 2007; Montello et al., 2003; Purves and Derungs, 2015; Winter and Freksa, 2012)
Place Social Bonding	(Almuzaini, 2017)
Positive Space	(Vasardani et al., 2016; Vasardani and Winter, 2016)
Purpose	(Papadakis et al., 2016)
Qualities (1)	(Edwardes and Purves, 2007a; Edwardes and Purves, 2007b; Purves et al., 2011)
Qualities (2)	(Ortmann and Kuhn, 2010)
Quality of Place	(Jonietz, 2016)
Roughness	(Vasardani et al., 2016; Vasardani and Winter, 2016)
Salience of Place (Absolute view)	(Almuzaini, 2017)
Salience of Place (Personal view)	(Almuzaini, 2017)
Saliency	(Almuzaini, 2017; Richter et al., 2013a; Richter et al., 2013b; Winter and Truelove, 2013; Winter and Freksa, 2012)
Scale of Interaction	(Canter, 1997; Goldberg et al., 2009; Purves and Derungs, 2015; Richter et al., 2013a; Richter et al., 2013b)
Self	(Gustafson, 2001)

Self-Others	(Gustafson, 2001)
Self-Others-Environment	(Gustafson, 2001)
Semantic Relation to Activities	(Scheider and Purves, 2013)
Semantic Relation to Objects	(Scheider and Purves, 2013)
Semantic Relation to Qualities	(Scheider and Purves, 2013)
Sense of Place (1)	(Buttimer, 1976; Canter, 1977; Capineri, 2016; Cresswell, 2004; Edwardes, 2007; Entrikin, 1991; Harrison and Dourish, 1996; Massey, 1994; Purves and Derungs, 2015; Purves et al., 2008; Tuan, 1977)
Sense of Place (2)	(Vanclay, 2008)
Sentiment (Individual Reflection)	(ElGindy and Abdelmoty, 2014b; ElGindy and Abdelmoty, 2014a)
Service	(Alazzawi et al., 2012; Winter and Truelove, 2013)
Simplicity and Inner Calm	(Vasardani et al., 2016; Vasardani and Winter, 2016)
Social Elements	(Scannell and Gifford, 2010)
Social Role (Social Dimension)	(Ballatore, 2016)
Social Role Characteristics	(Calafiore, 2016)
Socioeconomic and Cultural Factors	(Jordan et al., 1998)
Space	(Canter, 1997)
Spatial Accessibility	(Jonietz, 2016; Winter and Freksa, 2012)
Spatial Band	(Adams, 2015; Janowicz, 2012; Janowicz et al., 2013; McKenzie and Adams, 2017)
Spatial Identity	(Quesnot and Roche, 2015; Roche, 2015)
Spatial Level	(Scannell and Gifford, 2010)
Spatial Properties	(Papadakis et al., 2016)
Spatial Reference	(Adams, 2015; Alazzawi et al., 2012; Almuzaini, 2017; Bennett and Agarwal, 2007; Jones et al., 2001; Kim et al., 2016; Montello et al., 2003; Richter et al., 2013b; Vasardani et al., 2013; Winter, 2009; Winter and Freksa, 2012)

Spatial Relation to Other Places	(Scheider and Purves, 2013)
Spatial Value	(Quesnot and Roche, 2015)
Specificity	(Scannell and Gifford, 2010)
Spirit of Place	(Vanclay, 2008)
Strong Centres	(Vasardani et al., 2016; Vasardani and Winter, 2016)
Structural Properties	(Vasardani et al., 2016; Vasardani and Winter, 2016)
Suitability (of an activity)	(Jonietz, 2016; Jonietz and Timpf, 2015)
Symbolic Representation	(Jordan et al., 1998)
Telic Dimension	(Couclelis, 2010)
Temporal Band	(Adams and Janowicz, 2015; Janowicz, 2012; Janowicz et al., 2013)
The Void	(Vasardani et al., 2016; Vasardani and Winter, 2016)
Thematic Band	(Janowicz, 2012; Janowicz et al., 2013; McKenzie and Adams, 2017)
Type	All
Typologies	(Jordan et al., 1998; Tversky and Hemenway, 1983)
Valuation (of an action)	(Hockenberry, 2006)

A.2 FACET DEFINITIONS

Table 42 shows the definitions of the facets which are extracted from the literature.

Table 42: Extracted facets from the literature with their definitions

Facet name	Definition

Accessibility	"In one sense the word 'accessibility' means capable of being reached, thus, implying a measure of the proximity between two points. Alternatively, 'accessibility' is related to the ability of a transportation system to provide a low cost and/or quick method of overcoming the distance between different locations" (Ingram, 1971). Accessibility is a term related to the spatial configuration of places with a focus on connectivity of places, whether places are connected or disconnected and an agent can move from one place to another. Accessibility is a contextual property of places – i.e., in some context, a place can be accessible from another place, while in other contexts it may not be accessible (Winter and Freksa, 2012). Accessibility is highly related to the mode of transportation, and the contextual information about the agents.
Action	Action is defined as a type of activity that is intentional and conscious and oriented towards the fulfillment of a particular goal. In Jordan's place model, it is defined as: "People perform actions in places. As we have seen, actions are one of the most important aspects that give meaning to a place. Personal actions can create personal histories. Personal actions allow users to participate in accepted patterns of behavior, as well as to personalize the place by performing individualizing actions. For example: Whenever I'm at this mall, I go to store X because I always found nice and cheap clothes there" (Jordan et al., 1998).
Activity	Activity is defined as a behavior or action from an agent in a specific place. As a part of a model (attributes, activities, and parts) for describing environmental scenes (Tversky and Hemenway, 1983), activities are defined as what people can do in the scenes such as 'shopping' for stores.
Address	Address is a reference to a place which describes where the place is in (usually hierarchically organized based on the part-of relationship).

Affordance	Affordance is described as what an environment (e.g., a place) can afford to an agent (e.g, human, animal). By capturing affordance in a place, one can consider the following questions to be answered: What can be done (even potentially) in the place? What can a place afford to a human?
Agentive Dimension	Agentive dimension is one of the four dimensions of objects of discourse (Couclelis, 2010). The four dimensions can be considered as dimensions of place as an object of discourse. It is defined as: "The processes by which things come to be or their roles as agents in other processes or their function relative to some end" (Couclelis, 2010). For example, stopping illegal immigration is the agentive dimension of Gibraltar (Couclelis, 2010).
Alternating Repetition	"Alternating repetitions are common in places with self-similarity patterns such as coastal places or forests." This property is one of the six properties (good shape, local symmetries, echoes, roughness, alternating repetitions, and positive space) about the internal structural organization of places which makes places distinctive. (Vasardani and Winter, 2016)
Aspect of Design	Aspect of design is a facet of place about physical characteristics of a place. It is defined as a fact that "gives rise to a whole basket of hypotheses about how the structure of place experience may take on different forms in relation to the aspects of design that are being considered. For example personal comfort at the immediate scale, say in a room, may be most clearly revealed by examining the functions of a particular space and how readily these are achieved. This is probably what most building evaluations seek to do. To take another example, the form (or style) of a place may be examined to see whether it reveals emphases at the personal, social or cultural level of differentiation" (Canter, 1997).

Attributes (Features)	It is a part of a place model (attributes, activities, and parts) for describing environmental scenes (Tversky and Hemenway, 1983). Attributes are described as features of the scenes. For example, 'high' can be considered as one of the attributes describing mountains.
Boundaries	Boundary defines which part of space belong a place. Places have boundaries, either crisp or fuzzy.
Classification	In place reference system (Scheider and Janowicz, 2014), classification is described as kind of places and wants to answer: "what kind of place it is".
Composition	"The dimension of composition illustrates the constitution level, representing the spatial organization of a place as a composite object formed by simple interrelated components" (Papadakis et al., 2016). Composition is related to spatial objects with specific spatial associations which enables some functions for that place and build specific spatial properties to the place.
Constitutive Dimension	Constitutive dimension is one of the four dimensions of objects of discourse. The four dimensions can be considered as dimensions of place as an object of discourse. It is defined as: "The constitutive dimension has to do with what the object is made of and in particular, its parts (material or abstract), and how these are connected" (Couclelis, 2010). For example, rock is related to the constitutive dimension of mountains.
Containment	In place reference system, things (places and other objects) that exist in a place are described as place containment. They defined this element as "An essential function of places is that they serve to localize other things" (Scheider and Janowicz, 2014).

Contrast	Contrast is a way that people differentiate places from each other and at the same time is a way to capture the notion of place (Winter and Freksa,2012). Contrast is one of the properties mentioned and defined by Vasardani and Winter (2016): "as a whole then, the entities that form a place and places themselves, exhibit centres, that can be contrasted to each other, may have distinct boundaries or gradients that create harmonious transitions between contrasting centres and they all belong to the same whole structure..."
Deep Interlock and Ambiguity	Deep interlock is related to ambiguity between boundaries of elements (parts) that constitute the whole (place). In case of deep interlock and ambiguity, the navigation from one part of a place to another part is not an abrupt transition and it is not clear that exactly when and where the person moves from one part to its adjacent part.
Echoes	"Echo is the property of a place when its whole structure is reminiscent of another place, or the sense of familiarity that people some times experience in a place, even when they visit it for the first time, due to its similarity with other places of similar properties. Echoes can be internal, as similarity of the elements within a place, or external, characterizing a place as a whole when compared with other places" (Vasardani and Winter, 2016).
Elements	Places can be described by the elements they are belonging to (e.g., a city has elements such as church and school). People use these elements as descriptive terms to describe a place using natural languages. Elements can be extracted from place descriptions. Elements are part of a place model which includes three parts: 1- elements, 2- qualities, and 3- activities (Edwardes and Purves, 2007a; Edwardes and Purves, 2007b).

Emotional Attachments	"Place can be understood, and represented, not only as an attribute of a location but also as the emotional attachments that characterize a relationship between an individual and a location". Emotional attachment can be modeled as a combination of mood state, behaviors, and social interactions (Mennis and Mason, 2016).
Environment	Environment is a part of a three-pole model which is presented by Gustafson (2001). Environment related to physical characteristics and natural conditions. "Very often, meanings of place depend neither on the self, nor on the relations with or perceptions of others. I have tentatively labelled this pole 'environment'". This pole includes "the physical environment, including the natural environment and various natural conditions (weather, seasons), as well as the built environment" (Gustafson, 2001).
Environment-Self	Environment-Self is part of Gustafson (2001) model which is defined in the relation between a person and a place. "Meanings of place may also concern the relationship between self and environment. This relationship is often based on the respondents' knowledge of the place. Some refer to a formal knowledge (geographical, historical), others value their familiarity with their lived-in physical environment" (Gustafson, 2001).
Equipment	Equipment described as objects associated with a place. This object can be inside or outside of the place, but they play a role in the activities related to the place. For example, a place to see a specific statue, the statue may not be in the place, but the act of seeing is doable in that place. This part of the model answers the following question: What is a place equipped with?
Event	"Roche (2012) formalizes this conception of place with this function: P = f (N, E, L), where P is the place, N the Name, E the Event, and L the Location." "Event refers to a large spectrum of meanings, the space within which humans carry out habitual aspects of their lives, such as shopping, work, recreation, and sleeping" (Roche, 2015).

Feasibility (of an activity)	Feasibility of an activity is related to whether the activity can be done in any quality or not. Feasibility of activities are mostly restricted to binary statements, either activity is feasible or not in a place.
Footprint	Footprint is the location of a place in the space. Footprint is one part of the gazetteers' model. Footprint of a place defines where the place is.
Form	Introduced as a part of an architectural view to place which includes 1-function, 2- form, and 3-space. Form is related to geometric properties and styles in buildings (places) (Canter, 1997).
Formal Dimension	Formal dimension is one of the four dimensions of objects of discourse. The four dimensions can be considered as dimensions of place as an object of discourse. It is defined as: "What kind of object something is: it is concerned with those properties that distinguish one category of things from another" (Couclelis, 2010). For example, the formal dimension of "Perejil is 'a tiny, uninhabited island near Ceuta'" (Couclelis, 2010).
Function	"The functional perspective points to the semantic level of functions that provides a sense of context by depicting the set of operations that the place supports" (Papadakis et al., 2016). Function is determining the agentive dimension of place as an object of discourse. In other words, a function suggests a specific composition for a place and determines the functionality of the place. Also, defined as part of architectural aspects of building (function, form, and space). This aspect is directly related to activities that happen in buildings.
Functional Differentiation	This facet is related to activities that happen in places –e.g., home: eating, resting
Functional Properties	Functional properties are defined as an affordance that a place can afford (Vasardani et al., 2016).

Generalized Properties (Place Properties)	"A generalized property is a place property that is assigned to the place as a whole. Generalized attribute values result from some kind of inferential process, such as a statistical inference on a sample, simulation, or algorithm, performed on place observations or other interpreted data" (Adams, 2015). The properties can be inferred based on the observed properties of the location. Are stored as a set of key/value pairs. "A generalization is similar to an observation with the following exceptions. A generalization is associated with one and only one place, unlike observations which can have multiple place associations. Instead of an observation procedure, a generalization is generated by an inference mechanism. Generalizations do not have a location, only a place association" (Adams, 2015) – e.g., Population count, temperature seasonality, median income.
Geographic Location	"A place is a unique spot in the universe. Place is the distinction between here and there, and it is what allows people to appreciate near and far. Places have finitude, but they nest logically because the boundaries are (analytically and phenomenologically) elastic. A place could be your favourite armchair, a room, building, neighbourhood, district, village, city, county, metropolitan area, region, state, province, nation, continent, planet–or a forest glade, the seaside, a mountaintop" (Gieryn, 2000). Geographic location is part of a model introduced by Gieryn (2000) which includes: 1- geographic location, 2- material 3- form, and 4- invested meaning and value.
Good shape	"The property of good shape of a place is the result of the cumulative good shape of its parts, such that if the individual centres and their surroundings are characterized by good shape by experiencing a few or all the aforementioned properties, then a place as the sum of its parts is also in good shape" (Vasardani and Winter, 2016).

Gradients	Gradient or graded variations are related to a change or a contrast which is happening in a smooth way. It is a part of place properties introduced by Vasardani and Winter (2016), and described as “[...] gradients that create harmonious transitions between contrasting centres and they all belong to the same whole structure".
Identification	In place reference system, place identification is described as a part of the model which can be used as an id to determine whether an object is a place or not. The id can be used to check whether two places are the same. This part of the model answers these two questions: Is it a place? Are these places the same? (Scheider and Janowicz, 2014)
Investment with Meaning and Value	“Without naming, identification, or representation by ordinary people, a place is not a place. Places are doubly constructed: most are built or in some way physically carved out. They are also interpreted, narrated, perceived, felt, understood, and imagined". “[...], the meaning or value of the same place is labile—flexible in the hands of different people or cultures, malleable over time, and inevitably contested" (Gieryn, 2000).
Level of Scale	“Formations within and across places can exhibit different scales" (Vasardani and Winter, 2016).
Local Symmetries	“Local symmetries are experienced in both natural and artificial places where different foci create symmetrical neighborhoods around them" (Vasardani and Winter, 2016).
Locale	“Locales are effectively affordances, that is to say the properties of a location which allow a particular activity to take place. Locales need not be tied to a fixed location, for example a bus or train allows particular sets of activities and interactions, in the same way as a church or indeed a mountain" (Purves and Derungs, 2015). “Apart from the location of a place, however, one of its most apparent and observable characteristics is its material content, also referred to as its landscape, materiality, setting or locale" (Jonietz, 2016).

Localization		In place reference system, localization is described as the location of a place at a certain time. This part aimed at answering the following 'where question': Where is a place? (Scheider and Janowicz, 2014)
Location		"Location is defined as a location in space which can be named and thus, at least implicitly, assigned coordinates" (Purves and Derungs, 2015). "A place and its location are inseparably connected. A location can be interpreted as an absolute point in space which can be clearly determined by an x, y, z coordinate tuple and allows geometrical analysis such as measurements of distances to other locations" (Jonietz, 2016).
Material (Physicality)	Form	"Place has physicality. Whether built or just come upon, artificial or natural, streets and doors or rocks and trees, place is stuff. It is a compilation of things or objects at some particular spot in the universe. Places are worked by people: we make places and probably invest as much effort in making the supposedly pristine places of Nature as we do in cities or buildings" (Gieryn, 2000). Material form is part of a model introduced by Gieryn (2000) which includes: 1- geographic location, 2- material 3- form, and 4- invested meaning and value.
Meaning		"Individual and group meanings (of a place) created through people's experience and intentions in regard to that place" (Relph, 1976).
Narrative Descriptions		"Stories are told in order to help characterize the uniqueness of a place as well as to define normative/acceptable behavior, by revealing the past actions of others. On a continuum with a place defined by an extensive narrative history, is a place defined by a single event (e.g. Chernobyl, Lockerbee, Three Mile Island)" (Jordan et al., 1998). Narrative descriptions is part of a model which includes: 1- physical features, 2- actions, 3- narrative descriptions, 4- symbolic representation, 5- socioeconomic and cultural factors, 6- typologies/categorisations.

Not Separateness	Not separateness is defined as “some of them (borders of places) distinct, some others gradient, but together they blend in, creating non-separated places" (Vasardani and Winter, 2016).
Object of Action	Object of action is a part of an experiential account for place which includes 1- action, 2- objects of action, and 3- valuation (Hockenberry, 2006). This part of the model defines what objects are related to a specific action. For example, the action of drinking happens in a bar, and objects of action such as beer can be found in the place.
Observed Properties (Environment Properties)	“An observed property is the result of a measurement of the environment by a sensor, whether it be a mechanical sensor such as a temperature gauge or a human sensor who records a written description of a place, at a specific location and time" (Adams, 2015).
Others	Others is a part of a three-pole model which is presented by Gustafson (2001). Others is a facet that is related to social relations and norms. “Places may also be associated with ‘others’ without reference to any social relations or encounters. In these cases, places are attributed meaning through the perceived characteristics, traits and behaviours of their inhabitants" (Gustafson, 2001).
Others-Environment	Others-environment is rarely mentioned in Gustafson’s empirical study. He described it as meanings related to place regarding others and environments such as the atmosphere of a place, and street-life of a place (Gustafson, 2001).
Parts	Parts is introduced in a place-related model which includes: 1- attributes, 2- activities, and 3- parts (Tversky and Hemenway, 1983). Parts are described as elements of the scenes such as ‘bird’, ‘sand’ for describing beaches.
Physical Dimension	A dimension of place that represents the physical structure of the place – e.g., a shopping mall: a set of buildings

Physical Elements	Places have physical elements which have an impact on the bonding of people to them. The dimension of places in place attachment is measured by the prominence of social and physical elements – e.g., density, proximity, presence of amenities.
Physical Features	Physical features can be described as a collection of objects. "Places consist of collections of objects. Each person perceives some set of affordances for a given small-scale object (e.g., a cup, a door handle, or a coffee pot) or collection of objects in large-scale space (e.g., a room, a house, or a restaurant)" (Jordan et al., 1998). Physical features is part of a model which includes: 1- physical features, 2- actions, 3- narrative descriptions, 4- symbolic representation, 5- socioeconomic and cultural factors, 6- typologies/categorisations.
Physical Landscape	Physical landscape is a part of place elements (location, physical landscape, and sense of place). Physical landscape is described as the total sum of material (both natural and cultural).
Physical Setting	Physical setting is described as answers to the following questions: Where is it? What are the place physical characteristics?
Place Affect	"Reflecting the importance of social relationships and the context within which they occur. The specific settings of the place share the meanings attributed to them by the individual's social environment" (Almuzaini, 2017).
Place Attachment (1)	Place attachment is the bonding that occurs between individuals and their meaningful environments (places). Place attachment has three dimensions: person (an individual or a group), place (with spatial level, specificity, and prominence of social and physical elements), and process (which can view as affective, cognitive, and behavioural processes) (Scannell and Gifford, 2010).

Place Attachment (2)	Place attachment is the closest component part of a sense of place. “It has often been said that place attachment is the environmental psychologist’s term for the geographer’s concept of sense of place" (Vanclay, 2008). “Place attachment refers specifically to the extent to which an individual has positive feeling about their local environment and/or community" (Vanclay, 2008). Place attachment is also called place connectedness, connection to place, or place bonding. Here, place attachment is considered as a component part of a sense of place.
Place Commitment	“Place commitment refers to the extent to which individuals are willing to contribute to their local place" (Vanclay, 2008) –e.g., contribution such as repairing houses in a neighbourhood
Place Dependence (1)	“How far the place satisfies the individual’s behavioural goals as compared to other alternative" (Almuzaini, 2017)
Place Dependence (2)	“I consider place dependence to be more-or-less similar (to place identity), although strictly speaking it might be seen as the self-perceived strength of association between an individual and a specific place" (Vanclay, 2008). Here, place dependence is defined as a component of a sense of place.
Place Familiarity	“Place familiarity and place awareness are similar concepts which relate to the extent of knowledge an individual has about a specific place or local environment" (Vanclay, 2008). Here, it is defined as a part (component) of a sense of place.
Place Identity (1)	Relph (1976) has defined identity of place as “persistent sameness and unity which allows that [place] to be differentiated from others." “There are three aspects which determine a place’s identity: its physical setting, the particular actions which can be conducted there, and any additional meaning which it is allocated, such as shared cultural values or personal memories" (Jonietz, 2016).

Place Identity (2)	Place identity is defined as a component part of a sense of place that "refers more specifically to the extent to which a person's identity is vested in the local place" (Vanclay, 2008).
Place Name (Toponym)	A name or names associated with a place. It can be an official name or a vernacular one. A place can have one or more names which can be used a reference to the place. A name of a place can be changed during the time.
Place Objectives	The facet is about "individual, social and cultural aspects of place" (Canter, 1997). "In other words, different aspects of the goals that a person has in a place. The distinct constituents each lead to a proposed distinct element i.e., individual, social and cultural" (Canter, 1997).
Place Reference (Verbal Reference)	"Verbal reference to named and unnamed places, or more complex verbal place descriptions containing several such references linked by spatial relationships are different from the common localization approach of geographic information systems and spatial databases" (Winter and Freksa,2012). "People may refer to locations by places' names, such as 'Cardiff University', or by associating different references with spatial relationships explicitly, such as 'in front of the cinema', or implicitly: 'Hilton, Cardiff' implying the Hilton hotel in Cardiff. 'In front of the cinema' is a directional relationship with the place, while 'Hilton, Cardiff' has a topological relationship that implies hierarchy. People can also refer to location by providing relatively exact descriptions of locations, for example, by stating an address, for example, 'the place is at 5 High St., near Liberty Square', or pin point it on a map" (Almuzaini, 2017).
Place Social Bonding	Place social bonding is defined as "reflecting the emotional or affective bond between an individual and a place" (Almuzaini, 2017).
Positive Space	"Positive spaces are what differentiate the figure from the ground in Gestalt theory and they are usually associated with the convex footprints of spatial features especially in the built environment, vs. the background space that surrounds them" (Vasardani and Winter, 2016).

Purpose	Purpose is determining the telic dimension of place as an object of discourse. In other words, it defines why a place as an object of discourse exists. Considering places as objects of discourse one may consider the purpose as a facet which defines why the place exists.
Qualities (1)	A place can be described using adjective terms defining the qualities about the place. People use the qualities as descriptive terms to describe a place using natural languages. Qualities are introduced as part of a model for place descriptions, i.e., elements, qualities, and activities (Edwardes and Purves, 2007a; Edwardes and Purves, 2007b). Qualities are related to the sense of the place or the subjective idea of a person regarding the place.
Qualities (2)	Qualities are defined as a superset of affordances. It includes observations (producing values), and affordances (producing actions). While Ortmann and Kuhn (2010) define qualities and affordances not only for places but for objects. However, it can be used as a part of place definition.
Quality of Place	The quality of a place is defined as a combination of two factors, spatial accessibility, and suitability of the place for a specific activity (Jonietz, 2016).
Roughness	Roughness is defined as “Roughness has to do with the irregularities and imperfections in any place configuration, as the result of organic generative process of development that do not adhere to perfect geometric symmetries" (Vasardani and Winter, 2016).
Salience of Place (Absolute view)	“On an absolute level, the salience of a place can be defined as being irrelevant to the attachment to specific individuals" (Almuzaini, 2017). Hence, the salience of place in absolute level is an objective measure that defines how important a place is.
Salience of Place (Personal view)	“The salience of a place can be described from a personal or from an absolute point of view. On a personal level, many factors can influence the importance of a place to an individual. This includes: place dependence, place affect, and place social bonding" (Almuzaini, 2017).

Salience	"Several properties contribute to the salience of an object: its singularity, its prominence, its accessibility, its meaning or cultural significance, and prototypicality. This element helps to distinguish landmarks in an environment" (Almuzaini, 2017).
Scale of Interaction	Scale of interaction is related to the scale in which the place belongs to (in which people interact with places).
Self	Self is a part of a three-pole model which is presented by Gustafson (2001). Self is related to personal meaning about a place and the place's self identification. The theme of 'life-path', 'emotion', and 'activity' belongs to this pole (Gustafson, 2001).
Self-Others	Self-others is related to a type of information that their meaning is related to both self and others. "Places often become meaningful because of the respondents' relations with people living there - friends, acquaintances, relatives - and the sense of community that such social relations create" (Gustafson, 2001).
Self-Others-Environment	Meaning related to places which involve self, others, and environment. "Traditions, festivals and anniversaries often implicate self, others and various environments (local as well as national). Similarly, when the respondents' membership in spatially defined associations or organizations makes the place meaningful, it is clear that self, others (other members) and the environment (geographical and sometimes institutional) contribute to the overall meaning of place" – e.g., citizenship (Gustafson, 2001).
Semantic Relation to Activities	Semantic relation to activities is part of a model designed to capture references to places in textual information (Scheider and Purves, 2013). Activities can be used in/as a reference to localize places in a description – e.g., a biking site.
Semantic Relation to Objects	Semantic relation to objects is part of a model designed to capture references to places in textual information (Scheider and Purves, 2013). Objects can be used in/as a reference to localize places in a description – e.g., a coffee place.

Semantic Relation to Qualities	Semantic relation to objects is part of a model designed to capture references to places in textual information (Scheider and Purves, 2013). Qualities can be used in/as a reference to localize places in a description – e.g., a warm place.
Sense of Place (1)	"Sense of place relates to the ways in which we might (or indeed might not) identify with a place and is necessarily based on our associations and experiences, or lack thereof, of a place" (Purves and Derungs, 2015). "In general, however, the notion of a sense of place involves, apart from its activities, additional associations to place which are typically of affective and intangible nature" (Jonietz, 2016).
Sense of Place (2)	Sense of place is defined by its components such as place attachment, place identity, place dependence, place familiarity, place commitment, and place satisfaction (Vanclay, 2008).
Sentiment (Individual Reflection)	An individual reflection is defined as a part of a place model which clarifies what people feel about a specific place (ElGindy and Abdelmoty, 2014b; ElGindy and Abdelmoty,2014a).
Service	Service is defined as an answer to what kind of services a place provides for an individual person.

Simplicity and Inner Calm	"This property is a more abstract one perhaps, but one well understood and agreed upon, nonetheless, when describing places of preference, of which it more often than not, is a characteristic. It abides to the notion that simplicity offers a feeling of calmness, or that places function better when they are not overloaded with elements or functions, or both. Think of the words one uses when describing shopping malls. Usually, and depending on the previous experiences, expressions such as 'crowded', 'busy', 'difficult to navigate through', 'confusing', or 'tiring', often appear in the descriptions for such places where many needs and many functions are attempted to be accommodated, such as shopping for various categories, eating, resting, recreation and others in one single place" (Vasardani and Winter, 2016).
Social Elements	Place plays a key role in human social life, and consequently place attachment is the results of this relationship between people and places. Place attachment is measured by the prominence of social and physical elements. Hence, this property is related to the social aspects of place attachment.
Social Role (Social Dimension)	"Representing what patterns of social interaction occurs there" (Ballatore, 2016) – e.g., a shopping mall: trading activities.
Social Role Characteristics	The social role characteristics are related to activities happen in a place. For example, one may consider the role of a beach as Place-Where-Sun-Bathing. The social role characteristics of a place are anti-rigid, emergent, and dynamic.

Socioeconomic and Cultural Factors	"People identify themselves with places socioeconomically. For example, sea ports are special socioeconomic places since they afford transportation and trading, therefore, they afford a certain type of economic climate. Ports need people to work there (e.g., loading and unloading ships), but at the same time they attract those who are buying and selling goods. Similarly, different cultures afford different behaviour in places. For example, black is the colour of mourning in the west, whereas in China it is white. Williams (1981) views culture as a system through which a certain order is communicated and experienced" (Jordan et al., 1998). Socioeconomic and cultural factors is part of a model which includes: 1- physical features, 2- actions, 3- narrative descriptions, 4- symbolic representation, 5- socioeconomic and cultural factors, 6- typologies/categorisations.
Space	Space is the term used as a part of place (building) model based on an architectural view. The model includes 1- function, 2- form, and 3- space. Space is related to the location of the places (buildings) (Canter, 1977).
Spatial Accessibility	"The relative accessibility of a place's location, which can be defined as the inherent characteristic (or advantage) of a place with respect to overcoming some form of spatially operating source of friction (for example, time and/or distance) is of relevance. In fact, a place's central location in the movement network might be a better indicator for high usage frequencies than its design" (Jonietz, 2016).
Spatial Band	Spatial band is part of a model which includes: 1- spatial band, 2- temporal band, and 3- thematic band (Janowicz, 2012). Spatial band represents the spatial patterns of a place. A specific place type has a specific spatial pattern (considering the other place with similar and different type).

Spatial Identity	"A place has to be identifiable to exist: it has a spatial identity that allows it to differentiate itself from other Places. This first definition of spatial identity is associated with the Relph's 'identity of Place', who explains that some Places may share joint attributes. The spatial identity of a Place may also be shared by a number of individuals. On the other hand, spatial identity can also be defined as the identification of an operator in a particular Space. This second acceptance refers to Relph's 'identity with Place'. As highlighted by Stock, identity is 'an arbitrary link between Place and people' and 'is not only linked to representations but also related to the practice of Places'" (Quesnot and Roche, 2015).
Spatial Level	Spatial level determines the scale of interaction in which the people are bonded to places –e.g., home, neighbourhood, city.
Spatial Properties	"The dimension of spatial properties reflects the semantic level of classification, describing place as a set of properties (e.g., geometries)" (Papadakis et al., 2016). Spatial properties of a place define spatial patterns which describe the relationship between the place and space.
Spatial Reference	"Spatial references are locating something in the world, as in 'I am at the bus stop'" (Winter, 2009). Spatial reference is one of the ways to describe a place in a place description.
Spatial Relation to Other Places	Spatial relation to other places is part of a model designed to capture references to places in textual information (Scheider and Purves, 2013). Other places can be used in/as a reference to localize a place in a description – e.g., "the location of land parcels may be described by relative spatial relations to surviving buildings, such as churches" (Scheider and Purves, 2013).

Spatial Value	"Lussault defines the spatial value as the totality of socially valuable qualities of a space" . "Within the field of geography, the value of a Place is defined by its relationships maintained with individuals. These relations are both physical (i.e. concrete practices) and philosophical (i.e. everything that is related to the imagination and representations)". "This notion of spatial identity is closely linked to the spatial value" (Quesnot and Roche, 2015).
Specificity	Specificity is a measure which is defined to answer: how a place, in the sense of attachment, is different from other places in the same type.
Spirit of Place	"Spirit of place, or genius loci, is a more appropriate (compared to sense of place) term when referring to the qualities of a place that make it special" (Vanclay, 2008). This definition is similar to the definition of sense of place in human geography literature. In here, the author claims that sense of place is more about individuals rather than places. He also mentioned that spirit of place is similar to the term 'topophilia' (love of place) mentioned by Yi-Fu Tuan.
Strong Centres	"In any 'place' formation, there exists at least one, and in most cases multiple centres of varying sizes, which support each other. This concept is clear in the central-place theory in geography and its later modifications" (Vasardani and Winter, 2016) – e.g., central district for a city.
Structural Properties	Structural properties includes 15 properties of place: 1- level of scale, 2- boundaries, 3- positive space, 4- local symmetries, 5- contrast, 6- roughness, 7- the void, 8- not separateness, 9- strong centres, 10- alternating repetition, 11- good shape, 12- deep interlock and ambiguity, 13- gradients, 14- echoes, 15- simplicity and inner calm (Vasardani et al., 2016).

Suitability (of an activity)	Suitability of activities is related to "influences the allocation of place utility, the level of satisfaction expected from interacting with this particular place, which correlates with the probability of it being selected and actually used". "[...] its [a place's] suitability, meaning its appropriateness with regards to a particular activity, tend to vary among its prospective users" (Jonietz, 2016).
Symbolic Representation	"Certain places are referenced by symbols (e.g., New York City is often referenced as the 'Big Apple') having symbolic and/or mythical meanings. Similarly, the Statue of Liberty is a common symbol for New York, related to its history as point of entry for many U.S. immigrants" (Jordan et al., 1998). Symbolic representation is part of a model which includes: 1- physical features, 2- actions, 3- narrative descriptions, 4- symbolic representation, 5- socioeconomic and cultural factors, 6- typologies/categorisations.
Telic Dimension	Telic dimension is one of the four dimensions of objects of discourse. The four dimensions can be considered as dimensions of place as an object of discourse. It is defined as "The purpose of things or the reasons why things happen" (Couclelis, 2010).
Temporal Band	Temporal band is part of a model which includes: 1- spatial band, 2- temporal band, and 3- thematic band (Janowicz, 2012). Temporal band represents the temporal behaviour of a place as a pattern – e.g., a restaurant is crowded in specific hours in a day (around 12 pm, and 8 pm).

The Void	"The property of void is the one that exists in all places, natural and artificial, rural or urban, in micro- or macro-scales. It is that of empty space, a requirement for any place configuration through which people can move around, where space is not completely filled. It is also one of the properties that allows for people to use each place according to their needs at any given context." "The void, emptiness or openness of the space in such places allows for its creative and individualized use, be it for recreation, sport, a meeting, or a picnic among others" (Vasardani and Winter, 2016).
Thematic Band	Thematic band is part of a model which includes: 1- spatial band, 2- temporal band, and 3- thematic band (Janowicz, 2012). Thematic band represents the theme or non-spatiotemporal aspects of a place.
Type	Type of a place is a part of a gazetteer's model. The type of place defines what kind of place it is.
Typologies	"People categorize places in order to understand what is new, in terms of what is already understood. This represents an important mental strategy for dealing with complexity and new situations." "Our model of place would allow for comparisons of places based on the means-end hierarchy. If two places have very different physical features, and yet provide the required affordances for a given agent and task, then they can be classified as similar places" (Jordan et al., 1998). Typologies/Categorisations is part of a model which includes: 1- physical features, 2- actions, 3- narrative descriptions, 4- symbolic representation, 5- socioeconomic and cultural factors, 6- typologies/categorisations.
Valuation (of an action)	Valuation of action is a part of the experiential model for place which contains (action, objects of action, and valuation) designed by Hockenberry (2006). This part of model defines the quality of an action which happens in a place – e.g., either 'good', 'greater than average' or 'not bad' for a meal at a restaurant.

APPENDIX B

B.1 PART-OF-SPEECH TAGS

Table 43 shows a list of selected part-of-speech tags. More detail can be found in https://www.ling.upenn.edu/courses/Fall_2003/ling001/penn_treebank_pos.html.

B.2 TYPED DEPENDENCIES

Table 44 shows the dependency types and their descriptions. More details can be found in https://universaldependencies.org/u/dep/index.html.

Tag	Description
CC	Coordinating conjunction
CD	Cardinal number
DT	Determiner
IN	Preposition or subordinating conjunction
JJ	Adjective
JJR	Adjective, comparative
JJS	Adjective, superlative
MD	Modal
NN	Noun, singular or mass
NNS	Noun, plural
NNP	Proper noun, singular
NNPS	Proper noun, plural
POS	Possessive ending
RB	Adverb
RBR	Adverb, comparative
RBS	Adverb, superlative
VB	Verb, base form
VBD	Verb, past tense
VBN	Verb, past participle
WDT	Wh-determiner
WP	Wh-pronoun
WRB	Wh-adverb

Table 43: A selected list of part-of-speech tags

Dependency type	Description
acl	clausal modifier of noun (adjectival clause)
advcl	adverbial clause modifier
advmod	adverbial modifier
amod	adjectival modifier
appos	appositional modifier
aux	auxiliary
case	case marking
cc	coordinating conjunction
ccomp	clausal complement
conj	conjunct
csubj	clausal subject
dep	unspecified dependency
nmod	nominal modifier
nsubj	nominal subject
nummod	numeric modifier
obj	object
punct	punctuation

Table 44: A selected list dependency types

APPENDIX C

C.1 IMPLEMENTATION

C.1.1 *Scale Encoding*

In scale encoding, the scale values of OpenStreetMap (OSM) are mapped into the proposed scale schema. The mapping is shown in Table 45.

C.1.2 *Prominence Encoding*

To derive prominence levels from OSM importance values, 'natural breaks' classification is used. First, all the importance values are collected, then based on the histogram the natural break method calculates the boundaries of each prominence level. Figure 29 shows the histogram and the derived boundaries for the prominence levels.

C.2 COMPARING SIMPLE-WHERE QUESTIONS AND DETAILED WHERE-QUESTIONS

Figure 30 shows the distributions of place types in SWQs and DWQs. In DWQs, administrative divisions are almost twice as frequent as in SWQs. The extra detail (additional toponyms) included in the content of DWQs dominantly refers to administrative divisions. For example, *Where in California is Disneyland?* can be viewed as a SWQ, *Where is Disneyland?*, with extra details that shows what the inquirer already knows about its coarse location (*in California*). Hence, the differences

Table 45: Scale mapping from OSM values to proposed schema

Scale level (proposed schema)	OSM scale levels	Description
1	value $\geqslant 27$	buildings and houses
2	$22 \leqslant$ value < 27	airports, roads and streets
3	$18 \leqslant$ value < 22	suburbs and districts
4	$16 \leqslant$ value < 18	cities, islands and villages
5	$12 \leqslant$ value < 16	counties
6	$8 \leqslant$ value < 12	states and regions
7	$4 \leqslant$ value < 8	countries
8	value < 4	oceans, seas and continents

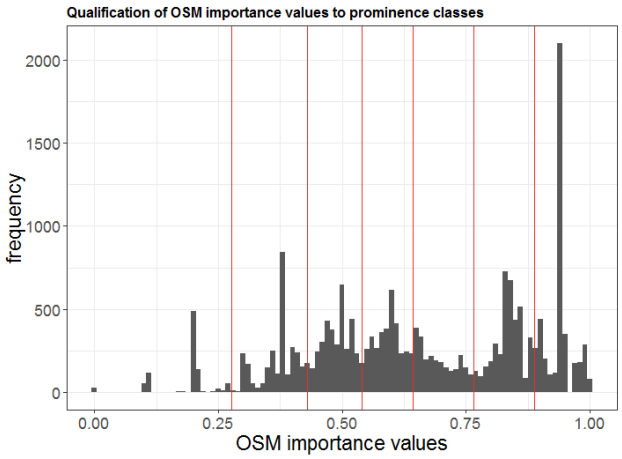

Figure 29: Histogram of importance values and boundaries of prominence levels.

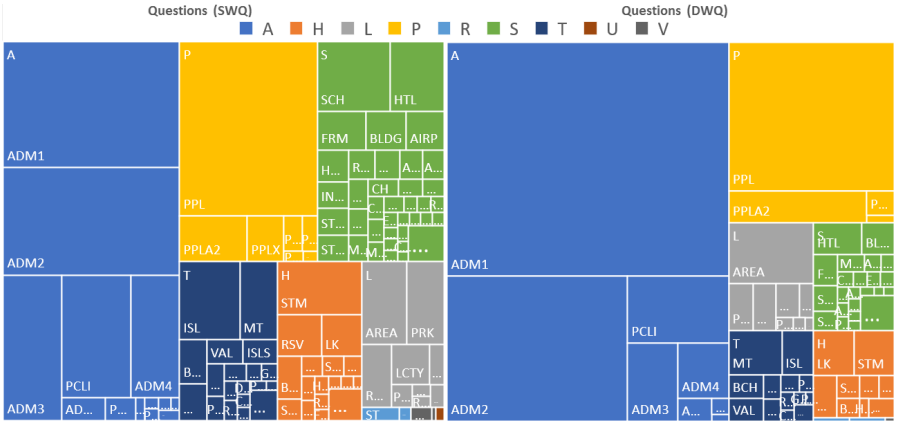

Figure 30: Distribution of types in the SWQs versus DWQs.

between the two distributions illustrate where the additional detail is included.

Answers to SWQs (Figure 31) have a linearly incremental distribution (except for the coarsest level of scale). The figure shows the answers are generated with toponyms that belong to *coarser levels* of scale compared to places mentioned in the SWQs. The distribution of scale in the answers to DWQs shows a dominance of references to places of scale *between* the two mean values of the questions' bi-modal distribution. Thus, people tend to generate responses at a level of scale greater than that of the place that is asked for, and lower than the scale of the places mentioned as additional details. This is because places at a same or coarser level of scale than the places in questions would lead to obvious, and hence irrelevant, answers. The clear distinction of the difference in answering SWQs and DWQs are shown in Table 46, enabling the comparison of the levels of scale of places identified in the questions and their answers.

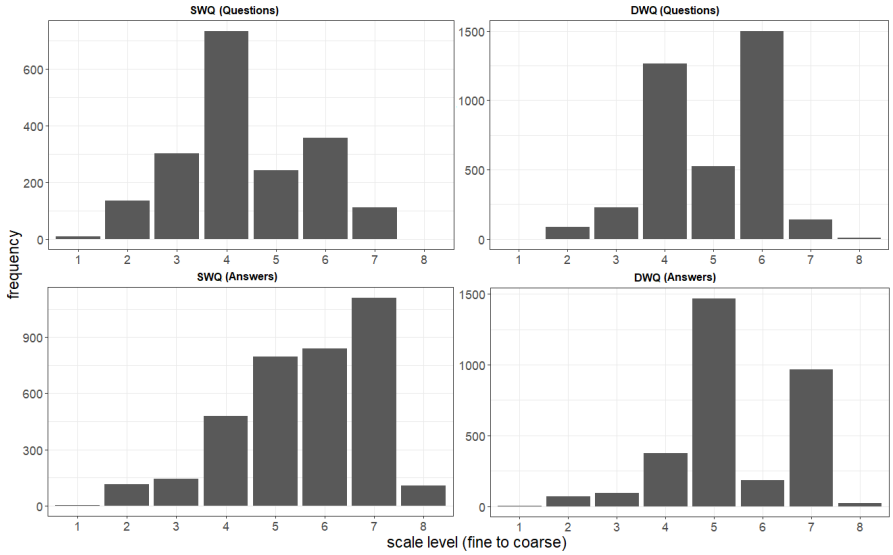

Figure 31: Distribution of the levels of scale in the SWQs versus DWQs.

Table 46: Comparison of levels of scale in the answers compared to their questions

Level of scale in answers	Lower than	Equal (Between)	Greater than
SWQ			
each value	17.0%	13.1%	**69.9%**
min value	25.8%	16.0%	**58.2%**
median value	14.5%	14.9%	**70.6%**
max value	10.6%	5.9%	**83.5%**
DWQ			
each value	7.2%	**59.6%**	33.2%
min value	10.5%	**77.7%**	11.8%
median value	6.7%	**77.3%**	16.0%
max value	3.9%	40.8%	**55.3%**

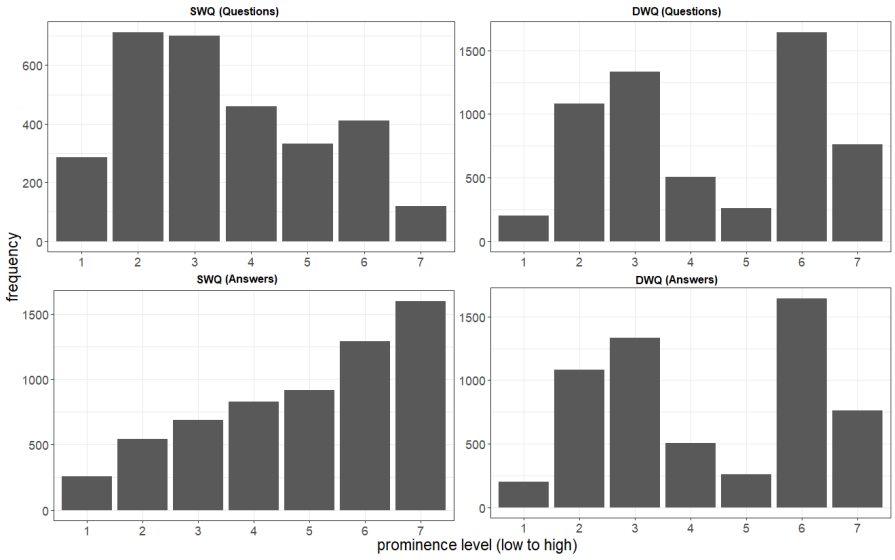

Figure 32: Distribution of prominence levels in the SWQs versus DWQs

Figure 32 shows the distribution of prominence in SWQs, DWQs, and their answers. Based on the figure, people ask about less-known places and answer referring to well-known ones. In DWQs, both less-known and well-known places are frequently observed. The well-known places in DWQs are the details presented in the content of the questions. In DWQs, inquirer's state of knowledge can be estimated better, because the question contains what they already know and what they want to find out. Similar to the pattern observed based on scale, prominence levels of places in the answers to DWQs are mostly between the levels of prominence of the toponyms mentioned in their corresponding questions (Table 47).

C.3 SEQUENCE DISTRIBUTIONS AND FREQUENT PATTERNS

C.3.1 *Sequence Distributions*

Sequence distributions capture not only the content of sequences, but the way that the content is generated in the sequence (i.e., *style*). The length of sequences can be further investigated using sequence distributions. Figure 33 shows the sequence distributions of type, scale and prominence in the questions and answers. Here, only the first five positions of answer-sequences are visualised. While this long-tailed distribution contains answers up to a length of 13, the vast majority (95.84%) of answers have less than five toponyms. Most of the questions have only one or two toponyms, and the answers are mostly short as well (usually with less than three toponyms in their content). In Figure 33, only the top ten types are visualized, with the rest grouped as *OTHER*.

Table 47: Comparison of prominence levels in the answers compared to their questions

Prominence level in answers	Lower than	Equal (Between)	Greater than
SWQ			
each value	19.8%	12.5%	**67.7%**
min value	30.7%	14.7%	**54.6%**
median value	19.4%	9.5%	**71.1%**
max value	9.7%	7.9%	**82.4%**
DWQ			
each value	8.7%	**68.4%**	22.9%
min value	13.0%	**77.9%**	9.1%
median value	6.8%	**79.7%**	13.5%
max value	3.5%	**55.6%**	40.9%

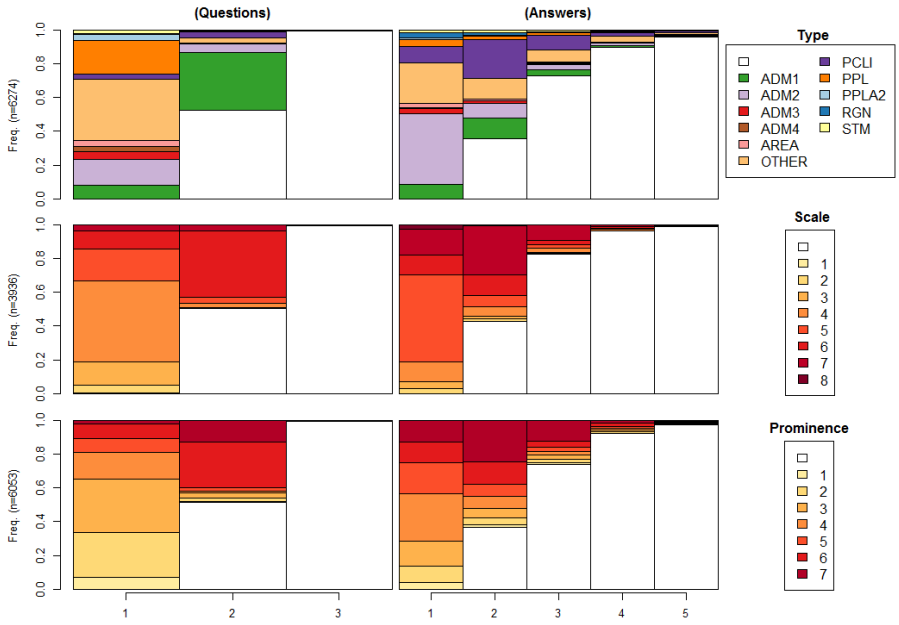

Figure 33: Sequence distributions of the questions and answers

Based on the distributions of type, scale and prominence in the questions, the details in DWQs are likely to be found as the second toponym. In Figure 33, one can observe that these details usually belong to the type of ADM1[1] (first-order administrative divisions – i.e., states) or PCLI (independent political entities – i.e., countries). Similar patterns can be observed using scale and prominence – i.e., the second elements in questions belong to well-known places which are in coarse levels of scale.

Figure 33 shows the strong role of administrative places in the human-generated answers in this corpus. The answers are generated with mid-levels to coarse-levels of scale and presented starting with mid-levels followed by coarse levels in most cases. Similar patterns can be observed in terms of prominence; however, places in lowest and highest prominence levels can be found in the answers. When comparing the questions and answers based on type, the questions are asked about a diverse range of place types while the answers are mostly generated using the top ten types (i.e., the less frequent types labelled as *OTHER* are almost twice as frequent in the questions compared to answers).

C.3.2 *Frequent Patterns*

Frequent patterns describe the generalized patterns in human answering behavior. Stronger patterns of a generic class (e.g., scale) show the usefulness of the generic information to describe the human-generated answers. Figures 34, 35 and 36 illustrate the top ten patterns based on type, scale and prominence, respectively. These patterns are shown for all answers, answers to SWQs, and answers to DWQs.

Figure 34 shows that the top ten patterns in the answers are mostly generated with first and second orders administrative divisions and independent political entities. In the generic class, *type*, almost 60% of all answers are described through only ten sequence patterns. Ten patterns also cover 70% and 50% of the answers to DWQs and SWQs, respectively. The differences shows that answers to DWQs are more describable based on place type than the responses to SWQs. This is because in DWQs, some details are provided and people can therefore infer the inquirers' information needs better. In SWQs, due to lack of context leads to more ambiguity and likely requires more subjective judgments in responses.

Figure 35 illustrates the top ten patterns based on scale. The coarse levels of scale constitute most of the frequent patterns. Answers covered by these patterns show that scale is an important generic information characterizing well human-generated answers.

1 The codes for place types are described in Section C.5.

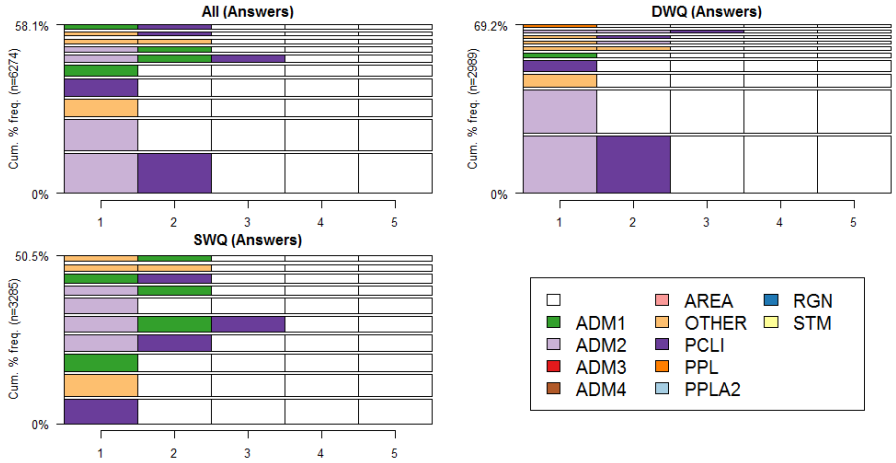

Figure 34: Top ten frequent pattern in type-sequences of the answers

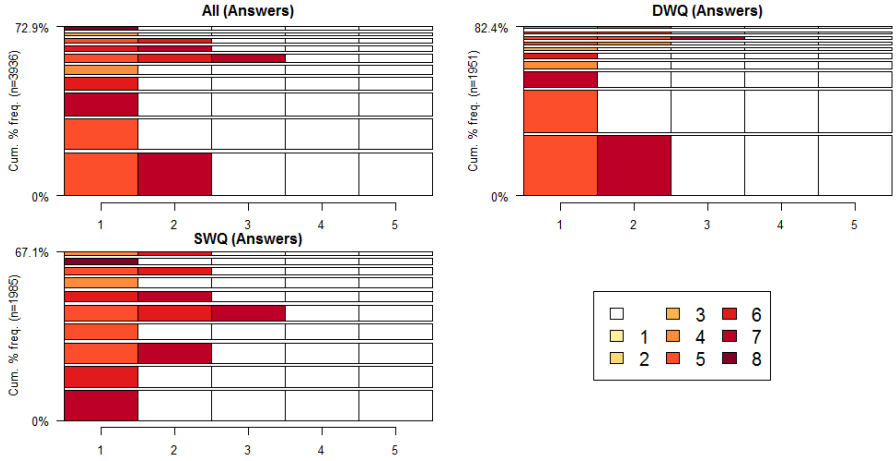

Figure 35: Top ten frequent pattern in scale-sequences of the answers

The comparison of the answers to DWQs and SWQs shows similar results to types. More than 80% of the answers to DWQs can be described using ten patterns based on scale. The style of these patterns shows that the answers are hierarchically presented (starting from finer levels followed by coarser levels of scale).

Figure 36 shows the ten most frequent patterns based on prominence levels. Most of the patterns are constructed with high levels of prominence. Similar to scale, a strong pattern in the style of answers can be observed – i.e., starting with less-known places followed by well-known ones. The patterns based on prominence have, however, less support in the data compared to the patterns derived from scale or type sequences. This observation can be explained by prominence being more tightly related to the specific context of the questions, compared to scale or type. For example, finding a highly prominent place reference may not be always possibleas prominent places are not uniformly distributed in the world.

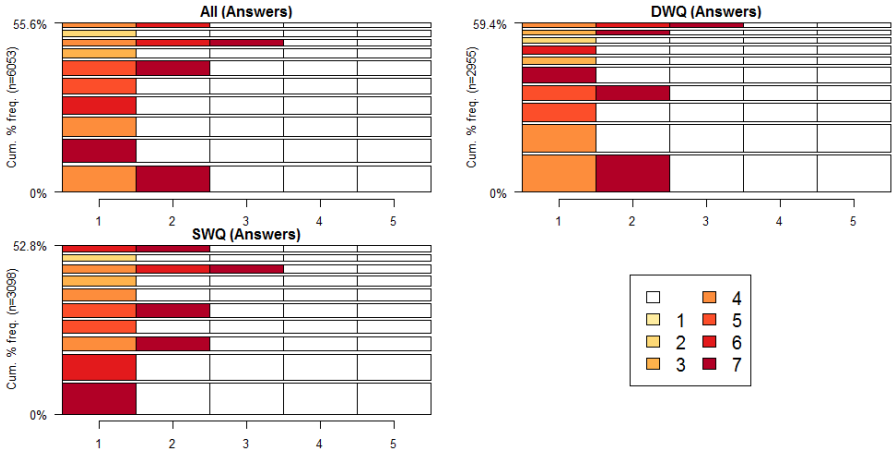

Figure 36: Top ten frequent pattern in prominence-sequences of the answers

C.4 DEMONSTRATION EXAMPLES

In this section, four examples are provided to demonstrate how generic form of answers (e.g., type-sequence) can be translated into their specific form (toponym-sequence).

Example 1: *Where is Nagasaki? In Japan*. The specific form and TSP encoding of this question-answer pair are shown below:

- specific representation (question): [Nagasaki];
- TSP encoding (question): type-sequence [ADM1], scale-sequence [6], prominence-sequence [4];
- TSP encoding (answer): type-sequence [PCLI], scale-sequence [7], prominence-sequence [7];
- specific representation (answer): [Japan]

Queries 12 and 13 show the SPARQL queries to derive the specific form of the answer using the information available in DBPedia and Geonames, respectively. The result of both queries is the unified resource identifier (URI) of Japan which is the correct specific representation of the answer. The query results are shown in Table 48.

```
PREFIX dbo: <http://dbpedia.org/ontology/>

SELECT distinct ?q1 ?a1 WHERE {
  VALUES ?q1 {<http://dbpedia.org/resource/Nagasaki>}

  ?a1 a dbo:Country .
  ?q1 ?r ?a1
}
```

Query 12: SPARQL query of Example 1 (DBPedia)

Table 48: SPARQL results to Example 1

Knowledge Base	Q1	A1
DBPedia	Nagasaki	**Japan**
Geonames	Nagasaki	**Japan**

```
PREFIX gn: <http://www.geonames.org/ontology#>

SELECT distinct ?q1 ?a1 WHERE  {
  VALUES ?q1 {<http://sws.geonames.org/1856156/>} .

  ?a1 gn:featureCode gn:A.PCLI .
  ?q1 ?r ?a1
}
```

Query 13: SPARQL query of Example 1 (Geonames)

Example 2: *Where in Illinois is Cahokia? In St. Clair County, Illinois, United States*:

- specific representation (question): [Illinois, Cahokia];
- TSP encoding (question): type-sequence [ADM1, PPL], scale-sequence [6, 4], prominence-sequence [6, 3];
- TSP encoding (answer): type-sequence [ADM2, PCLI], scale-sequence [5, 7], prominence-sequence [3, 7];
- specific representation (answer): [St. Clair County, United States]

Based on the SPARQL template, two queries (Queries 18 and 19) are used to translate the generic form of the answer to a specific form using DBPedia and Geonames. The results of the queries are shown in Table 51. Using DBPedia, the correct answer is among one of the three retrieved responses {[Collinsville, United States], [Illinois, United States], [St. Clair County, United States]}. The second response in the results can be easily filtered due to repetitive content (i.e., Illinois) considering the content of the question. However, the first one which is not mentioned in the human-generated answer cannot be avoided using only the type-sequence of the answer. Consequently, by considering the predicted scale-sequence of the answer, the correct response [St. Clair County, United State] can be derived (Table 51).

Table 49: SPARQL results to Example 2

Knowledge Base	Q1	Q2	A1	A2
DBPedia	Cahokia	Illinois	Collinsville	**United States**
			Illinois	**United States**
			St. Clair County	**United States**
Geonames	Cahokia	Illinois	**St. Clair County**	**United States**

```
PREFIX dbo: <http://dbpedia.org/ontology/>

SELECT distinct ?q1 ?q2 ?a1 ?a2
WHERE {
    VALUES ?q1 {<http://dbpedia.org/resource/Cahokia>}
    VALUES ?q2 {<http://dbpedia.org/resource/Illinois>}

    ?a1 a dbo:PopulatedPlace .
    ?q1 ?r ?a1 .
    ?a1 ?r2 ?q2 .

    ?a2 a dbo:Country .
    ?q2 ?r3 ?a2 .
}
```

Query 14: SPARQL query of Example 2 (DBPedia)

```
PREFIX gn: <http://www.geonames.org/ontology#>

SELECT distinct ?q1 ?q2 ?a1 ?a2
WHERE  {
    VALUES ?q1 {<http://sws.geonames.org/4234969/>} .
    VALUES ?q2 {<http://sws.geonames.org/4896861/>} .

    ?a1 gn:featureCode gn:A.ADM2 .
    ?q1 ?r ?a1 .
    ?a1 ?r2 ?q2 .

    ?a2 gn:featureCode gn:A.PCLI .
    ?q2 ?r3 ?a2 .
}
```

Query 15: SPARQL query of Example 2 (Geonames)

Example 3: *Where is the Danube River, Europe? It originates in Germany's Black Forest, and flows in a southeasterly direction through central and eastern Europe to the Black Sea*:

- specific representation (question): [Danube River, Europe];
- TSP encoding (question): type-sequence [STM, CONT], scale-sequence [4, 8], prominence-sequence [6, 7];
- TSP encoding (answer): type-sequence [PCLI, MTS, SEA], scale-sequence [7, 3, 7], prominence-sequence [7, 4, 5];

- specific representation (answer): [Germany, Black Forest, Black Sea]

The SPARQL queries for finding the specific forms of answers of Example 2 are presented in Queries 16 and 17. The results of these queries are shown in Table 50. Using DBPedia, the results for the mountain range and the sea that are related to the Danube are the Black Forest and the Black Sea, while the country is not unique and ten countries (including the right one, Germany, based on the human-generated answer) are found in relation to the river. Here, only the type-sequence of the answer is used, and using prominence the country lists would narrow down to Germany (the only country in Level 7 of prominence in the results). Interestingly, the results using Geonames are incorrect. In Geonames, the river is stored as a coordinate point, which belongs to one country. Moreover, the relationships between Black Forest (*originates*) and Black Sea (*flows to*) to the river are not stored in Geonames due to its limited list of supported spatial relations – i.e., Geonames supports only containment.

```
PREFIX dbo: <http://dbpedia.org/ontology/>

SELECT distinct ?q1 ?a1 ?a2 ?a3 WHERE {
  VALUES ?q1 {<http://dbpedia.org/resource/Danube>}

  ?a1 a dbo:Country .
  ?q1 ?r1 ?a1 .

  ?a2 a dbo:MountainRange .
  ?q1 ?r2 ?a2 .

  ?a3 a dbo:Sea .
  ?a3 ?r3 ?q1
}
```

Query 16: SPARQL query of Example 3 (DBPedia)

```
PREFIX gn: <http://www.geonames.org/ontology#>

SELECT distinct ?q1 ?a1 ?a2 ?a3 WHERE {
  VALUES ?q1 {<http://sws.geonames.org/791630/>}

  ?a1 gn:featureCode gn:A.PCLI .
  ?q1 ?r1 ?a1 .

  ?a2 gn:featureCode gn:T.MTS .
  ?q1 ?r2 ?a2 .

  ?a3 gn:featureCode gn:H.SEA .
  ?q1 ?r3 ?a3
}
```

Query 17: SPARQL query of Example 3 (Geonames)

Example 4: *Where is Golden Gate Bridge? It is located between San Francisco and Marin County, in the U.S. state of California.*:

- specific representation (question): [Golden Gate Bridge];

Table 50: SPARQL results to Example 3

Knowledge Base	Q1	A1	A2	A3
DBPedia	Danube	Moldova	**Black Forest**	**Black Sea**
		Ukraine	**Black Forest**	**Black Sea**
		Germany	**Black Forest**	**Black Sea**
		Romania	**Black Forest**	**Black Sea**
		Croatia	**Black Forest**	**Black Sea**
		Serbia	**Black Forest**	**Black Sea**
		Austria	**Black Forest**	**Black Sea**
		Bulgaria	**Black Forest**	**Black Sea**
		Hungary	**Black Forest**	**Black Sea**
		Slovakia	**Black Forest**	**Black Sea**
Geonames	Danube River	Romania	–	–

- TSP encoding (question): type-sequence [BDG], scale-sequence [5], prominence-sequence [4];
- TSP encoding (answer): type-sequence [ADM2, ADM2, ADM1], scale-sequence [5, 5, 7], prominence-sequence [4, 4, 7];
- specific representation (answer): [San Francisco, Marin County, U.S. state of California]

```
PREFIX dbo: <http://dbpedia.org/ontology/>

SELECT distinct ?q1 ?a1 ?a2 ?a3 WHERE {
    VALUES ?q1 {<http://dbpedia.org/resource/Golden_Gate_Bridge>}

    ?a1 a dbo:PopulatedPlace .
    {?a1 ?r1 ?q1} UNION {?q1 ?r1 ?a1} .

    ?a2 a dbo:PopulatedPlace .
    {?a2 ?r2 ?q1} UNION {?q1 ?r2 ?a2} .

    ?a3 a dbo:PopulatedPlace .
    {?a3 ?r3 ?q1} UNION {?q1 ?r3 ?a3} .
}
```

Query 18: SPARQL query of Example 4 (DBPedia)

Table 51: SPARQL results to Example 4

Knowledge Base	Q1	A1	A2	A3
DBPedia	Golden Gate Bridge	**San Francisco**	**Marin County**	**California**
Geonames	Golden Gate Bridge	**San Francisco**	–	**California**

```
PREFIX gn: <http://www.geonames.org/ontology#>

SELECT distinct ?q1 ?a1 ?a2 WHERE {
    VALUES ?q1 {<http://sws.geonames.org/5352844/>}

    ?a1 gn:featureCode gn:A.ADM2 .
    {?a1 ?r1 ?q1} UNION {?q1 ?r1 ?a1} .

    ?a2 gn:featureCode gn:A.ADM2 .
    {?a2 ?r2 ?q1} UNION {?q1 ?r2 ?a2} .


    ?a3 gn:featureCode gn:A.ADM1 .
    {?a3 ?r3 ?q1} UNION {?q1 ?r3 ?a3} .

}
```

Query 19: SPARQL query of Example 4 (Geonames)

C.5 FEATURE CODES

Table 52 shows the types which are mentioned in Chapter 6. The complete list can be found in the Geonames website.

Table 52: Feature codes used in Chapter 6 (extracted from Geonames documentation)

Code	Description	Example
ADM1	first-order administrative division (states, and provinces)	Oklahoma
ADM2	second-order administrative division (counties)	Brevard County
ADM3	third-order administrative division (cities)	City of Alhambra
ADM4	fourth-order administrative division (towns)	Newburgh
AREA	a part of land without homogeneous character/boundaries	Theresienwiese
BDG	a bridge	Putney Bridge
FRM	a part of land dedicated to agricultural purposes	Branksome
HTL	hotels	The Carriage House
MT	mountains	Eagles Nest
PCLI	independent political entity	Paraguay
PPL	diverse type of populated places (e.g., cities, and villages)	El Granada
PPLA2	seat of a second-order administrative division	Lake City
PRK	parks and recreational places	Franklin Square Park
RGN	an area with particular cultural character	Central Africa
SCH	schools and universities	Stuyvesant High School
STM	streams	Withlacoochee River

BIBLIOGRAPHY

Acheson, Elise, Stefano De Sabbata, and Ross S. Purves (2017). "A quantitative analysis of global gazetteers: Patterns of coverage for common feature types." In: *Computers, Environment and Urban Systems* 64, pp. 309 –320. ISSN: 0198-9715. DOI: https://doi.org/10.1016/j.compenvurbsys.2017.03.007.

Adams, Benjamin (2015). "Finding similar places using the observation-to-generalization place model." In: *Journal of Geographical Systems* 17.2, pp. 137–156. ISSN: 1435-5949. DOI: 10.1007/s10109-015-0209-3.

Adams, Benjamin and Krzysztof Janowicz (2015). "Thematic signatures for cleansing and enriching place-related linked data." In: *International Journal of Geographical Information Science* 29.4, pp. 556–579. ISSN: 1365-8816. DOI: 10.1080/13658816.2014.989855.

Agnew, John A. (2011). "Space and Place." In: *The SAGE Handbook of Geographical Knowledge*. Ed. by John Agnew and David N. Livingstone. London: SAGE Publications Ltd, pp. 316–330. DOI: 10.4135/9781446201091.

Agrawal, Rakesh and Ramakrishnan Srikant (1994). "Fast Algorithms for Mining Association Rules in Large Databases." In: *Proceedings of the 20th International Conference on Very Large Data Bases*. VLDB '94. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., pp. 487–499. ISBN: 1-55860-153-8.

Alazzawi, Ahmed N., Alia I. Abdelmoty, and Christopher B. Jones (2012). "What can I do there? Towards the automatic discovery of place-related services and activities." In: *International Journal of Geographical Information Science* 26.2, pp. 345–364. ISSN: 1365-8816. DOI: 10.1080/13658816.2011.595954.

Alexander, Christopher (2002). *The Nature of Order: An Essay on the Art of Building and the Nature of the Universe, Book 1- The Phenomenon of Life*. Routledge.

Almuzaini, Khalid Obaid S (2017). "Qualitative Modelling of Place Location on the Linked Data Web and GIS." Thesis.

Alonso, Omar, Vasileios Kandylas, Serge-Eric Tremblay, and Stewart Whiting (2020). "Answering recreational Web searches with relevant things to do results." In: *Information Processing & Management* 57.2, p. 102184. ISSN: 0306-4573. DOI: https://doi.org/10.1016/j.ipm.2019.102184.

Aloteibi, Saad and Mark Sanderson (2014). "Analyzing geographic query reformulation: An exploratory study." In: *J. Am. Soc. Inf. Sci. Technol.* 65.1, pp. 13–24. DOI: 10.1002/asi.

Andogah, G., G. Bouma, and J. Nerbonne (2012). "Every document has a geographical scope." In: *Data & Knowledge Engineering* 81-82, pp. 1 –20. ISSN: 0169-023X. DOI: https://doi.org/10.1016/j.datak.2012.07.002.

Auer, Sören, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives (2007). "DBpedia: A Nucleus for a Web of Open Data." In: *Proceedings of the 6th International The Semantic Web and 2nd Asian Conference on Asian Semantic Web Conference*. ISWC'07/ASWC'07. Busan, Korea: Springer-Verlag, pp. 722–735. ISBN: 3-540-76297-3, 978-3-540-76297-3.

Backstrom, Lars, Jon Kleinberg, Ravi Kumar, and Jasmine Novak (2008). "Spatial Variation in Search Engine Queries." In: *Proceedings of the 17th International Conference on World Wide Web*. WWW '08. Beijing, China: ACM, pp. 357–366. ISBN: 978-1-60558-085-2. DOI: 10.1145/1367497.1367546.

Bajgar, Ondrej, Rudolf Kadlec, and Jan Kleindienst (2016). "Embracing Data Abundance: BookTest Dataset for Reading Comprehension." In: *arXiv preprint arXiv:1610.00956*.

Ballatore, Andrea (2014). "The search for places as emergent aggregates." In: *Spatial Search-Specialist meeting. Center for Spatial Studies, University of California, Santa Barbara. Santa Barbara (CA), USA*.

— (2016). "Prolegomena for an ontology of place." In: *Advancing Geographic Information Science: The Past and Next Twenty Years*. GSDI Association Press Needham, MA, pp. 91–103.

— (2019). "A Context Frame for Interactive Maps." In: *AGILE Conference on Geographical Information Science: Short Papers*, 1—5.

Behrangi, Ehsan, Hamed Ghasemzadeh, Kyumars Sheykh Esmaili, and Behrouz Minaei-Bidgoli (2007). "A Geographical Question Answering System." In: *Proceedings of the Third International Conference on Web Information Systems and Technologies*, pp. 308–314. ISBN: 978-972-8865-78-8. DOI: 10.5220/0001287203080314.

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

Bennett, David C. (1975). *Spatial and temporal uses of English prepositions: An essay in stratificational semantics*. English. Longman London, 235 pages. ISBN: 0582524539.

Bilhaut, Frédérik, Thierry Charnois, Patrice Enjalbert, and Yann Mathet (2003). "Geographic reference analysis for geographic document querying." In: *Proceedings of the*

HLT-NAACL 2003 Workshop on Analysis of Geographic References, pp. 55–62.

Blaschke, Thomas, Helena Merschdorf, Pablo Cabrera-Barona, Song Gao, Emmanuel Papadakis, and Anna Kovacs-Györi (2018). "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, p. 452. ISSN: 2220-9964. DOI: 10.3390/ijgi7110452. URL: http://www.mdpi.com/2220-9964/7/11/452.

Bordes, Antoine, Jason Weston, and Nicolas Usunier (2014a). "Open Question Answering with Weakly Supervised Embedding Models." In: *Machine Learning and Knowledge Discovery in Databases*. Ed. by Toon Calders, Floriana Esposito, Eyke Hüllermeier, and Rosa Meo. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 165–180. ISBN: 978-3-662-44848-9.

— (2014b). "Open question answering with weakly supervised embedding models." In: *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, pp. 165–180.

Bouma, Gosse (2006). "Linguistic knowledge and question answering." In: *Proceedings of the Workshop KRAQ'06: Knowledge and Reasoning for Language Processing*.

Brun, Geoffrey, Catherine Dominguès, and M-D Van Damme (2015). "TEXTOMAP: determining geographical window for texts." In: *Proceedings of the 9th Workshop on Geographic Information Retrieval*, pp. 1–2.

Buck, Christian, Jannis Bulian, Massimiliano Ciaramita, Wojciech Gajewski, Andrea Gesmundo, Neil Houlsby, and Wei Wang (2018). "Ask the Right Questions: Active Question Reformulation with Reinforcement Learning." In: *International Conference on Learning Representations*.

Buscaldi, Davide (2011). "Approaches to Disambiguating Toponyms." In: *SIGSPATIAL Special* 3.2, 16–19. DOI: 10.1145/2047296.2047300.

Buscaldi, Davide, Yassine Benajiba, Paolo Rosso, and Emilio Sanchis (2006). "The UPV at QA@ CLEF 2006." In: *CLEF (Working Notes)*.

Buttimer, Anne (1976). "Grasping the Dynamism of Lifeworld." In: *Annals of the Association of American Geographers* 66.2, pp. 277–292. ISSN: 0004-5608. DOI: 10.1111/j.1467-8306.1976.tb01090.x.

Calafiore, Alessia (2016). "Designing an Ontology of Social Place." In: *DC@AI*IA*.

Caliński, Tadeusz and Jerzy Harabasz (1974). "A dendrite method for cluster analysis." In: *Communications in Statistics* 3.1, pp. 1–27. DOI: 10.1080/03610927408827101.

Canter, David (1977). *The Psychology of Place*. Architectural Press.

Canter, David (1997). "The Facets of Place." In: *Toward the Integration of Theory, Methods, Research, and Utilization*. Ed. by Gary T. Moore and Robert W. Marans. Boston, MA: Springer US, pp. 109–147. ISBN: 978-1-4757-4425-5. DOI: 10.1007/978-1-4757-4425-5_4.

— (2012). *Facet theory: Approaches to social research*. Springer Science & Business Media. ISBN: 1461250420.

Capineri, Cristina (2016). "Kilburn High Road Revisited." In: *Urban Planning* 1.2, p. 13. ISSN: 2183-7635. DOI: 10.17645/up.v1i2.614.

Carmona, Matthew, Tim Heath, Taner Oc, and Steve Tiesdell (2012). *Public places-Urban spaces*. Routledge. ISBN: 1136020500.

Casey, Edward (1997). *The Fate of Place: A Philosophical History*. University of California Press.

Cheema, Muhammad Aamir, Wenjie Zhang, and Lijun Chang (2016). *Databases Theory and Applications: 27th Australasian Database Conference, ADC 2016, Sydney, NSW, September 28-29, 2016, Proceedings*. Vol. 9877. Springer.

Chen, Hao, Maria Vasardani, Stephan Winter, and Martin Tomko (2018). "A Graph Database Model for Knowledge Extracted from Place Descriptions." In: *ISPRS International Journal of Geo-Information* 7.6, p. 221. ISSN: 2220-9964. DOI: 10.3390/ijgi7060221.

Chen, W. (2014). "Parameterized Spatial SQL Translation for Geographic Question Answering." In: *2014 IEEE International Conference on Semantic Computing*, pp. 23–27.

Chen, W., E. Fosler-Lussier, N. Xiao, S. Raje, R. Ramnath, and D. Sui (2013). "A Synergistic Framework for Geographic Question Answering." In: *Proceedings of IEEE 7th International Conference on Semantic Computing*, pp. 94–99.

Chen, Wenhu, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyou Zhou, and William Yang Wang (2019). "TabFact: A Large-scale Dataset for Table-based Fact Verification." In: *International Conference on Learning Representations*.

Chen, Wenhu, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, and William Wang (2020). "HybridQA: A Dataset of Multi-Hop Question Answering over Tabular and Textual Data." In: *arXiv preprint arXiv:2004.07347*.

Chesnokova, Olga, Joanna E. Taylor, Ian N. Gregory, and Ross Purves (2018). "Hearing the silence: finding the middle ground in the spatial humanities? Extracting and comparing perceived silence and tranquillity in the English Lake District." In: *International Journal of Geographical Information Science*, pp. 1–25. ISSN: 1365-8816. DOI: 10.1080/13658816.2018.1552789.

Chierchia, Gennaro (2006). "Broaden your views: Implicatures of domain widening and the "logicality" of language." In: *Linguistic Inquiry* 37.4, pp. 535–590.

Chung, Hoojung, Young-In Song, Kyoung-Soo Han, Do-Sang Yoon, Joo-Young Lee, Hae Chang Rim, and Soo-Hong Kim (2004). "A practical QA system in restricted domains." In: *Proceedings of the Conference on Question Answering in Restricted Domains*, pp. 39–45.

Church, Karen, Joachim Neumann, Mauro Cherubini, and Nuria Oliver (2010). "The "Map Trap"? An Evaluation of Map versus Text-Based Interfaces for Location-Based Mobile Search Services." In: *Proceedings of the 19th International Conference on World Wide Web*. WWW '10. Raleigh, North Carolina, USA: Association for Computing Machinery, 261–270. ISBN: 9781605587998. DOI: 10.1145/1772690.1772718.

Cleary, John and Ian Witten (1984). "Data compression using adaptive coding and partial string matching." In: *IEEE transactions on Communications* 32.4, pp. 396–402.

Cohen, Daniel and W Bruce Croft (2018). "A hybrid embedding approach to noisy answer passage retrieval." In: *European Conference on Information Retrieval*. Springer, pp. 127–140.

Cohen, Daniel, Liu Yang, and W. Bruce Croft (2018a). "WikiPassageQA: A Benchmark Collection for Research on Non-Factoid Answer Passage Retrieval." In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*. SIGIR '18. Ann Arbor, MI, USA: Association for Computing Machinery, 1165–1168. ISBN: 9781450356572. DOI: 10.1145/3209978.3210118.

Cohen, Daniel, Liu Yang, and W Bruce Croft (2018b). "WikiPassageQA: A benchmark collection for research on non-factoid answer passage retrieval." In: *The 41st International ACM SIGIR Conference on Research & Development in Information Retrieval*, pp. 1165–1168.

Coleman, Simon and Peter Collins (2006). *Locating the field: space, place and context in anthropology*. Berg.

Couclelis, H., R.G. Golledge, N. Gale, and W. Tobler (1987). "Exploring the anchor-point hypothesis of spatial cognition." In: *Journal of Environmental Psychology* 7.2, pp. 99 –122. ISSN: 0272-4944. DOI: https://doi.org/10.1016/S0272-4944(87)80020-8.

Couclelis, Helen (1992). "Location, place, region, and space." In: *Geography's inner worlds*. Vol. 2. Rutgers University Press, New Jersey, pp. 15–233.

— (2010). "Ontologies of geographic information." In: *International Journal of Geographical Information Science* 24.12, pp. 1785–1809. ISSN: 1365-8816. DOI: 10.1080/13658816.2010.484392.

Cresswell, Tim (2004). *Place: a short introduction*. John Wiley & Sons. ISBN: 0470655623.

Cresswell, Tim (2014). *Place : An Introduction*. Hoboken, United Kingdom: John Wiley & Sons, Incorporated. ISBN: 9781118574164.

Cummins, Steven, Sarah Curtis, Ana V Diez-Roux, and Sally Macintyre (2007). "Understanding and representing 'place' in health research: A relational approach." In: *Social Science & Medicine* 65.9, pp. 1825–1838. ISSN: 0277-9536. DOI: https://doi.org/10.1016/j.socscimed.2007.05.036.

Curran, James R and Stephen Clark (2003). "Language independent NER using a maximum entropy tagger." In: *Proceedings of the Seventh Conference on Natural Language Learning at HLT-NAACL 2003*, pp. 164–167.

Das, Rahul Deb and Stephan Winter (2016). "A context-sensitive conceptual framework for activity modeling." In: *Journal of Spatial Information Science* 12.1, pp. 45–85. DOI: 10.5311/JOSIS.2016.12.260.

De Marneffe, Marie-Catherine, Timothy Dozat, Natalia Silveira, Katri Haverinen, Filip Ginter, Joakim Nivre, and Christopher D Manning (2014). "Universal Stanford dependencies: A cross-linguistic typology." In: *Proceedings of the Ninth International Conference on Language Resources and Evaluation*. Vol. 14, pp. 4585–4592.

Delmonte, Rodolfo (2007). *Computational Linguistic Text Processing: Logical Form, Semantic Interpretation, Discourse Relations and Question Answering*. Nova Publishers.

Derungs, Curdin and Ross S. Purves (2014). "From text to landscape: locating, identifying and mapping the use of landscape features in a Swiss Alpine corpus." In: *International Journal of Geographical Information Science* 28.6, pp. 1272–1293. ISSN: 1365-8816. DOI: 10.1080/13658816.2013.772184.

Di Gennaro, Giovanni, Amedeo Buonanno, Antonio Di Girolamo, and Francesco AN Palmieri (2020). "Intent Classification in Question-Answering Using LSTM Architectures." In: *arXiv preprint arXiv:2001.09330*.

Diefenbach, Dennis, Vanessa Lopez, Kamal Singh, and Pierre Maret (2018). "Core techniques of question answering systems over knowledge bases: a survey." In: *Knowledge and Information systems* 55.3, pp. 529–569.

Dimitrakis, Eleftherios, Konstantinos Sgontzos, and Yannis Tzitzikas (2019). "A survey on question answering systems over linked data and documents." In: *Journal of Intelligent Information Systems*, pp. 1–27.

Dong, Li and Mirella Lapata (2018). "Coarse-to-Fine Decoding for Neural Semantic Parsing." In: *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 731–742.

Du Bois, John W. (2014). "Towards a dialogic syntax." In: *Cognitive Linguistics* 25.3, pp. 359–410. ISSN: 16133641. DOI: 10.1515/cog-2014-0024. URL: https://www.degruyter.com/view/journals/cogl/25/3/article-p359.xml.

Dubey, Mohnish, Debayan Banerjee, Abdelrahman Abdelkawi, and Jens Lehmann (2019). "C-quAD 2.0: A large dataset for complex question answering over wikidata and dbpedia." In: *International Semantic Web Conference*. Springer, pp. 69–78.

Dunn, Stuart (2017). "Praxes of "The Human" and "The Digital": Spatial Humanities and the Digitization of Place." In: *GeoHumanities* 3.1, pp. 88–107. DOI: 10.1080/2373566X.2016.1245107.

Dwivedi, Sanjay K. and Vaishali Singh (2013). "Research and Reviews in Question Answering System." In: *Procedia Technology* 10, pp. 417–424. ISSN: 22120173. DOI: 10.1016/j.protcy.2013.12.378.

Edwardes, Alistair J (2007). "Re-placing location: Geographic perspectives in location based services." Thesis. Faculty of Science.

Edwardes, Alistair J. and Ross Purves (2007a). "A Theoretical Grounding for Semantic Descriptions of Place." In: *Web and Wireless Geographical Information Systems*. Ed. by J. Mark Ware and George E. Taylor. Springer Berlin Heidelberg, pp. 106–120. ISBN: 978-3-540-76925-5.

Edwardes, Alistair J and Ross Purves (2007b). "Eliciting concepts of place for text-based image retrieval." In: *Proceedings of the 4th ACM workshop on Geographical Information Retrieval*. ACM, pp. 15–18. ISBN: 1595938281.

Egenhofer, Max J (1991). "Reasoning about binary topological relations." In: *Symposium on Spatial Databases*. Springer, pp. 141–160.

ElGindy, Ehab and Alia Abdelmoty (2014a). "Capturing Place Semantics on the GeoSocial Web." In: *Journal on Data Semantics* 3.4, pp. 207–223. ISSN: 1861-2040. DOI: 10.1007/s13740-014-0034-8.

— (2014b). "Enriching user profiles using geo-social place semantics in geo-folksonomies." In: *International Journal of Geographical Information Science* 28.7, pp. 1439–1458. ISSN: 1365-8816. DOI: 10.1080/13658816.2014.894194.

Entrikin, J. Nicholas (1991). "The Betweenness of Place." In: *The Betweenness of Place: Towards a Geography of Modernity*. London: Macmillan Education UK, pp. 6–26. ISBN: 978-1-349-21086-2. DOI: 10.1007/978-1-349-21086-2_2.

Entrikin, J Nicholas (1994). "Place and Region." In: *Progress in Human Geography* 18.2, pp. 227–233. ISSN: 0309-1325.

Entrikin, J Nicholas (1996). "Place and Region 2." In: *Progress in Human Geography* 20.2, pp. 215–221. ISSN: 0309-1325.

— (1997). "Place and Region 3." In: *Progress in Human Geography* 21.2, pp. 263–268. ISSN: 0309-1325.

Ferrés Domènech, Daniel (2017). "Knowledge-based and data-driven approaches for geographical information access." PhD thesis.

Ferrés, Daniel and Horacio Rodríguez (2006). "Experiments Adapting an Open-domain Question Answering System to the Geographical Domain Using Scope-based Resources." In: *Proceedings of the Workshop on Multilingual Question Answering*. MLQA '06. Stroudsburg, PA, USA: Association for Computational Linguistics, pp. 69–76. ISBN: 2-9524532-4-1.

Ferrés, Daniel and Horacio Rodríguez (2010). "TALP at GikiCLEF 2009." In: *Multilingual Information Access Evaluation I. Text Retrieval Experiments*. Ed. by Carol Peters, Giorgio Maria Di Nunzio, Mikko Kurimo, Thomas Mandl, Djamel Mostefa, Anselmo Peñas, and Giovanna Roda. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 322–325. ISBN: 978-3-642-15754-7.

Finkel, Jenny Rose, Trond Grenager, and Christopher D Manning (2005). "Incorporating Non-local Information into Information Extraction Systems by Gibbs Sampling." In: *Proceedings of the 43rd Annual Meeting of the Association for Computational Linguistics (ACL'05)*, pp. 363–370.

Fournier-Viger, Philippe, Jerry Chun-Wei Lin, Antonio Gomariz, Ted Gueniche, Azadeh Soltani, Zhihong Deng, and Hoang Thanh Lam (2016). "The SPMF Open-Source Data Mining Library Version 2." In: *Machine Learning and Knowledge Discovery in Databases*. Cham: Springer International Publishing, pp. 36–40. ISBN: 978-3-319-46131-1.

Frank, Andrew U (1992). "Qualitative spatial reasoning about distances and directions in geographic space." In: *Journal of Visual Languages & Computing* 3.4, pp. 343–371. ISSN: 1045-926X.

Galton, Antony (2010). "The formalities of affordance." In: *Proceedings of ECAI2010 Workshop on Spatio-Temporal Dynamics*, pp. 1–6.

Gan, Qingqing, Josh Attenberg, Alexander Markowetz, and Torsten Suel (2008). "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: ACM, pp. 49–56. ISBN: 978-1-60558-160-6. DOI: 10.1145/1367798.1367806.

Gazdar, Gerald (1977). *Implicature, presupposition and logical form*. Indiana University Linguistics Club.

Gibson, James J (1979). *The Ecological Approach to Visual Perception*. Boston: Houghton Mifflin Company.

Gieryn, Thomas F. (2000). "A Space for Place in Sociology." In: *Annual Review of Sociology* 26.1, pp. 463–496. ISSN: 0360-0572. DOI: 10.1146/annurev.soc.26.1.463.

Goldberg, Daniel W., John P. Wilson, and Craig A. Knoblock (2009). "Extracting geographic features from the Internet to automatically build detailed regional gazetteers." In: *International Journal of Geographical Information Science* 23.1, pp. 93–128. ISSN: 1365-8816. DOI: 10.1080/13658810802577262.

Goodchild, Michael F. (2011). "Formalizing Place in Geographic Information Systems." In: *Communities, Neighborhoods, and Health: Expanding the Boundaries of Place*. New York, NY: Springer New York, pp. 21–33. ISBN: 978-1-4419-7482-2. DOI: 10.1007/978-1-4419-7482-2_2.

Goodchild, Michael and Linna Li (2012). "Formalizing space and place." In: *Fonder les Sciences du Territoire*. Paris: Éditions Karthala, pp. 83–94.

Graham, Mark and Stefano De Sabbata (2015). "Mapping information wealth and poverty: the geography of gazetteers." In: *Environment and Planning A* 47.6, pp. 1254–1264.

Grice, H. P. (1975). "Logic and Conversation." In: Leiden, The Netherlands: Brill, pp. 41 –58. ISBN: 9789004368576. DOI: https://doi.org/10.1163/9789004368811_003.

Groenendijk, Jeroen and Martin Stokhof (1980). "A pragmatic analysis of specificity." In: *Ambiguities in Intensional Contexts*. Springer, pp. 153–190.

Gueniche, Ted, Philippe Fournier-Viger, Rajeev Raman, and Vincent S Tseng (2015). "CPT+: Decreasing the time/space complexity of the Compact Prediction Tree." In: *Pacific-Asia Conference on Knowledge Discovery and Data Mining*. Springer, pp. 625–636.

Gueniche, Ted, Philippe Fournier-Viger, and Vincent S Tseng (2013). "Compact prediction tree: A lossless model for accurate sequence prediction." In: *International Conference on Advanced Data Mining and Applications*. Springer, pp. 177–188.

Guerzoni, Elena (2003). "Why even ask?: on the pragmatics of questions and the semantics of answers." PhD thesis. Massachusetts Institute of Technology.

Gupta, Poonam and Vishal Gupta (2012). "A survey of text question answering techniques." In: *International Journal of Computer Applications* 53.4.

Gustafson, Per (2001). "Meanings of place: Everyday experience and theoretical conceptualizations." In: *Journal of Environmental Psychology* 21.1, pp. 5–16. ISSN: 0272-4944. DOI: https://doi.org/10.1006/jevp.2000.0185.

Guy, Ido (2016). "Searching by Talking: Analysis of Voice Queries on Mobile Web Search." In: *Proceedings of the 39th International*

ACM SIGIR Conference on Research and Development in Information Retrieval. SIGIR '16. Pisa, Italy: ACM, pp. 35–44. ISBN: 978-1-4503-4069-4. DOI: 10.1145/2911451.2911525.

Guy, Ido (Mar. 2018). "The Characteristics of Voice Search: Comparing Spoken with Typed-in Mobile Web Search Queries." In: *ACM Transactions on Information Systems (TOIS)* 36.3, 30:1–30:28. ISSN: 1046-8188. DOI: 10.1145/3182163.

Hales, Thomas C. (2007). "The Jordan Curve Theorem, Formally and Informally." In: *The American Mathematical Monthly* 114.10, pp. 882–894. ISSN: 00029890, 19300972.

Hamzei, Ehsan and Farshad Hakimpour (2017). "Entity recognition and disambiguation for natural-language spatial search queries." In: *2017 3th International Conference on Web Research (ICWR)*. IEEE, pp. 32–37.

Hamzei, Ehsan, Haonan Li, Maria Vasardani, Timothy Baldwin, Stephan Winter, and Martin Tomko (2020). "Place Questions and Human-Generated Answers: A Data Analysis Approach." In: *Geospatial Technologies for Local and Regional Development*. Ed. by Phaedon Kyriakidis, Diofantos Hadjimitsis, Dimitrios Skarlatos, and Ali Mansourian. Cham: Springer International Publishing, pp. 3–19. ISBN: 978-3-030-14745-7. DOI: https://doi.org/10.1007/978-3-030-14745-7_1.

Hamzei, Ehsan, Stephan Winter, and Martin Tomko (2019). "Initial Analysis of Simple Where-Questions and Human-Generated Answers (Short Paper)." In: *14th International Conference on Spatial Information Theory (COSIT 2019)*. Ed. by Sabine Timpf, Christoph Schlieder, Markus Kattenbeck, Bernd Ludwig, and Kathleen Stewart. Vol. 142. Leibniz International Proceedings in Informatics (LIPIcs). Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 12:1–12:8. ISBN: 978-3-95977-115-3. DOI: 10.4230/LIPIcs.COSIT.2019.12.

— (2020). "Place facets: a systematic literature review." In: *Spatial Cognition & Computation* 20.1, pp. 33–81. DOI: 10.1080/13875868.2019.1688332.

— (2021). "Templates of generic geographic information for answering where-questions." In: *International Journal of Geographical Information Science*, pp. 1–27. DOI: 10.1080/13658816.2020.1869977.

Harrison, Steve and Paul Dourish (1996). "Re-place-ing Space: The Roles of Place and Space in Collaborative Systems." In: *Proceedings of the 1996 ACM Conference on Computer Supported Cooperative Work*. CSCW '96. Boston, Massachusetts, USA: ACM, pp. 67–76. ISBN: 0-89791-765-0. DOI: 10.1145/240080.240193.

Hartrumpf, Sven and Johannes Leveling (2010). "Recursive Question Decomposition for Answering Complex Geographic Questions." In: *Multilingual Information Access Evaluation I. Text Retrieval*

Experiments. Ed. by Carol Peters, Giorgio Maria Di Nunzio, Mikko Kurimo, Thomas Mandl, Djamel Mostefa, Anselmo Peñas, and Giovanna Roda. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 310–317. ISBN: 978-3-642-15754-7.

Henrich, Andreas and Volker Luedecke (2007). "Characteristics of Geographic Information Needs." In: *Proceedings of the 4th ACM Workshop on Geographical Information Retrieval*. New York, NY, USA: ACM, pp. 1–6.

Hermann, Karl Moritz, Tomas Kocisky, Edward Grefenstette, Lasse Espeholt, Will Kay, Mustafa Suleyman, and Phil Blunsom (2015). "Teaching machines to read and comprehend." In: *Advances in Neural Information Processing Systems*, pp. 1693–1701.

Hermjakob, Ulf (2001). "Parsing and Question Classification for Question Answering." In: *Proceedings of the Workshop on Open-Domain Question Answering - Volume 12*. ODQA '01. Toulouse, France: Association for Computational Linguistics, 1–6. DOI: 10.3115/1117856.1117859.

Hill, Felix, Antoine Bordes, Sumit Chopra, and Jason Weston (2015). "The Goldilocks principle: Reading children's books with explicit memory representations." In: *arXiv preprint arXiv:1511.02301*.

Hill, Linda, James Frew, and Qi Zheng (1999). *Geographic Names: The Implementation of a Gazetteer in a Georeferenced Digital Library*. Tech. rep.

Hillier, Bill (2007). *Space is the machine: a configurational theory of architecture*. Space Syntax.

Hockenberry, Matthew Curtis (2006). "Grounding for a computational model of place." Thesis. School of Architecture and Planning.

Hollenstein, Livia and Ross Purves (2010). "Exploring place through user-generated content: Using Flickr tags to describe city cores." In: *Journal of Spatial Information Science* 1.1, pp. 21–48. DOI: 10.5311/JOSIS.2010.1.3.

Huilin, G., G. Tong, W. Fan, and M. Chao (2019). "Bidirectional Attention for SQL Generation." In: *2019 IEEE 4th International Conference on Cloud Computing and Big Data Analysis (ICCCBDA)*, pp. 676–682.

Ingram, D. R. (1971). "The concept of accessibility: A search for an operational form." In: *Regional Studies* 5.2, pp. 101–107. ISSN: 0034-3404. DOI: 10.1080/09595237100185131.

Jang, Gwan, Keun-Chan Park, Kyung-min Kim, Yoonjae Jeong, and Sung-Hyon Myaeng (2010). "What Aspects Do People Search in Geo-Referenced Text?" In: *Proceedings of the 3rd International Workshop on Location and the Web*. LocWeb '10. Tokyo, Japan: Association for Computing Machinery. ISBN: 9781450304122. DOI: 10.1145/1899662.1899666.

Janowicz, Krzysztof (2012). "Observation-driven geo-ontology engineering." In: *Transactions in GIS* 16.3, pp. 351–374. ISSN: 1361-1682.

Janowicz, Krzysztof, Simon Scheider, and Benjamin Adams (2013). "A Geo-semantics Flyby." In: *Reasoning Web. Semantic Technologies for Intelligent Data Access: 9th International Summer School 2013, Mannheim, Germany, July 30 – August 2, 2013. Proceedings*. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 230–250. ISBN: 978-3-642-39784-4. DOI: 10.1007/978-3-642-39784-4_6.

Jaro, Matthew A. (1989). "Advances in Record-Linkage Methodology as Applied to Matching the 1985 Census of Tampa, Florida." In: *Journal of the American Statistical Association* 84.406, pp. 414–420. DOI: 10.1080/01621459.1989.10478785.

Jenks, G. F. (1967). "The Data Model Concept in Statistical Mapping." In: *International Yearbook of Cartography* 7, pp. 186–190.

Jiang, Bin and Xiaobai Yao (2006). "Location-based services and GIS in perspective." In: *Computers, Environment and Urban Systems* 30.6. Location Based Services, pp. 712 –725. ISSN: 0198-9715. DOI: https://doi.org/10.1016/j.compenvurbsys.2006.02.003.

Johnson, Stephen C (1967). "Hierarchical clustering schemes." In: *Psychometrika* 32.3, pp. 241–254. ISSN: 0033-3123.

Jones, Christopher B., Harith Alani, and Douglas Tudhope (2001). "Geographical Information Retrieval with Ontologies of Place." In: *Spatial Information Theory*. Ed. by Daniel R. Montello. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 322–335. ISBN: 978-3-540-45424-3.

Jones, Rosie, Wei V. Zhang, Benjamin Rey, Pradhuman Jhala, and Eugene Stipp (2008). "Geographic intention and modification in web search." In: *International Journal of Geographical Information Science* 22.3, pp. 229–246. ISSN: 1365-8816. DOI: 10.1080/13658810701626186.

Jones, Stephen B. (1959). "BOUNDARY CONCEPTS IN THE SETTING OF PLACE AND TIME." In: *Annals of the Association of American Geographers* 49.3, pp. 241–255. ISSN: 0004-5608. DOI: 10.1111/j.1467-8306.1959.tb01611.x.

Jonietz, David (2016). "From Space to Place: A Computational Model of Functional Place." PhD Thesis. Faculty of Applied Computer Science.

Jonietz, David and Sabine Timpf (2015). "On the relevance of Gibson's affordance concept for geographical information science (GISc)." In: *Cognitive Processing* 16.1, pp. 265–269. ISSN: 1612-4790. DOI: 10.1007/s10339-015-0719-4.

Jordan, Troy, Martin Raubal, Bryce Gartrell, and M Egenhofer (1998). "An affordance-based model of place in GIS." In: *8th International Symposium on Spatial Data Handling, SDH*. Vol. 98, pp. 98–109.

Jorgensen, Bradley S. and Richard C. Stedman (2001). "SENSE OF PLACE AS AN ATTITUDE: LAKESHORE OWNERS ATTITUDES TOWARD THEIR PROPERTIES." In: *Journal of Environmental Psychology* 21.3, pp. 233–248. ISSN: 0272-4944. DOI: https://doi.org/10.1006/jevp.2001.0226.

Karttunen, Lauri (1977). "Syntax and semantics of questions." In: *Linguistics and Philosophy* 1.1, pp. 3–44. ISSN: 01650157. DOI: 10.1007/BF00351935.

Khan, Arbaz, Maria Vasardani, and Stephan Winter (2013). "Extracting Spatial Information From Place Descriptions." In: *Proceedings of The First ACM SIGSPATIAL International Workshop on Computational Models of Place*. COMP '13. Orlando FL, USA: Association for Computing Machinery, 62–69. ISBN: 9781450325356. DOI: 10.1145/2534848.2534857.

Kim, Hyeonji, Byeong-Hoon So, Wook-Shin Han, and Hongrae Lee (2020). "Natural language to SQL: Where are we today?" In: *Proceedings of the VLDB Endowment* 13.10, pp. 1737–1750.

Kim, Junchul, Maria Vasardani, and Stephan Winter (2016). "From descriptions to depictions: A dynamic sketch map drawing strategy." In: *Spatial Cognition & Computation* 16.1, pp. 29–53. ISSN: 1387-5868. DOI: 10.1080/13875868.2015.1084509.

— (2017). "Similarity matching for integrating spatial information extracted from place descriptions." In: *International Journal of Geographical Information Science* 31.1, pp. 56–80. ISSN: 1365-8816. DOI: 10.1080/13658816.2016.1188930.

Kordjamshidi, Parisa, Martijn Van Otterlo, and Marie-Francine Moens (2011). "Spatial role labeling: Towards extraction of spatial relations from natural language." In: *ACM Transactions on Speech and Language Processing (TSLP)* 8.3, pp. 1–36. ISSN: 1550-4875. DOI: 10.1145/2050104.2050105.

Kravi, Elad, Eugene Agichtein, Ido Guy, Yaron Kanza, Avihai Mejer, and Dan Pelleg (2015). "Searcher in a Strange Land: Understanding Web Search from Familiar and Unfamiliar Locations." In: *The 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM, pp. 855–858. DOI: 10.1145/2766462.2767782.

Krifka, Manfred (1996). "Pragmatic strengthening in plural predications and donkey sentences." In: *Semantics and Linguistic Theory*. Vol. 6, pp. 136–153.

Kuhn, Werner (2001). "Ontologies in support of activities in geographical space." In: *International Journal of Geographical Information Science* 15.7, pp. 613–631. ISSN: 1365-8816. DOI: 10.1080/13658810110061180.

— (2012). "Core concepts of spatial information for transdisciplinary research." In: *International Journal of Geographical Information Science* 26.12, pp. 2267–2276.

Kuhn, Werner and Andrea Ballatore (2015). "Designing a Language for Spatial Computing." In: *AGILE 2015: Geographic Information Science as an Enabler of Smarter Cities and Communities*. Ed. by Fernando Bacao, Maribel Yasmina Santos, and Marco Painho. Cham: Springer International Publishing, pp. 309–326. ISBN: 978-3-319-16787-9. DOI: 10.1007/978-3-319-16787-9_18.

Kuhn, Werner, Ehsan Hamzei, Martin Tomko, Maria Vasardani, Stephan Winter, and Haonan Li (2021, Under Review). "Neural Geospatial Question Answering." In:

Kyle, Gerard, Kelly Bricker, Alan Graefe, and Thomas Wickham (2004). "An Examination of Recreationists' Relationships with Activities and Settings." In: *Leisure Sciences* 26.2, pp. 123–142. ISSN: 0149-0400. DOI: 10.1080/01490400490432019.

Kyle, Gerard, Alan Graefe, and Robert Manning (2005). "Testing the Dimensionality of Place Attachment in Recreational Settings." In: *Environment and Behavior* 37.2, pp. 153–177. ISSN: 0013-9165. DOI: 10.1177/0013916504269654.

Laird, Philip and Ronald Saul (1994). "Discrete sequence prediction and its applications." In: *Machine Learning* 15.1, pp. 43–68.

Lakshmi, K and K.M. Arivuchelvan (2019). "A Survey on Datasets for Machine Reading Comprehension." In: *SSRN Electronic Journal*. DOI: 10.2139/ssrn.3429211.

Lample, Guillaume, Miguel Ballesteros, Sandeep Subramanian, Kazuya Kawakami, and Chris Dyer (2016). "Neural Architectures for Named Entity Recognition." In: *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 260–270.

Lee, Chia-Jung, Nick Craswell, and Vanessa Murdock (2015). "Inter-Category Variation in Location Search." In: *The 38th International ACM SIGIR Conference on Research and Development in Information Retrieval*. 2767797: ACM, pp. 863–866. DOI: 10.1145/2766462.2767797.

Leidner, Jochen L (2007). "Toponym resolution in text: annotation, evaluation and applications of spatial grounding." In: *ACM SIGIR Forum*. Vol. 41. ACM, pp. 124–126. ISBN: 0163-5840.

Leidner, Jochen L. and Michael D. Lieberman (July 2011). "Detecting Geographical References in the Form of Place Names and Associated Spatial Natural Language." In: *SIGSPATIAL Special* 3.2, 5–11. DOI: 10.1145/2047296.2047298. URL: https://doi.org/10.1145/2047296.2047298.

Leidner, Jochen L., Gail Sinclair, and Bonnie Webber (2003). "Grounding Spatial Named Entities for Information Extraction and Question Answering." In: *Proceedings of the HLT-NAACL 2003 Workshop on Analysis of Geographic References - Volume 1*. HLT-NAACL-GEOREF '03. Stroudsburg, PA, USA: Association

for Computational Linguistics, pp. 31–38. DOI: 10.3115/1119394.1119399.

Lewicka, Maria (2010). "What makes neighborhood different from home and city? Effects of place scale on place attachment." In: *Journal of Environmental Psychology* 30.1, pp. 35–51. ISSN: 0272-4944. DOI: https://doi.org/10.1016/j.jenvp.2009.05.004.

Li, Haonan, Ehsan Hamzei, Ivan Majic, Hua Hua, Jochen Renz, Martin Tomko, Maria Vasardani, Stephan Winter, and Timothy Baldwin (2021, Under Review). "Neural Geospatial Question Answering." In:

Li, Xiaoya, Fan Yin, Zijun Sun, Xiayu Li, Arianna Yuan, Duo Chai, Mingxin Zhou, and Jiwei Li (2019). "Entity-Relation Extraction as Multi-Turn Question Answering." In: *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pp. 1340–1350.

Li, Xin and Dan Roth (2002). "Learning Question Classifiers." In: *Proceedings of the 19th International Conference on Computational Linguistics - Volume 1*. COLING '02. Taipei, Taiwan: Association for Computational Linguistics, 1–7. DOI: 10.3115/1072228.1072378.

— (2006). "Learning question classifiers: The role of semantic information." In: *Natural Language Engineering* 12.3, pp. 229–249.

Lieberman, Michael D. and Hanan Samet (2012). "Adaptive Context Features for Toponym Resolution in Streaming News." In: *Proceedings of the 35th International ACM SIGIR Conference on Research and Development in Information Retrieval*. SIGIR '12. Portland, Oregon, USA: ACM, pp. 731–740. ISBN: 978-1-4503-1472-5. DOI: 10.1145/2348283.2348381.

Luque, Jordi, Daniel Ferrés, Javier Hernando, José B Mariño, and Horacio Rodríguez (2006). "GEOVAQA: A VOICE ACTIVATED GEOGRAPHICAL QUESTION ANSWERING SYSTEM." In: *IV Jornadas en Tecnologia del Habla*. Zaragoza, Spain, pp. 309–314.

Lussault, Michel (2003). "Valeur Spatiale." In: *Dictionnaire de géographie et de l'espace des sociétés*. La Documentation Française, pp. 973–974.

MacRoberts, Michael H. and Barbara R. MacRoberts (1989). "Problems of Citation Analysis: A Critical Review." In: *Journal of the American Society for Information Science* 40.5, pp. 342–349. ISSN: 00028231.

Madabushi, Harish Tayyar and Mark Lee (2016). "High accuracy rule-based question classification using question syntax and semantics." In: *Proceedings of COLING 2016, the 26th International Conference on Computational Linguistics: Technical Papers*, pp. 1220–1230.

Maechler, Martin (2018). ""Finding Groups in Data": Cluster Analysis Extended Rousseeuw et. al." In: *Documentation for software package. The Comprehensive R Archive Network (CRAN): Wien*.

Mai, Gengchen, Krzysztof Janowicz, Cheng He, Sumang Liu, and Ni Lao (2018). "POIReviewQA: A Semantically Enriched POI Retrieval and Question Answering Dataset." In: *Proceedings of the 12th Workshop on Geographic Information Retrieval*. GIR'18. Seattle, WA, USA: Association for Computing Machinery. ISBN: 9781450360340. DOI: 10.1145/3281354.3281359.

Mai, Gengchen, Bo Yan, Krzysztof Janowicz, and Rui Zhu (2020). "Relaxing Unanswerable Geographic Questions Using A Spatially Explicit Knowledge Graph Embedding Model." In: *Geospatial Technologies for Local and Regional Development*. Ed. by Phaedon Kyriakidis, Diofantos Hadjimitsis, Dimitrios Skarlatos, and Ali Mansourian. Cham: Springer International Publishing, pp. 21–39. ISBN: 978-3-030-14745-7.

Malamud, Sophia A (2006). "(Non) maximality and Distributivity: A Decision Theory Approach." In: *Semantics and Linguistic Theory*. Vol. 16, pp. 120–137.

Manning, Christopher D., Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J. Bethard, and David McClosky (2014). "The Stanford CoreNLP Natural Language Processing Toolkit." In: *Association for Computational Linguistics (ACL) 2014 System Demonstrations*, pp. 55–60.

Markus, T. A. (1987). "Buildings as Classifying Devices." In: *Environment and Planning B: Planning and Design* 14.4, pp. 467–484. ISSN: 0265-8135. DOI: 10.1068/b140467.

Marx, Edgard, Ricardo Usbeck, Axel-Cyrille Ngonga Ngomo, Konrad Höffner, Jens Lehmann, and Sören Auer (2014). "Towards an Open Question Answering Architecture." In: *Proceedings of the 10th International Conference on Semantic Systems*. SEM '14. Leipzig, Germany: ACM, pp. 57–60. ISBN: 978-1-4503-2927-9. DOI: 10.1145/2660517.2660519.

Massey, Doreen (1994). *Space, Place, and Gender*. Vol. 35. University of Minnesota Press., pp. 24–29.

McClure, Charles R, R David Lankes, Melissa Gross, and Beverly Choltco-Devlin (2002). "Statistics, Measures, and Quality Standards for Assessing Digital Reference Library Services: Guidelines and Procedures." In:

McKenzie, Grant and Benjamin Adams (2017). "Juxtaposing thematic regions derived from spatial and platial user-generated content." In: *Proceedings of the 13th International Conference on Spatial Information Theory*. Vol. 86. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik. ISBN: 395977043X.

Mendes, Ana Cristina and Luisa Coheur (2013). "When the answer comes into question in question-answering: survey and open issues." In: *Natural Language Engineering* 19.1, pp. 1–32.

Mennis, Jeremy and Michael J Mason (2016). "Modeling Place as a Relationship between a Person and a Location." In: *International Conference on GIScience*. Montreal, Canada.

Merschdorf, Helena and Thomas Blaschke (2018). "Revisiting the Role of Place in Geographic Information Science." In: *ISPRS International Journal of Geo-Information* 7.9, p. 364. ISSN: 2220-9964.

Mishra, A., N. Mishra, and A. Agrawal (2010). "Context-Aware Restricted Geographical Domain Question Answering System." In: *2010 International Conference on Computational Intelligence and Communication Networks*, pp. 548–553.

Mondschein, Andrew, Evelyn Blumenberg, and Brian Taylor (2010). "Accessibility and cognition: the effect of transport mode on spatial knowledge." In: *Urban Studies* 47.4, pp. 845–866. ISSN: 0042-0980.

Monteiro, Bruno R, Clodoveu A Davis Jr, and Fred Fonseca (2016). "A survey on the geographic scope of textual documents." In: *Computers & Geosciences* 96, pp. 23–34.

Montello, Daniel R (1993). "Scale and multiple psychologies of space." In: *European Conference on Spatial Information Theory*. Springer, pp. 312–321.

Montello, Daniel R., Michael F. Goodchild, Jonathon Gottsegen, and Peter Fohl (2003). "Where's Downtown?: Behavioral Methods for Determining Referents of Vague Spatial Queries." In: *Spatial Cognition & Computation* 3.2-3, pp. 185–204. ISSN: 1387-5868. DOI: 10.1080/13875868.2003.9683761.

Mooney, Martha T (1991). *Sears list of subject headings*. New York: Wilson.

Mountantonakis, Michalis and Yannis Tzitzikas (2016). "On measuring the lattice of commonalities among several linked datasets." In: *Proceedings of the VLDB Endowment* 9.12, pp. 1101–1112.

Mourelatos, Alexander PD (1978). "Events, processes, and states." In: *Linguistics and Philosophy* 2.3, pp. 415–434.

Nguyen, Tri, Mir Rosenberg, Xia Song, Jianfeng Gao, Saurabh Tiwary, Rangan Majumder, and Li Deng (2016). "MS MARCO: A human generated machine reading comprehension dataset." In: *Proceedings of the Workshop on Cognitive Computation: Integrating neural and symbolic approaches 2016 co-located with the 30th Annual Conference on Neural Information Processing Systems (NIPS 2016)*. Barcelona, Spain.

Nielsen, Rodney D, Jason Buckingham, Gary Knoll, Ben Marsh, and Leysia Palen (2008). "A taxonomy of questions for question

generation." In: *Proceedings of the Workshop on the Question Generation Shared Task and Evaluation Challenge*.

Norman, Don (2013). *The design of everyday things: Revised and expanded edition*. Basic books.

Norman, Donald A (1973). "Memory, knowledge, and the answering of questions." In: *Contemporary Issues in Cognitive Psychology: The Loyola Symposium*, pp. 135–165.

Nothegger, Clemens, Stephan Winter, and Martin Raubal (2004). "Selection of Salient Features for Route Directions." In: *Spatial Cognition & Computation* 4.2, pp. 113–136. ISSN: 1387-5868. DOI: 10.1207/s15427633scc0402_1.

Nouvel, Damien, Maud Ehrmann, and Sophie Rosset (2016). *Named Entities for Computational Linguistics*. John Wiley & Sons. ISBN: 1119268575.

Okoli, Chitu (2015). "A Guide to Conducting a Standalone Systematic Literature Review." In: *Communications of the Association for Information Systems* 37.

Ortmann, Jens and Werner Kuhn (2010). "Affordances As Qualities." In: *Proceedings of the 2010 Conference on Formal Ontology in Information Systems: Proceedings of the Sixth International Conference (FOIS 2010)*. Amsterdam, The Netherlands, The Netherlands: IOS Press, pp. 117–130. ISBN: 978-1-60750-534-1.

Padmanabhan, Venkata N and Jeffrey C Mogul (1996). "Using predictive prefetching to improve world wide web latency." In: *ACM SIGCOMM Computer Communication Review* 26.3, pp. 22–36.

Papadakis, Emmanuel, Bernd Resch, and Thomas Blaschke (2016). "A Function-based model of Place." In: *GIScience*. Vol. 1.

Perry, Matthew and John Herring (2012). "OGC GeoSPARQL-A geographic query language for RDF data." In: *OGC Implementation Standard* 40.

Peters, Matthew E, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer (2018). "Deep contextualized word representations." In: *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. New Orleans, USA, pp. 2227–2237.

Pitkow, James and Peter Pirolli (1999). "Mining Longest Repeating Subsequences to Predict World Wide Web Surfing." In: *Proceedings of the 2Nd Conference on USENIX Symposium on Internet Technologies and Systems - Volume 2*. USITS'99. Boulder, Colorado: USENIX Association, pp. 13–13.

Pomerantz, Jeffrey (2005). "A linguistic analysis of question taxonomies." In: *Journal of the American Society for Information Science and Technology* 56.7, pp. 715–728. ISSN: 15322882. DOI: 10.1002/asi.20162. URL: http://doi.wiley.com/10.1002/asi.20162.

Punjani, Dharmen, K Singh, Andreas Both, Manolis Koubarakis, Ioannis Angelidis, Konstantina Bereta, Themis Beris, Dimitris Bilidas, T Ioannidis, Nikolaos Karalis, et al. (2018). "Template-Based Question Answering over Linked Geospatial Data." In: *Proceedings of the 12th Workshop on Geographic Information Retrieval*. New York, NY, USA: ACM, p. 7.

Punjani, Dharmen et al. (2021). *Template-Based Question Answering over Linked Geospatial Data*. arXiv: 2007.07060 [cs.IR].

Purves, Ross S, Paul Clough, Christopher B Jones, Mark H Hall, and Vanessa Murdock (2018). "Geographic information retrieval: progress and challenges in spatial search of text." In: *Foundations and Trends in Information Retrieval* 12.2-3, pp. 164–318.

Purves, Ross S., Stephan Winter, and Werner Kuhn (2019). "Places in Information Science." In: *Journal of the Association for Information Science and Technology* 70.11, pp. 1173–1182. DOI: 10.1002/asi.24194.

Purves, Ross S. et al. (2007). "The design and implementation of SPIRIT: a spatially aware search engine for information retrieval on the Internet." In: *International Journal of Geographical Information Science* 21.7, pp. 717–745. DOI: 10.1080/13658810601169840.

Purves, Ross (2014). "Geographic information Retrieval: Are we making progress?" In: *NCGIA Specialist Meeting on Spatial Search*, pp. 1–6.

Purves, Ross and Curdin Derungs (2015). "From Space to Place: Place-Based Explorations of Text." In: *International Journal of Humanities and Arts Computing* 9.1, pp. 74–94. ISSN: 1753-8548. DOI: 10.3366/ijhac.2015.0139.

Purves, Ross, Alistair Edwardes, and Mark Sanderson (2008). "Describing the where--improving image annotation and search through geography." In: *Proceedings of the workshop on Metadata Mining for Image Understanding (MMIU 2008)*. Sheffield.

Purves, Ross, Alistair Edwardes, and Jo Wood (2011). "Describing place through user generated content." In: *First Monday*. ISSN: 13960466. DOI: 10.5210/fm.v16i9.3710.

Quesnot, T. and S. Roche (2015). "Platial or Locational Data? Toward the Characterization of Social Location Sharing." In: *48th Hawaii International Conference on System Sciences*, pp. 1973–1982. ISBN: 1530-1605. DOI: 10.1109/HICSS.2015.236.

Radlinski, Filip and Nick Craswell (2017). "A theoretical framework for conversational search." In: *Proceedings of the 2017 Conference on Conference Human Information Interaction and Retrieval*, pp. 117–126.

Rajpurkar, Pranav, Jian Zhang, Konstantin Lopyrev, and Percy Liang (2016). "SQuAD: 100,000+ Questions for Machine

Comprehension of Text." In: *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pp. 2383–2392.

Raubal, Martin and Stephan Winter (2002). "Enriching Wayfinding Instructions with Local Landmarks." In: *Geographic Information Science*. Ed. by Max J. Egenhofer and David M. Mark. Springer Berlin Heidelberg, pp. 243–259. ISBN: 978-3-540-45799-2.

Rauch, Erik, Michael Bukatin, and Kenneth Baker (2003). "A confidence-based framework for disambiguating geographic terms." In: *Proceedings of the HLT-NAACL 2003 Workshop on Analysis of Geographic References*, pp. 50–54.

Ravichandran, Deepak and Eduard Hovy (2002). "Learning surface text patterns for a question answering system." In: *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pp. 41–47.

Relph, E (1976). *Place and placelessness*. London: Pion.

Ren, Y., M. Tomko, F. Salim, K. Ong, and M. Sanderson (2017). "Analyzing Web Behavior in Indoor Retail Spaces." In: *Journal of the American Society for Information Science* 68.1, pp. 62–76. DOI: 10.1002/asi.23587.

Richardson, Matthew, Christopher JC Burges, and Erin Renshaw (2013). "MCtest: A challenge dataset for the open-domain machine comprehension of text." In: *Proceedings of the 2013 Conference on Empirical Methods in Natural Language Processing*, pp. 193–203.

Richter, Daniela, Stephan Winter, Kai-Florian Richter, and Lesley Stirling (2013). "Granularity of locations referred to by place descriptions." In: *Computers, Environment and Urban Systems* 41, pp. 88–99. ISSN: 0198-9715. DOI: https://doi.org/10.1016/j.compenvurbsys.2013.03.005.

Roberts, Kirk, Cosmin Adrian Bejan, and Sanda Harabagiu (2010). "Toponym disambiguation using events." In: *Twenty-Third International FLAIRS Conference*.

Robinson, William Peter and Susan J Rackstraw (1972). *A question of answers*. Vol. 5. Routledge and K. Paul.

Roche, Stéphane (2015). "Geographic information science II: Less space, more places in smart cities." In: *Progress in Human Geography* 40.4, pp. 565–573. ISSN: 0309-1325. DOI: 10.1177/0309132515586296.

Rooy, Robert (2004). "Utility of mention-some questions." In: *Research on Language and Computation* 2.3, pp. 401–416.

Sallaberry, Christian, Mauro Gaio, Damien Palacio, and Julien Lesbegueries (2008). "Fuzzying GIS topological functions for GIR needs." In: *Proceedings of the 5th Workshop on Geographic Information Retrieval*, pp. 1–8.

Sanderson, Mark and Janet Kohler (2004). "Analyzing geographic queries." In: *SIGIR Workshop on Geographic Information Retrieval*. Vol. 2, pp. 8–10.

Santos, Diana and Marcirio Chaves (2006). "The place of place in geographical IR." In: *quot; In Proceedings of GIR06 the 3rd Workshop on Geographic Information Retrieval (GIR 2006)(Seattle 10 th August 2006)*.

Santos, Rui, Patricia Murrieta-Flores, Pável Calado, and Bruno Martins (2018). "Toponym matching through deep neural networks." In: *International Journal of Geographical Information Science* 32.2, pp. 324–348. ISSN: 1365-8816. DOI: 10.1080/13658816.2017.1390119.

Saquete, Estela, J Luis Vicedo, Patricio Martínez-Barco, Rafael Munoz, and Hector Llorens (2009). "Enhancing QA systems with complex temporal question processing capabilities." In: *Journal of Artificial Intelligence Research* 35, pp. 775–811.

Scannell, Leila and Robert Gifford (2010). "Defining place attachment: A tripartite organizing framework." In: *Journal of Environmental Psychology* 30.1, pp. 1–10. ISSN: 0272-4944. DOI: https://doi.org/10.1016/j.jenvp.2009.09.006.

Scheider, Simon and Krzysztof Janowicz (2010). "Places as media of containment." In: *International Conference on GIScience*. Zurich, Switzerland.

— (2014). "Place reference systems: A constructive activity model of reference to places." In: *Applied Ontology* 9.2, pp. 97–127. ISSN: 15705838. DOI: 10.3233/AO-140134.

Scheider, Simon, Rogier Meerlo, Vedran Kasalica, and Anna-Lena Lamprecht (2020a). "Ontology of core concept data types for answering geo-analytical questions." In: *Journal of Spatial Information Science* 2020.20, pp. 167–201.

Scheider, Simon, Enkhbold Nyamsuren, Han Kruiger, and Haiqi Xu (2020b). "Geo-analytical question-answering with GIS." In: *International Journal of Digital Earth* 0.0, pp. 1–14. DOI: 10.1080/17538947.2020.1738568.

Scheider, Simon and Ross Purves (2013). "Semantic Place Localization from Narratives." In: *Proceedings of The First ACM SIGSPATIAL International Workshop on Computational Models of Place*. COMP '13. Orlando FL, USA: ACM, pp. 16–19. ISBN: 978-1-4503-2535-6. DOI: 10.1145/2534848.2534858.

Searle, John R, Herman Parret, and Jef Verschueren (1992). *(On) Searle on Conversation: Compiled and introduced by Herman Parret and Jef Verschueren*. Vol. 21. John Benjamins Publishing.

Seuren, Pieter AM (1984). "The comparative revisited." In: *Journal of Semantics* 3.1, pp. 109–141.

Shanon, Benny (1983). "Answers to where-questions." In: *Discourse Processes* 6.4, pp. 319–352.

Sharma, Yashvardhan and Sahil Gupta (2018). "Deep Learning Approaches for Question Answering System." In: *Procedia Computer Science* 132. International Conference on Computational Intelligence and Data Science, pp. 785 –794. ISSN: 1877-0509. DOI: https://doi.org/10.1016/j.procs.2018.05.090.

Singh, Kuldeep, Andreas Both, Arun Sethupat, and Saeedeh Shekarpour (2018). "Frankenstein: A Platform Enabling Reuse of Question Answering Components." In: *The Semantic Web*. Ed. by Aldo Gangemi, Roberto Navigli, Maria-Esther Vidal, Pascal Hitzler, Raphaël Troncy, Laura Hollink, Anna Tordai, and Mehwish Alam. Cham: Springer International Publishing, pp. 624–638. ISBN: 978-3-319-93417-4.

Smith, Barry and Achille C Varzi (2000). "Fiat and bona fide boundaries." In: *Philosophical and Phenomenological Research*, pp. 401–420. ISSN: 0031-8205.

Sneiders, Eriks (2002). "Automated Question Answering Using Question Templates That Cover the Conceptual Model of the Database." In: *Natural Language Processing and Information Systems*. Ed. by Birger Andersson, Maria Bergholtz, and Paul Johannesson. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 235–239. ISBN: 978-3-540-36271-5.

Song, Hongya, Zhaochun Ren, Shangsong Liang, Piji Li, Jun Ma, and Maarten de Rijke (2017). "Summarizing answers in non-factoid community question-answering." In: *Proceedings of the Tenth ACM International Conference on Web Search and Data Mining*, pp. 405–414.

Spencer, Donna (2009). *Card sorting: Designing usable categories*. Rosenfeld Media. ISBN: 1933820071.

Srihari, Rohini and Wei Li (Apr. 2000). "A Question Answering System Supported by Information Extraction." In: *Sixth Applied Natural Language Processing Conference*. Seattle, Washington, USA: Association for Computational Linguistics, pp. 166–172. DOI: 10.3115/974147.974170.

Stadler, Claus, Jens Lehmann, Konrad Höffner, and Sören Auer (Oct. 2012). "LinkedGeoData: A Core for a Web of Spatial Open Data." In: *Semantic Web* 3.4, 333–354. ISSN: 1570-0844.

Stefano De, Sabbata, Mizzaro Stefano, and Reichenbacher Tumasch (2015). "Geographic dimensions of relevance." In: *Journal of Documentation* 71.4, pp. 650–666. ISSN: 0022-0418. DOI: 10.1108/JD-12-2013-0167.

Sui, Daniel and Michael Goodchild (2011). "The convergence of GIS and social media: challenges for GIScience." In: *International Journal of Geographical Information Science* 25.11, pp. 1737–1748.

Sun, Huan, Hao Ma, Wen-tau Yih, Chen-Tse Tsai, Jingjing Liu, and Ming-Wei Chang (2015). "Open Domain Question Answering

via Semantic Enrichment." In: *Proceedings of the 24th International Conference on World Wide Web*. WWW '15. Florence, Italy: International World Wide Web Conferences Steering Committee, 1045–1055. ISBN: 9781450334693. DOI: 10.1145/2736277.2741651.

Suomela, Riku, Harri Lakkala, and Ilkka Salminen (2009). *Displaying a map having a close known location*. US Patent 7,480,567.

Tang, Lappoon R. and Raymond J. Mooney (2001). "Using Multiple Clause Constructors in Inductive Logic Programming for Semantic Parsing." In: *Machine Learning: ECML 2001*. Ed. by Luc De Raedt and Peter Flach. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 466–477. ISBN: 978-3-540-44795-5.

Tax, Niek, Irene Teinemaa, and Sebastiaan J. van Zelst (2020). "An interdisciplinary comparison of sequence modeling methods for next-element prediction." In: *Software and Systems Modeling*. DOI: 10.1007/s10270-020-00789-3.

Thalhammer, Andreas and Achim Rettinger (2016). "PageRank on Wikipedia: towards general importance scores for entities." In: *European Semantic Web Conference*. Springer, pp. 227–240.

Tomko, Martin and Ross S Purves (2008). "Categorical prominence and the characteristic description of regions." In: *Proceedings of the Semantic Web Meets Geospatial Application Workshop, held in conjunction with AGILE*.

Tomko, Martin and Stephan Winter (2009). "Pragmatic Construction of Destination Descriptions for Urban Environments." In: *Spatial Cognition & Computation* 9.1, pp. 1–29. ISSN: 1387-5868. DOI: 10.1080/13875860802427775.

Trivedi, Priyansh, Gaurav Maheshwari, Mohnish Dubey, and Jens Lehmann (2017). "Lc-QuAD: A corpus for complex question answering over knowledge graphs." In: *International Semantic Web Conference*. Springer, pp. 210–218.

Tuan, Yi-Fu (1977). *Space and place: The perspective of experience*. U of Minnesota Press. ISBN: 1452905533.

— (1990). *Topophilia: A study of environmental perceptions, attitudes, and values*. Columbia University Press. ISBN: 0231513283.

Turing, Alan Mathison (1950). "Computing machinery and intelligence-AM Turing." In: *Mind* 59.236, p. 433.

Turner, Phil and Susan Turner (2006). "Place, sense of place, and presence." In: *Presence: Teleoperators & Virtual Environments* 15.2, pp. 204–217. ISSN: 1054-7460.

Tversky, Barbara and Kathleen Hemenway (1983). "Categories of environmental scenes." In: *Cognitive Psychology* 15.1, pp. 121–149. ISSN: 0010-0285. DOI: https://doi.org/10.1016/0010-0285(83)90006-3.

Vahedi, Behzad, Werner Kuhn, and Andrea Ballatore (2016). "Question-Based Spatial Computing—A Case Study." In: *Geospatial Data in a Changing World*. Ed. by Tapani Sarjakoski,

Maribel Yasmina Santos, and L. Tiina Sarjakoski. Cham: Springer International Publishing, pp. 37–50. ISBN: 978-3-319-33783-8.

Vanclay, Frank (2008). "Place matters." In: *Making Sense of Place: Exploring Concepts and Expressions of Place through Different Senses and Lenses*. National Museum of Australia Press, Canberra, pp. 3–11.

Varges, Sebastian, Fuliang Weng, and Heather Pon-Barry (July 2006). "Interactive Question Answering and Constraint Relaxation in Spoken Dialogue Systems." In: *Proceedings of the 7th SIGdial Workshop on Discourse and Dialogue*. Sydney, Australia: Association for Computational Linguistics, pp. 28–35.

Vasardani, Maria, Sabine Timpf, Stephan Winter, and Martin Tomko (2013). "From Descriptions to Depictions: A Conceptual Framework." In: *Spatial Information Theory*. Ed. by Thora Tenbrink, John Stell, Antony Galton, and Zena Wood. Springer International Publishing, pp. 299–319. ISBN: 978-3-319-01790-7.

Vasardani, Maria, Martin Tomko, and Stephan Winter (2016). "The Cognitive Aspect of Place Properties." In: *International Conference on GIScience*. Montreal, Canada.

Vasardani, Maria and Stephan Winter (2016). "Place properties." In: *Advancing Geographic Information Science: The Past and Next Twenty Years*. GSDI Association Press Needham, MA, pp. 243–254.

Vasardani, Maria, Stephan Winter, Kai-Florian Richter, Lesley Stirling, and Daniela Richter (2012). "Spatial Interpretations of Preposition "At"." In: *Proceedings of the 1st ACM SIGSPATIAL International Workshop on Crowdsourced and Volunteered Geographic Information*. GEOCROWD '12. Redondo Beach, California: Association for Computing Machinery, 46–53. ISBN: 9781450316941. DOI: 10.1145/2442952.2442961.

Voorhees, E. (1999). "The TREC-8 Question Answering Track Report." In: *Proceedings of the 8th Text REtrieval Conference (TREC-8), 1999*.

Wang, Lee, Chuang Wang, Xing Xie, Josh Forman, Yansheng Lu, Wei-Ying Ma, and Ying Li (2005). "Detecting dominant locations from search queries." In: *Proceedings of the 28th Annual International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 424–431.

Wang, Mengqiu, Noah A Smith, and Teruko Mitamura (2007). "What is the Jeopardy model? A quasi-synchronous grammar for QA." In: *Proceedings of the 2007 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning (EMNLP-CoNLL)*, pp. 22–32.

Wang, W., Y. Tian, H. Wang, and W. Ku (2020). "A Natural Language Interface for Database: Achieving Transfer-learnability Using Adversarial Method for Question Understanding." In: *2020 IEEE*

36th International Conference on Data Engineering (ICDE), pp. 97–108.

Wang, X., Y. Zhang, M. Chen, X. Lin, H. Yu, and Y. Liu (2010). "An evidence-based approach for Toponym Disambiguation." In: *Proceedings of 18th International Conference on Geoinformatics*, pp. 1–7. DOI: 10.1109/GEOINFORMATICS.2010.5567805.

Werlen, Benno (2003). *Society, action and space*. Routledge. ISBN: 1134913966.

White, Ryen W., Matthew Richardson, and Wen-tau Yih (2015). "Questions vs. Queries in Informational Search Tasks." In: *Proceedings of the 24th International Conference on World Wide Web*. WWW '15 Companion. Florence, Italy: Association for Computing Machinery, 135–136. ISBN: 9781450334730. DOI: 10.1145/2740908.2742769.

Wilson, Deirdre and Dan Sperber (2002). "Relevance theory." In: *Handbook of Pragmatics*. Blackwell.

Wilson, James Arthur (2018). *Systems and methods for presenting personalized map labels*. US Patent App. 15/811,376.

Winter, Stephan (2009). "Spatial Intelligence: Ready for a Challenge?" In: *Spatial Cognition & Computation* 9.2, pp. 138–151. ISSN: 1387-5868. DOI: 10.1080/13875860902893868.

Winter, Stephan and Christian Freksa (2012). "Approaching the notion of place by contrast." In: *Journal of Spatial Information Science* 5.1, pp. 31–50. DOI: 10.5311/JOSIS.2012.5.90.

Winter, Stephan, Ehsan Hamzei, Nico Van de Weghe, and Kristien Ooms (2018). "A Graph Representation for Verbal Indoor Route Descriptions." In: *Spatial Cognition XI*. Cham: Springer International Publishing, pp. 77–91. ISBN: 978-3-319-96385-3.

Winter, Stephan and Marie Truelove (2013). "Talking About Place Where it Matters." In: *Cognitive and Linguistic Aspects of Geographic Space: New Perspectives on Geographic Information Research*. Ed. by Martin Raubal, David M. Mark, and Andrew U. Frank. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 121–139. ISBN: 978-3-642-34359-9. DOI: 10.1007/978-3-642-34359-9_7.

Xie, Zhiwen, Zhao Zeng, Guangyou Zhou, and Tingting He (2016). "Knowledge Base Question Answering Based on Deep Learning Models." In: *Natural Language Understanding and Intelligent Applications*. Ed. by Chin-Yew Lin, Nianwen Xue, Dongyan Zhao, Xuanjing Huang, and Yansong Feng. Cham: Springer International Publishing, pp. 300–311. ISBN: 978-3-319-50496-4.

Xu, Boyan, Ruichu Cai, Zhenjie Zhang, Xiaoyan Yang, Zhifeng Hao, Zijian Li, and Zhihao Liang (2019). "NADAQ: natural language database querying based on deep learning." In: *IEEE Access* 7, pp. 35012–35017.

Xu, Haiqi., Ehsan. Hamzei, Enkhbold. Nyamsuren, Han. Kruiger, Stephan. Winter, Martin. Tomko, and Simon. Scheider (2020). "Extracting interrogative intents and concepts from geo-analytic questions." In: *AGILE: GIScience Series* 1, pp. 1 –23. DOI: 10.5194/agile-giss-1-23-2020.

Xu, Kun, Siva Reddy, Yansong Feng, Songfang Huang, and Dongyan Zhao (2016). "Question Answering on Freebase via Relation Extraction and Textual Evidence." In: *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 2326–2336.

Yang, Xiao, Craig Macdonald, and Iadh Ounis (2018). "Using word embeddings in Twitter election classification." In: *Information Retrieval Journal* 21.2, pp. 183–207. ISSN: 1573-7659. DOI: 10.1007/s10791-017-9319-5. URL: https://doi.org/10.1007/s10791-017-9319-5.

Yang, Yi, Wen-tau Yih, and Christopher Meek (Sept. 2015). "WikiQA: A Challenge Dataset for Open-Domain Question Answering." In: *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*. Lisbon, Portugal: Association for Computational Linguistics, pp. 2013–2018. DOI: 10.18653/v1/D15-1237. URL: https://www.aclweb.org/anthology/D15-1237.

Ye, Mao, Krzysztof Janowicz, Christoph Mülligann, and Wang-Chien Lee (2011). "What you are is when you are: the temporal dimension of feature types in location-based social networks." In: *Proceedings of the 19th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems*. ACM, pp. 102–111. ISBN: 1450310311.

Yen, Show-Jane, Yu-Chieh Wu, Jie-Chi Yang, Yue-Shi Lee, Chung-Jung Lee, and Jui-Jung Liu (2013). "A support vector machine-based context-ranking model for question answering." In: *Information Sciences* 224, pp. 77 –87. ISSN: 0020-0255. DOI: https://doi.org/10.1016/j.ins.2012.10.014.

Yu, Tao, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, et al. (2018). "Spider: A Large-Scale Human-Labeled Dataset for Complex and Cross-Domain Semantic Parsing and Text-to-SQL Task." In: *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 3911–3921.

Yulianti, Evi, Ruey-Cheng Chen, Falk Scholer, and Mark Sanderson (2016). "Using semantic and context features for answer summary extraction." In: *Proceedings of the 21st Australasian Document Computing Symposium*, pp. 81–84.

Zelle, John M. and Raymond J. Mooney (1996). "Learning to Parse Database Queries Using Inductive Logic Programming." In: *Proceedings of the Thirteenth National Conference on Artificial*

Intelligence - Volume 2. AAAI'96. Portland, Oregon: AAAI Press, 1050–1055. ISBN: 026251091X.

Zhao, Yiming, Jin Zhang, Xue Xia, and Taowen Le (2019). "Evaluation of Google question-answering quality." In: *Library Hi Tech* 37.2, pp. 312–328. ISSN: 07378831. DOI: 10.1108/LHT-10-2017-0218.

Zhou, Ming (2018). "What Will Search Engines Be Changed by NLP Advancements." In: *Proceedings of the 2018 ACM SIGIR International Conference on Theory of Information Retrieval*. ICTIR '18. Tianjin, China: Association for Computing Machinery, p. 7. ISBN: 9781450356565. DOI: 10.1145/3234944.3241521.

Ziv, Jacob and Abraham Lempel (1978). "Compression of individual sequences via variable-rate coding." In: *IEEE Transactions on Information Theory* 24.5, pp. 530–536.

Zou, Lei, Ruizhe Huang, Haixun Wang, Jeffrey Xu Yu, Wenqiang He, and Dongyan Zhao (2014). "Natural Language Question Answering over RDF: A Graph Data Driven Approach." In: *Proceedings of the 2014 ACM SIGMOD International Conference on Management of Data*. SIGMOD '14. Snowbird, Utah, USA: Association for Computing Machinery, 313–324. ISBN: 9781450323765. DOI: 10.1145/2588555.2610525.