FAKULTÄT FÜR !NFORMATIK
Faculty of Informatics

Spatial Search & Computation in Urban Areas

DISSERTATION

zur Erlangung des akademischen Grades

Doktorin der Technischen Wissenschaften

eingereicht von

Dipl.-Ing. Heidelinde Hobel, BSc.
Matrikelnummer 0625620

an der Fakultät für Informatik
der Technischen Universität Wien

Betreuung: O.Univ.Prof. Dipl.-Ing. Dr.techn. Andrew U. Frank

Diese Dissertation haben begutachtet:

Andrew U. Frank

A Min Tjoa

Wien, 30. Juni 2016

Heidelinde Hobel

FAKULTÄT
FÜR !NFORMATIK
Faculty of Informatics

Spatial Search & Computation in Urban Areas

DISSERTATION

submitted in partial fulfillment of the requirements for the degree of

Doktorin der Technischen Wissenschaften

by

Dipl.-Ing. Heidelinde Hobel, BSc.
Registration Number 0625620

to the Faculty of Informatics
at the Vienna University of Technology

Advisor: O.Univ.Prof. Dipl.-Ing. Dr.techn. Andrew U. Frank

The dissertation has been reviewed by:

Andrew U. Frank	A Min Tjoa

Vienna, 30th June, 2016

Heidelinde Hobel

Erklärung zur Verfassung der Arbeit

Dipl.-Ing. Heidelinde Hobel, BSc.

Hiermit erkläre ich, dass ich diese Arbeit selbständig verfasst habe, dass ich die verwendeten Quellen und Hilfsmittel vollständig angegeben habe und dass ich die Stellen der Arbeit – einschließlich Tabellen, Karten und Abbildungen –, die anderen Werken oder dem Internet im Wortlaut oder dem Sinn nach entnommen sind, auf jeden Fall unter Angabe der Quelle als Entlehnung kenntlich gemacht habe.

Wien, 30. Juni 2016

Heidelinde Hobel

Acknowledgements

This thesis would not be possible without support of a number of people:

First, I would like to thank my advisors Prof. Andrew U. Frank and Prof. A Min Tjoa. I am grateful for all the insightful discussions and suggestions which shaped the form of this thesis. Since I was new in the field of Geographic Information Science, I appreciate all the help my advisors could provide me.

Special thanks go to my mentor Paolo Fogliaroni.

Thanks to my colleagues at the Technical University of Vienna: Amin Abdalla, Eva Maria Holy, Jürgen Hahn, Markus Mayr, Thomas Fuhrmann, Gerhard Navratil, Paul Weiser, and Christian Reuter.

Thanks to my colleagues at the my work: Johannes Heurix, Markus Rumpold, Stefan Fenz, and Thomas Neubauer.

Lastly, I would like to thank my family and by boyfriend. Without their support I would have been lost.

Kurzfassung

Geoinformationssuchmaschinen stellen dem Benutzer eine Reihe von Funktionen zur Verfügung wie zum Beispiel Geocoding oder die Suche nach Plätzen mit bestimmten kategorischen Tags. Während komplexe Algorithmen zur Verfügung stehen um Routen oder Trips zu planen, sind Suchalgorithmen vernachlässigt worden die kognitive Fragestellungen zu beantworten suchen. Zum Beispiel können Geoobjekte nur dann gefunden werden, wenn diese explizit in das fundamentale Geoinformationssystem eingetragen worden sind. Eine ähnliche, aber wohl noch größere Herausforderung stellt die Suche nach den zur Erholung dienenden Gebieten dar, welche sehr speziell von der Wahrnehmung der Menschen abhängen.

Diese Arbeit beschäftigt sich daher mit der Entwicklung der Theorie zur Integration von kognitiven Modellen in Geoinformationssystemen. Besonderes Augenmerk wurde auf rechner- und algorithmische Methoden zur Informationsgewinnung und Verarbeitung gelegt. Eine spezifisch angesprochene Fragestellung ist die Suche nach "Cognitive Regions", einem Begriff aus den Kognitionswissenschaften, welche aufgrund von wahrgenommenen möglichen Aktivitäten gebildet werden. Ein direkter Anwendungsbereich ergibt sich in der Integration der vorgeschlagenen Methoden in Geoinformationssuchmaschinen der Zukunft.

Als Gundlage wird somit eine Generalisierung von Orten und Gebieten in Form einer Segmentierungstechnik vorgeschlagen, welche aus der traditionellen Bildverarbeitung inspiriert ist. Basierend auf dieser Grundlage, werden Methoden untersucht um geografische Bereiche nur auf Grundlage von User Generated Content und Volunteered Geographic Information herauszulesen. Dabei stellt Natural Language Processing einen fundamentalen Baustein in der Datenverarbeitungskette dar. Abschließend wird die Integration von semantisch angereicherten Graphen in das entwickelte Framework untersucht.

Als Proof-of-Concept werden verschiedene Fragestellungen, die der Geoinformationswissenschaft entspringen, diskutiert und die vorgeschlagenen Ansätze evaluiert.

Abstract

Today's Geographic Information Systems and spatial search engines support people to search for spatial data by 'name' or categorical tags, or in contrast, by concrete address or location data. While sophisticated algorithms exist to compute complex routes or planning trips, spatial search is inadequately supported for answering nuanced and fuzzy questions such as searching for 'recreational' regions within a city.

To address this issue, spatial search engines have to incorporate cognitive models of spatial search behaviour, allowing sense-making of complex queries expressed according to human's conceptualization of place. In this thesis it is argued that cognitive areas bridge the gap between cognitive models and today's possibilities of spatial search engines.

The phenomena of cognitive regions capture the ability of humans to conceptualize and generalize space according to the activities they can carry out at a given place. Therefore, in this thesis cognitive areas are proposed, which are inspired by traditional image segmentation. Based on this foundation, methods are investigated that allow to infer the geometric extent of "cognitive regions" on the basis of User Generated Content and Volunteered Geographic Information. Hence, Natural Language Processing is one of the fundamental building blocks in the processing of huge amounts of data. Finally, the integration of semantically enriched conceptual graphs is investigated.

As proof-of-concept, different problems, originating from Geographic Information Science, are discussed and the proposed approaches are evaluated.

Contents

CHAPTER 1

Introduction

Volunteered Geographic Information has received much attention in recent years due to its potential for being a freely available Geographic Information System, where casual users can consume and produce geo(spatial) information. Goodchild [2007] coined the term "citizens as sensors" and draws attention to the benefits and disadvantages of bringing Geographic Information Systems into the reach of casual users. OpenStreetMap[1], the world's leading Volunteered Geographic Information platform, has produced a large amount of crowd-sourced geographic information. Wikimapia[2] allows users to add textual descriptions to spatial entities, while using FourSquare[3] allows users to check in at locations and share this information world-wide and among their friends and community. While the data drawn from these spatial knowledge bases is more or less accurate, Volunteered Geographic Information democratizes spatial systems which were for a long time under the control of governments, meaning access to them was moderated. Although Volunteered Geographic Information is publicly available and usable, little research has been done to exploit this vast amount of geographic information in the field of urban computational reasoning. In the coming age of smart cities, this voluntarily produced geographic information can be exploited for novel context- and location-based services, recommendation and urban planning systems, and much more. Despite the fact that geographic questions have been addressed in different fields of geographic information science, there remains a major gap of integrating the cognitive understanding into computational models for spatial search and reasoning. This thesis will therefore concern itself with the question of integrating spatial search and cognitive models for urban areas into the next generation of Geographic Information Systems.

[1] https://www.openstreetmap.org/
[2] http://wikimapia.org/
[3] https://de.foursquare.com/

1.1 Motivation

The notion of *place* and its representations have become prominent research topics within the field of geographic information science [Gao et al., 2013, Goodchild, 2011, Scheider and Janowicz, 2014, Kuhn, 2001]. Modern Geographic Information Systems are designed to deal with the notion of *space*, rather than with that of *place*. While space is an abstract concept that we can satisfactorily model through mathematical abstractions, the notion of place is strictly related to human conceptualizations of space, and how to represent and automatically process it within Geographic Information Systems remains an open question.

The notion of place may correspond to different things, such as e.g. Points of Interests, geographic regions, or settings (i.e. aggregations of spatial features) [Schatzki, 1991]. In this thesis, we regard a place as a region of space that is homogeneous with respect to certain criteria. We adopt the taxonomy for geographic regions proposed in [Montello, 2003] and focus on the category of so called *cognitive regions*. These regions are derived by people as they experience the world. Cognitive regions are vague (i.e. they do not have crisp boundaries), and their geographic interpretation may (and usually does) differ slightly among several individuals, as shown, for example, for the cognitive regions of downtown Santa Barbara [Montello et al., 2003], and Southern and Northern California, and Alberta [Montello et al., 2014].

As of today, the main support for place-based search in Geographic Information Retrieval systems relies on the use of place names as provided by gazetteers. However, this type of search falls short when it comes to capturing human conceptualization of places. For that reason, this thesis proposes a novel framework that facilitates capturing humans' conceptualization of place and preparing it for place-based questions by using computational models. One essential part is to model place based on mathematical abstractions and computational models according to the type of place discussed above, i.e. *cognitive regions*. The framework itself is tailored to extract place-related information from crowd-sourced spatial knowledge bases and User Generated Content, and process it for spatial needs. To obtain human conceptualizations, different techniques are investigated to model and process data obtained from the mentioned sources according to spatial problems such as segmentation of space and similarity of regions, Machine Learning to derive cognitive regions, and contextual analysis of User Generated Content for spatial question answering.

In conclusion, spatial search plays a fundamental role in our everyday life; from simple searches for restaurants in our vicinity, to more complex routing tasks. In a cognitive way, however, searching *should not* be restricted to geocoded entities or to computing shortest paths. In this work, synergies from different fields were investigated in order to build a framework for spatial search incorporating Natural Language Processing, Information Retrieval, Machine Learning, and Cognitive Science. The framework was implemented and evaluated based on different questions originating from geographic information science. The results achieved show the potential of the framework for use as a spatial recommendation system.

1.2 Methodology

As one of the first, Schatzki [1991] introduced the term settings as human-formed constellations of places. Recent research [Ballatore, 2014, Hobel et al., 2015, 2016] has taken up this new concept of place constellations as new research direction for Geographic Information Retrieval. In computational terms, the search for place constellations can be modeled in analogy to image segmentation and traditional Information Retrieval, where each *pixel, cell, segment*, or *document* contains a number of aggregated places. The essential parts of the framework are published in three papers, which are briefly introduced in the following and form the basis for the ideas presented in Chapter 3, Chapter 4, and Chapter 5, respectively.

A Semantic Region Growing Algorithm: Extraction of Urban Settings [Hobel et al., 2015] The vast amount of Volunteered Geographic Information allows for novel ways of understanding, analyzing and generalizing urban areas. In this chapter, a cognitive model for the understanding of urban areas is designed by extracting geographic areas that are conceptually uniform regarding certain activities. Thereby, *urban settings* are a more accurate way of generalizing cities, since they are more closely related to how humans make sense of urban space. To this end, a semantic region growing algorithm is formalized, allowing the segmentation of geographic areas, and, consequently, the similarity measurement of urban areas building on a semantic foundation becomes possible within a computational model.

Deriving the Geographic Footprint of Cognitive Regions [Hobel et al., 2016] The characterization of *place* and its representation in current Geographic Information Systems has become a prominent research topic. This chapter is concerned with places that are cognitive regions, and presents a computational framework to derive the geographic footprint of these regions. The main idea is to use Natural Language Processing tools to identify unique geographic features from User Generated Content sources consisting of textual descriptions of places. These features are used to detect an initial area that the descriptions refer to on a map. A semantic representation of this area is extracted from a Geographic Information System and passed over to a Machine Learning algorithm that locates other areas according to semantic similarity. As a case study, we employ the proposed framework to derive the geographic footprint of the *historic center of Vienna*, and validate the results by comparing the derived region against a historical map of the city. Furthermore, by using an ontological model we will show that a city can be categorized into functional regions.

Extracting Semantics of Places from User Generated Content [Hobel and Fogliaroni, 2016] The next generation of Geographic Information Systems should support place-based searches. As mentioned, the notion of place is a vague concept that strictly relates to human conceptualization of space. In this chapter, we regard places as cognitive regions affording activity opportunities and present a computational workflow to populate the model with information from User Generated Content available on the Web. An algorithmic realization is provided that relies on the Resource Description Framework along with a real example utilising an implementation of the proposed workflow that relies on OpenStreetMap and TripAdvisor[4] as data sources.

The following contributions discuss different problems in geographic information science, which gave inspiration to the concrete problems addressed in this thesis:

1. Visualisation of User-Generated Event Information: Towards Geospatial Situation Awareness Using Hierarchical Granularity Levels [Hobel et al., 2014]

2. Exploiting Linked Spatial Data and Granularity Transformations [Hobel and Frank, 2014]

3. Implementing Naïve Geography via Qualitative Spatial Relation Queries [Fogliaroni and Hobel, 2015]

1.2.1 Hypothesis

The hypothesis of this thesis is based on the following foundational theories and observations: (i) cognitive science describes cognitive regions as regions in the mind [Montello, 2003]; (ii) there is empirical evidence that cognitive regions have vague boundaries [Montello et al., 2003]; (iii) the notion of cognitive regions is related to human conceptualization and the activities they afford [Schatzki, 1991]. Following this argumentation, this work aims at demonstrating the following hypothesis:

> "The extraction and processing of cognitive regions can be formalized within a computational model."

1.2.2 Research Questions

The fundamental first research question of this thesis is given as follows:

> "Can cognitive regions be formalized and processed with the synergistic interplay of methods arising in different fields?"

This major research question is broken down into the following more specific research questions:

RQ_2**:** How to model cities as regions of functional areas?

[4] https://www.tripadvisor.at/

RQ_3: How can semantic similarity be integrated in spatial search engines?

RQ_4: How to incorporate statistics and Machine Learning in spatial Information Retrieval?

RQ_5: Can Natural Language Processing be used in Geographic Information Retrieval?

RQ_6: How to build the next generation of knowledge bases for spatial search?

1.2.3 Contribution of this Thesis

This thesis exploits synergies of Geoinformation and Informatics to advance Geographic Information Systems. In particular, the main contributions are:

- The formal framework for the 'semantic' generalization of urban areas is defined, which exploits synergies from image segmentation techniques and traditional Information Retrieval.
- A novel approach to identify cognitive regions is introduced, which operates on unstructured data, and is based on Natural Language Processing and Machine Learning.
- A semantic graph structure tailored for spatial search tasks, and an algorithmic solution to populate the introduced model are proposed. The algorithmic solution operates on textual descriptions of spatial areas.

1.3 Outline of this Thesis

The remainder of this thesis is structured as follows:

In Chapter 2 we provide an overview of the proposed framework and discuss the relationships between the different essential concepts developed in this thesis.

In Chapter 3, the concepts of place and cognitive regions are discussed, and some fundamental background information is provided. Furthermore, an approach of clustering into uniform and homogeneous regions in urban areas is introduced, and a partitioning technique inspired by image segmentation is proposed. Building upon cognitive regions, semantic similarity is discussed, and approaches to quantitatively compare different extracted regions are introduced.

An automated approach for extracting cognitive regions from unstructured text is the subject of Chapter 4. To that end, techniques from Natural Language Processing are utilized, and a combinatorial place matching procedure is introduced which is tailored to OpenStreetMap's knowledge base.

In Chapter 5, we discuss the creation of a knowledge base that aims at further enabling more "intelligent" spatial search engines. With this knowledge base, more complex and nuanced questions about urban areas can be answered.

Chapters 3-5 each conclude with a presentation of possible areas of application, as well as a case study.

Finally, a detailed discussion of the results is subject of Chapter 6, and conclusions are drawn in Chapter 7, where possible extensions and future work are also outlined.

CHAPTER 2

Overview

Human conceptualization of space is one of the main research questions in Geographic Information Science, Spatial Information Theory, Urban Planning, and many other disciplines [Karwan and Frank, 2012, Lynch, 1960, Tuan, 1979]. Many have studied the way humans navigate through or reason about space [Lynch, 1960, Raubal, 2001]. Building on the findings of such studies, computational models and applications have been developed that simulate human conceptualization in order to improve the usability of software, or to equip computer systems with basic reasoning capabilities for dealing with tasks involving a spatial component. This chapter provides an overview of the information system developed in this thesis, and how its contributions holistically advance spatial information systems.

2.1 Information needs

Spatial search is at the core of human activity. However, little research has been done in the field of cognitively supported spatial search. Spatial search is more or less restricted to names and categorical attributes in which semantics of places is expressed, and to spatial relations among the entities, e.g "restaurant in Vienna", or "nearby public transport". Modern Geographic Information Systems are designed to deal with the notion of *space*, rather than with the notion of *place*. While the term space is sufficiently represented by means of mathematical abstractions, the notion of place is related to human conceptualization, and is, therefore, a cognitive concept whose computational interpretation remains a challenging task. So far, spatial search engines support users in simple requests such as the search for Points of Interest in a given city. Consequently, gazetteers and ontologies are used as geographic knowledge bases containing attributes and topological information.

With the theory and computational models developed in this thesis, one important possible application lies in the advancement of spatial search engines to support more intuitive requests, tailored to the needs of the user and expressed in a way that is more natural to the way of human conceptualization of place than what is currently available. To that end, it is argued that traditional techniques of different fields can be combined to allow for the holistic processing of geographic information according to users' different spatial search needs. With the aid of computational models, such systems allow *more cognitive and intuitive questions* about the environment. In the past, places were represented and processed as single entities, disregarding the interwoven connections between them. As one of the first, Ballatore [2014] expressed the characteristics of places as follows:

> "Places are inescapably multi-faceted (comprising diverse processes), they are socially constructed (emerging as the result of human agency and practices), relational (emerging in a context, not in a vacuum), scale-dependent (different places exist at different scales), and they are dynamic (emerging, changing, and ultimately disappearing)."

In the sense of the above definition, places as aggregate concepts depend highly on the semantics that implicitly exists between places. Our developed approach in this thesis is to tackle this complexity by processing places based on the distribution of semantic attributes in order to model the interconnections of co-occurring objects. Computationally, these aggregated concepts can be processed by a methodology using a simplified "semantic" representation, specifically the bag-of-words model.

Definition 2.1.1: Bag-of-Words Model

The bag-of-words model is a representation of the frequencies of words used in text classification, disregarding the words' order of appearance in the text. The vocabulary is the set of all words occurring in the document $\mathcal{T} := \{t_i \colon i = 1, \ldots, n\}$, where n denotes the size of the vocabulary. Specifically, each document in the bag-of-words model is represented by the vector $(x_1, \ldots, x_n)$, where x_i is the frequency of the word t_i in this document.

The bag-of-words model allows for the processing of places by two different approaches. Firstly (i), processing the semantics of places according to techniques employed in image segmentation; and secondly (ii), utilizing Machine Learning techniques for identifying patterns of aggregated places. Considering techniques employed in image segmentation, a simplified representation of cities as homogeneous and functional regions can be created. Partitioning a map into functional regions allows for nuanced questions about the environment. For instance, this partitioning approach can be useful for tourists asking questions about which areas present similar characteristics to the areas they are familiar with, comparing them based on a user's preference model. Accordingly, Machine Learning based on the bag-of-words model has the potential to derive homogeneous and uniform areas automatically. The pattern recognition approach, which is based on Machine Learning, has several areas of application for aiding humans in finding places that are similar to the areas they are familiar with, automatically retrieving the geometric extent of cognitive regions. It can be argued that most cognitive regions have homogeneous characteristics, which leads humans to conceptualize these regions as cognitively coherent areas. To capture the human conceptualization of places that are composed of other places, User Generated Content can be exploited to infer the spatial footprint of cognitive regions. Moreover, User Generated Content has the potential to aid humans in complex questions. The proposed model, therefore, will preserve the semantics, and will be tailored to the spatial needs of the user. The model itself helps in asking for activities that can be done in combination, or in deriving implicit semantic relations, such as which places are conceptualized as *near* to each other. For the concept *near*, people often associate to one place multiple Points of Interest which are in the proximity, resulting in conceptualizations of *near* that are largely independent of the actual distance.

2.2 System Architecture

In Figure 2.1 a schematic conceptual model of the framework developed in this thesis is illustrated. It will be explained in detail in the following sections. The objective is to process and provide the information according to the spatial needs of the user. As a result, for a given spatial concept as input (e.g. shopping area), it is possible to combine all computation steps (e.g. computing the spatial extents of all regions of this type) into a pre-processing step, where the results are afterwards stored in the system, and remain available for later reuse. This scheduling of the workload has beneficial implications for the time and resources needed when the system is presented with a new query, since we can argue that in the area of tourist information many queries can be anticipated. A part of such a request will therefore already be available in the system when it is actually requested by a user (e.g. a tourist). The underlying processing and subsequent use of the results of queries can be classified into three conceptual classes: (i) segmentation and similarity, (ii) machine learning, and (iii) context analysis.

The foundational basis of the framework is a spatial database management system of Points of Interest and, to model places as aggregates, a grid-based clustering of the Points of Interest which yields the bag-of-words. This simplified representation in the form of a vector space model is exploited for a generalization method that partitions a map into “activity clusters” that further allow for finer-grained measurements of similarity. Questions the model can answer are, for example, “where are the activity centers of different cities?”, and “how similar are they according to my preferences?”. The segmentation and similarity model is a subject of Section 2.3 and of Chapter 3.

The grid-based clustering approach is reused for our Machine Learning approach, which allows questions such as which areas are similar to the areas a user is familiar with, or, by incorporating User Generated Content, makes it possible to automatically derive the geographic footprint of cognitive regions. In order to extract the training cells for

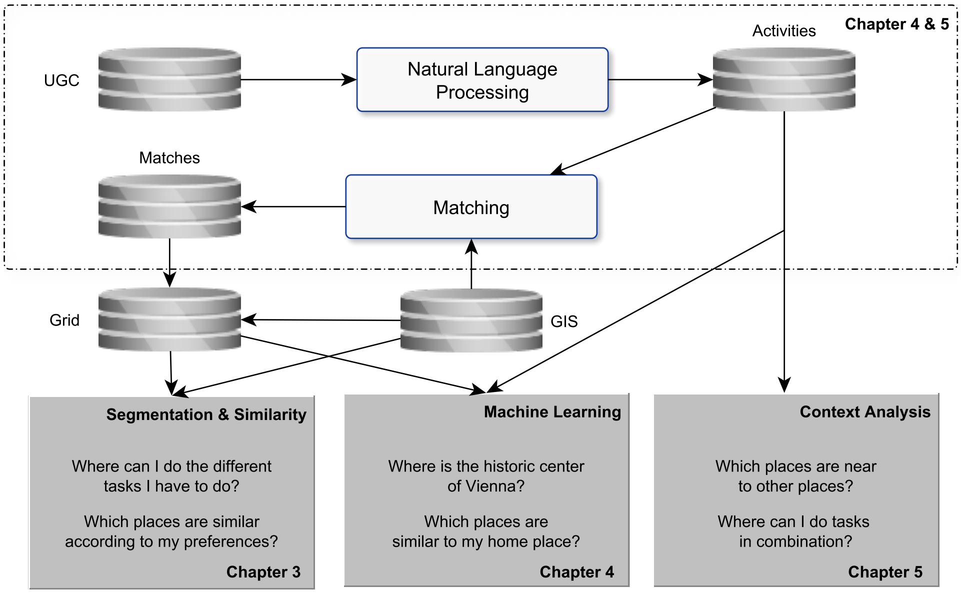

Figure 2.1: Conceptual illustration of the framework

automated reasoning, a place matching procedure is discussed that operates on textual place descriptions with the aid of Natural Language Processing. The Machine Learning model is a subject of Section 2.4 and of Chapter 4.

The Natural Language Processing approach is refined to automatically generate a model that allows questions tailored to spatial information needs. The proposed model will therefore preserve the spatial semantics, and the outlined algorithmic solution can automatically populate the model based on User Generated Content. Based on the information contained in the sources, a user can ask for implicit information, such as which places are near to other places, or implicit classifications, such as which spatial objects are related to different activities. The context analysis of User Generated Content is described in Section 2.5 and in Chapter 5.

2.3 Segmentation & Similarity

Generalization is one of the fundamental research areas of geographic information science [Hobel and Frank, 2014, Hobel et al., 2014]. The goal is to derive an abstraction of the physical reality and provide methods for reasoning. How to model place according to the conceptualization of human beings is a non-trivial question. A simplified model can be seen in analogy to an image, where different functional areas of a city are reflected as objects on a map, e.g. a shopping street or a park. While parks are often obviously characterized as green spots on a map, others, such as shopping areas, are not as easily perceivable.

To mimic human decision making, the model discussed in this thesis makes use of aggregated places (i.e. services an area provides) as the deciding characteristic for the partitioning of space. To answer questions such as "where can I do different tasks in combination?" and questions of tourists asking about the similarity of "activity clusters", an approach adapted from image segmentation is pursued in order to cluster a city map into cognitive regions. Image segmentation facilitates the process of analyzing an image by partitioning it into segments that share certain characteristics [Pinoli, 2014]. In analogy to traditional image segmentation, Hobel et al. [2015] argued that place affordances can be utilized to classify a city into functional regions. Partitioning a map into functional regions subsequently allows the application of distance functions to compare and measure how suitable an area is for certain tasks. The applicability of metrics and similarity functions of traditional Information Retrieval is investigated, and their relation to human preference models is discussed.

Example 2.3.1: Partitioning of City Maps into Cognitive Regions

Once homogeneous areas have been identified, a formal description of the area makes it possible to search, compare, or cluster such regions. Figure (a) depicts three conceptual shopping settings together with their respective frequency distribution of tags. While each one of them contains shops, places like parking spots and restaurants are also part of the constellation "shopping area". Figure (b) shows a schematic visualization of these cognitive regions as grid-based clusters.

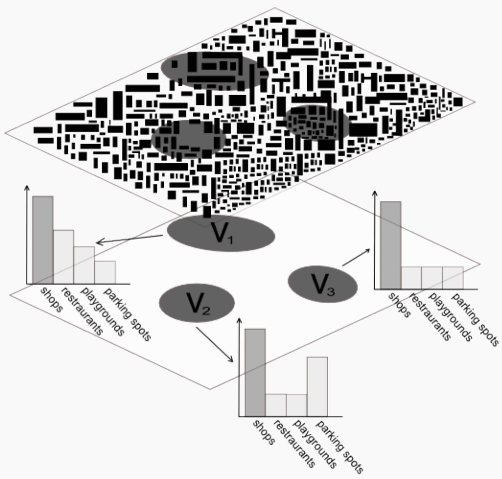

(a) Schematic Visualization

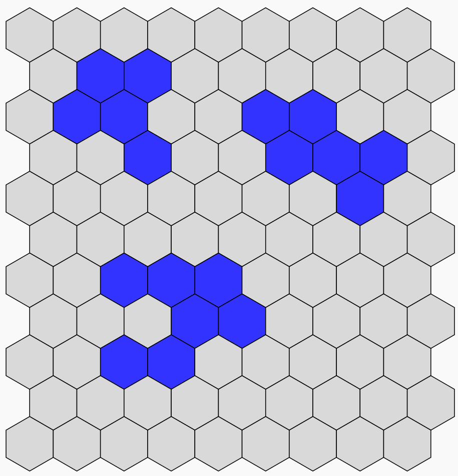
(b) Segmentation

2.4 Machine Learning

Hobel et al. [2016] extended the idea of clustering cities into functional regions by an automated approach, which exploits User Generated Content in the form of textual descriptions in order to capture human conceptualization of given areas, and a Machine Learning model in order to derive homogeneous and uniform areas according to the services a region offers.

Example 2.4.1: Deriving the Geographic Footprint of Cognitive Regions

To give an illustrative example, the geometric extent of the historic center of Vienna was derived by a Machine Learning model where the distribution of services is the underlying foundation for the semantic similarity of conceptual areas. Figure 4.9 approximately shows the encompassed area of a historic map of Vienna (from 1850). The outer boundary of the classified area coincides with a physical separation which is now a major road of the city.

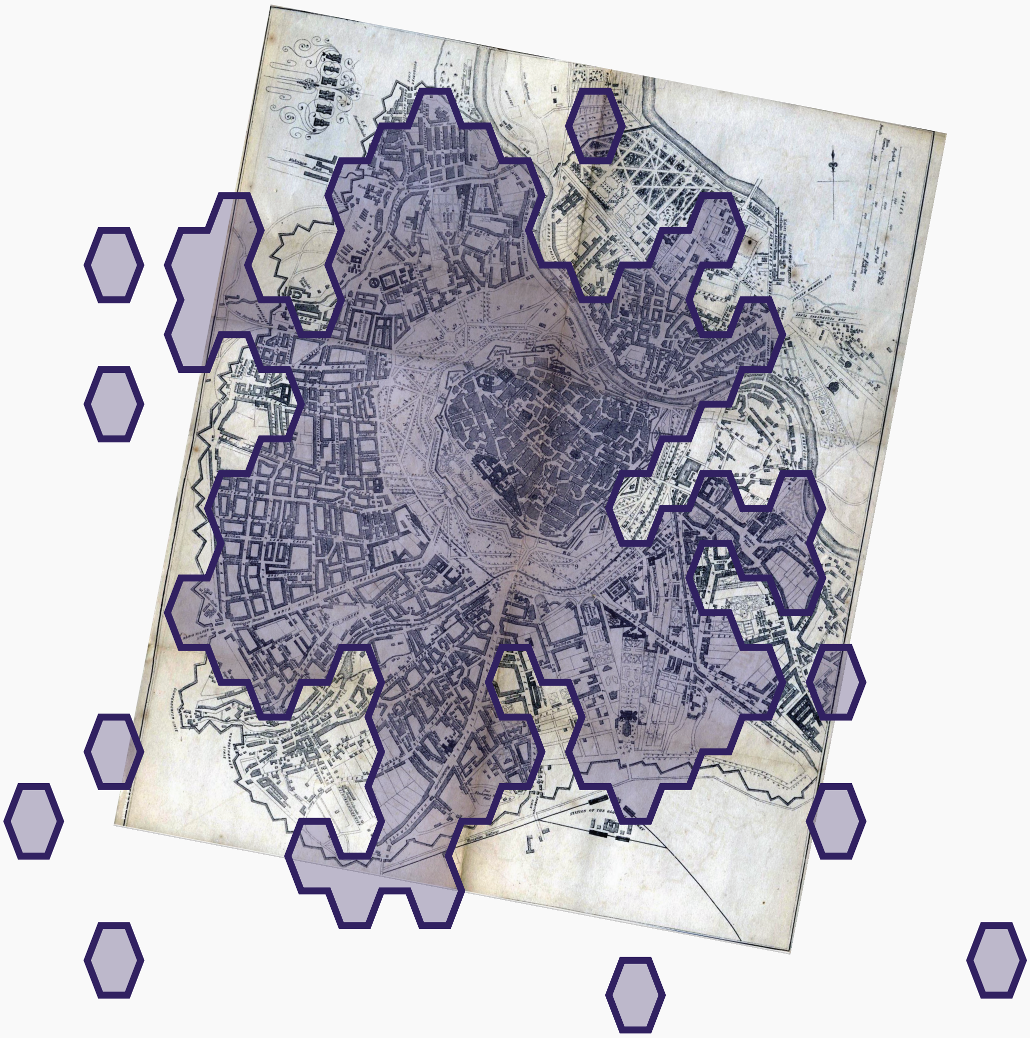

In order to derive the initial translation from textual descriptions expressed in natural language to the bag-of-words model, a place matching algorithm is discussed. The algorithmic solution is tailored to the OpenStreetMap dataset, and utilizes methods of Natural Language Processing. In the proposed process, Hobel et al. [2016] automatically extracted the places people mentioned in textual descriptions, and trained a classifier based on a grid-based clustering approach.

The proposed approach can be used for answering questions such as inferring the geometric extent of cognitive regions automatically, and questions such as "where are cognitive regions that correspond to the area I am familiar with?".

The introduced approach has significant implications for the next generation of routing services. It is comparable to today's indoor navigation, where cognitive areas are defined to ease the routing tasks of people. For instance, colored areas and navigation lines are defined, and routing instructions adhere to the functional properties which are available. With the presented approach, it will be possible to derive navigation and annotated mobile maps incorporating the semantic strength of cognitive regions for cognitively enhanced outdoor navigation. For example, the instruction "cross the historic center, and then keep right" is a more intuitively understandable formalization for outdoor navigation than a simple sequence of street names and crossings.

2.5 Context Analysis

Agnew [1987] defined place as a combination of three elements: location, locale, and sense of place. Of course, the sense of place can only be inferred from human descriptions. To that end, Hobel and Fogliaroni [2016] introduced a model, called the semantic representation of place in the Resource Description Framework, and an algorithmic solution to populate the ontological model based on User Generated Content by exploiting modern Natural Language Processing tools.

The introduced model complements the previously discussed approaches by adding the sense of place. The semantic representation of place unites the cognitive approach of place with affordance theory [Gibson, 1979, 1977] resulting in a coherent model for spatial search. In this thesis, it is argued that the semantic representation of place is suitable to answer spatial questions provided in natural language by processing the input in the same manner as User Generated Content. This means that we have a procedure of encoding User Generated Content into the semantic representation of place (i.e. activities), and spatial queries given in the form of natural language can be processed with the same procedure, resulting in a graph representation of the query. We have therefore reduced the problem of answering the query posed in natural language to the problem of matching two graphs. For simplification, the semantic similarity of verbs, context, and activities is disregarded.

The processed data gives an indication of human impressions, such as which geo-features are perceived as near to each other, and which activities are ascribed to which geo-features.

Although this has been addressed in literature before [Alazzawi et al., 2012], the presented approach is a novel one because it not only uses part-of-speech tags (e.g. verbs), but also exploits the grammatical structure (e.g. direct object) of sentences in the algorithmic solution.

Example 2.5.1: Illustration of Semantic Representation of Place

Obtained semantic representation for the place "Historic Center of Vienna" by processing the sentence "I even enjoyed walking down the beautiful Kärtnerstrasse admiring many nice, original shops".

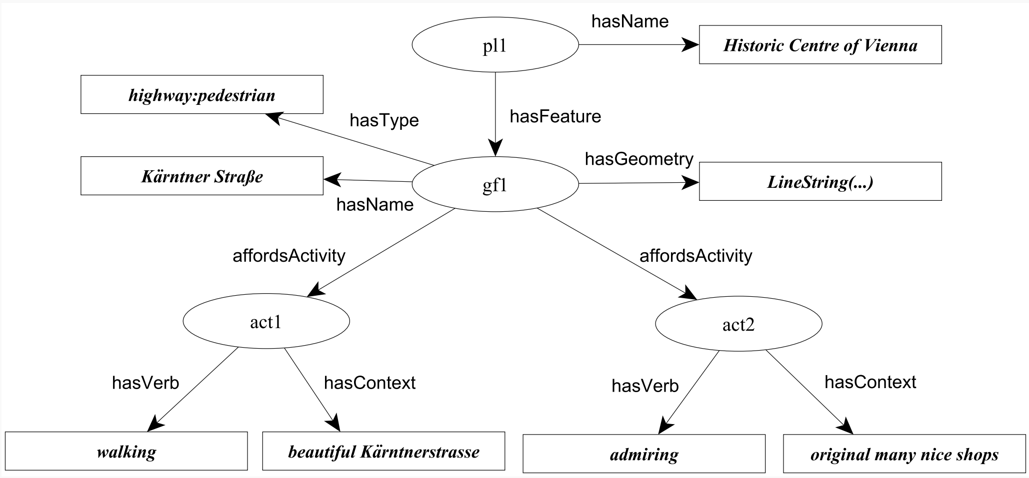

The created knowledge base also has beneficial implications for the technical realization of the previously introduced parts of the framework by complementing the proposed approaches. For the segmentation and similarity approach (see Section 2.3), the place affordances are manually mapped from the service to the actual activity, which can be done automatically with the proposed knowledge base. Another beneficial implication stems from the Machine Learning approach (see Section 2.4), as services (i.e. tags) that have no meaning for humans can be automatically excluded because no activities are mapped to the geo-feature.

2.6 Summary and Focus

So far, spatial search is insufficiently supported when dealing with questions about *place*. This thesis deals with techniques to equip Geographic Information Systems with the capability to process spatial queries in cognitively aware ways. Three approaches are proposed that facilitate spatial search by utilizing the semantics of places.

The first approach is inspired by image segmentation and Information Retrieval. The basic idea is to aggregate places according to the conceptualization of human beings, and to allow measurement of similarity based on distance functions.

The second approach enhances the first approach by utilizing place descriptions retrieved from User Generated Content and by computing cognitive regions, which are aggregations of places, based on Machine Learning. An analogy to traditional text classification is drawn to compare the semantic similarity of regions, and to cluster spatial areas according to semantic properties.

Finally, by further utilizing User Generated Content the models are extended to a full-fledged knowledge base that allows naïve question answering and reasoning. The model itself is based on affordance theory, and an algorithmic solution is proposed to populate the ontological model by human impressions retrieved from User Generated Content.

This chapter gave an overview of the investigated techniques and how they are combined to provide rich methodologies for the semantic modeling of spatial queries. To conclude, the focus of this thesis is the semantic representation of place. In the following chapters, the implementation details of the proposed approaches are elaborated in more detail. Each approach is evaluated based on use cases that illustrate the practical usefulness of the techniques applied.

CHAPTER 3

Segmentation & Information Retrieval

This chapter is primarily drawn from Hobel et al. [2015]. It discusses the cognitive aspects of a computational model for urban generalization, and also how cognitive homogeneous areas can be compared by means of Information Retrieval. The essential idea is to partition a map into homogeneous and uniform areas by focusing on the semantic aspects of space. It is shown that the derived regions can be compared by using their semantic description and human preference models.

3.1 Excursus

OpenStreetMap is a community-driven project whose main goal is to create a digital map of the entire world; it is essentially a prototype of Volunteered Geographic Information [Goodchild, 2007]. The *geometric footprint* of spatial features is represented by means of a simple and exceptionally flexible scheme consisting of

- *nodes*: pairs of coordinates (longitude and latitude) used to represent point features;
- *ways*: lists of nodes used to represent *line* and *surface* features;
- *relations*: sets of nodes, ways, or other relations mainly used to represent features consisting of several parts.

The thematic or semantic aspect of spatial features is managed through a tagging system, where each geometric feature is described by an arbitrary number of tags. As the OpenStreetMap project evolved and grew over time, its community developed a set of *tagging guidelines* which describe how tags should be used for a particular feature. Before contributing new information to the database, mappers are asked to carefully read these guidelines. However, users are neither obligated to respect the guidelines, nor are their contributions subject to rigorous control. It has been shown that in terms of geometry, the OpenStreetMap dataset is rapidly approaching the coverage and precision of commercial ones [Zielstra and Zipf, 2010]. The freedom granted by the tagging system yields a semantically very heterogeneous dataset [Mooney and Corcoran, 2012]. Consequently, different volunteers tag the same feature differently, or, conversely, use the same tag to annotate conceptually different features. Keßler and de Groot [2013], D'Antonio et al. [2014] show the possibility of assessing the trustworthiness of Volunteered Geographic Information data by analyzing the historical evolution of features in a dataset.

The ambiguous meaning of place poses a considerable challenge to knowledge engineers whose task is to design computational models of places. Today, the most commonly adopted strategy is to represent places by means of Points of Interests such as targeted by OpenStreetMap. The representation of places as Points of Interests, however, disregards many of the aspects that seem to characterize the human conceptualization of places: (*i*) there is empirical evidence [Montello et al., 2003] that people typically conceive a place as a region; (*ii*) different people tend to associate different spatial footprints to the same place [Montello et al., 2003]; (*iii*) there are indications [Schatzki, 1991, p.655] that conceptualization of a place relies on the activities available or possible at that spatial location – i.e. what some refer to as place affordances [Jordan et al., 1998]. The approach of representing places as Points of Interests suffers from several drawbacks: places are indicated as specific points rather than vague or approximated regions; while Points of Interest are associated with precise feature types, the place affordances are not explicitly indicated, and it is up to the user to map from an activity (e.g. to eat) to a feature type (e.g. a restaurant). Going even further and focusing our attention on activities, it is easy to see that activities are usually not restricted to a single place, and have an extent in space and time that involves several places of different kinds. Shopping, for example, can also involve sitting in a coffee shop, or going to a bank to withdraw money.

Humans are able to search for areas that *afford* [Gibson, 1977, 1979] an activity without having to specify the exact type of place they are looking for. For example, if the task is to "buy a pair of shoes and perhaps a coat", humans can, based on experience or knowledge, think of areas where they are most likely to find such things (e.g. a shopping street or shopping mall). In such a case, the individual shop is of less concern since the exact object to buy is not determined yet. Rather, it is the constellation or setting of shops and maybe restaurants that is of importance when attempting to find an area suitable for an activity, which corresponds to the cognitive notion of place.

The concept of *place* plays an increasingly important role in Geographic Information Science [Winter et al., 2009, Winter and Truelove, 2013] and the ontological discussion about how to model it is ongoing [Couclelis, 1992, Jones et al., 2001, Humayun and Schwering, 2012, Vasardani et al., 2013, Winter and Truelove, 2013]. Many suggest that the semantics of the term *place* is tightly bound to the idea of affordance and activities [Jordan et al., 1998, Scheider and Janowicz, 2014]. Drawing the connection of action to place is essential for the ability to plan [Abdalla and Frank, 2012]. Schatzki [1991] asserted that: "[...] places are defined by reference to human activity" [Schatzki, 1991, p.655]. He positions human activities as the central concept for understanding the construction of places. Furthermore, he explains that such representations of places organize into settings, local areas, and regions. This general notion of hierarchical structuring of space is relatively undisputed and supported by findings of other researchers [Montello, 1993, Couclelis and Gale, 1986, Richter et al., 2013, Freundschuh and Egenhofer, 1997]. How these levels of abstractions are formed, however, is unclear. For example, common administrative units of abstraction do not always correspond to what people have in mind about regions [Meegan and Mitchell, 2001].

The focus of this thesis lies on *settings* or cognitive regions which, according to Schatzki [1991], can either be demarcated by barriers (e.g. apartment building) or identified by bundles of activities that occur in them (e.g. a park, or a shopping street). Ontologically speaking, they can be categorized as entities of either bona fide (i.e. physical, sharp, crisp) or fiat (i.e. non-physical, imaginary, human-driven) type [Smith, 1995]. Since this thesis is concerned with entities larger than apartment buildings, such as shopping areas, fiat objects will be the the main type of inquiry. The entities are therefore of the vista-space scale [Montello, 1993], since they can be learned by human activity. In the following, settings are referred to as cognitive regions.

3.2 A Model for Cognitive Regions

City maps are cartographic representations of spatial data partitioning space into discrete chunks that represent physical or social (administrative) objects. These objects are either defined by their physical extent or by authoritative institutions. Schatzki [1991] asserted that there are places that fall into the same abstract category due to certain constellations of possible activities. The following approach is inspired by region growing used in image segmentation [Adams and Bischof, 1994].

A formalization of the proposed *semantic* region growing approach is summarized in the following steps:

1. The task is to partition a map according to the services an area offers. In analogy to the pixels of an image, the area of interest $\mathcal{M}$ (a city map in our case) is partitioned into a collection of n non–overlapping cells $C = \{c_i : i = 1, \ldots, n\}$ such that $\mathcal{M} = \bigcup_{i=1}^{n} c_i$. This method is analogous to spatial indexing strategies, where grid-based clustering is performed to hierarchically partition the space in uniform square or hexagonal grid cells.

2. A description D is a formula consisting of one or more predicates specifying the membership of a single cell c_i to a specified cognitive region R. An example of such a description could be: "contains at least one shop and restaurant".

3. What in image segmentation jargon is called a *segment*, is directly comparable to a cognitive region: a set of contiguous cells satisfying the same description D. A cognitive region $R \subseteq C$ is a subset of the cell partition C and is called *complete* iff it cannot be extended further with adjacent cells.

4. The segmentation of a map $\mathcal{M}$ according to a description D produces a (possibly empty) set $\mathcal{R}_D$ of cognitive regions such that by construction $\bigcup_{R \in \mathcal{R}_D} R \subseteq C$ holds. A segmentation $\mathcal{R}_D = \{R\}$ is called *complete* iff it consists of only one cognitive region such that $R = C$.

5. As image segmentation relies on a similarity function that is used to decide if two neighbor pixels are similar, our approach consequently relies on a Boolean function f_{sim} which, given a cell c and a description D, decides whether c adheres to D. Different cognitive regions identified through the same description D are pairwise disjoint, i.e. it holds that $\forall x, y \in \mathcal{R}_D \, (x \neq y \implies R_x \cap R_y = \emptyset)$. Cognitive regions that adhere to different descriptions can overlap, e.g. a park that crosses a shopping street.

3.3 Semantic Region Growing

Semantic region growing, as used here, is aimed at segmenting or extracting cognitive regions according to a description D and a set of m cells, referred to as *seeding cells* $C_{seed} = \{\tilde{c}_1, \tilde{c}_2, \ldots, \tilde{c}_m\}$.

In the case that a seeding cell $\tilde{c} \in C_{seed}$ matches a given description D – i.e., $f_{sim}(\tilde{c}, D) = TRUE$ – and it is not yet classified as a member of another cognitive region adhering to the same description D, $\tilde{c}$ will be the starting point of a new cognitive region: a recursive process extends the region identified so far until the elements of the adjacent neighborhood of a cell no longer adhere to the description D.

We can either process all cells as seeding cells ($C_{seed} = C$), or find all cells in C that adhere to the description D, and use them as C_{seed} – both cases yield a robust result in contrast to random seed generation. For instance, if C_{seed} contains only five seeding cells, then the result will be at most five segments/cognitive regions. Note that a cognitive region R may not be identified by the algorithm if $R \cap C_{seed} = \emptyset$, depending on the neighborhood function.

It is possible that, during the growing process, starting from a seeding cell $\tilde{c}_i$ and building a cognitive region R_i, another seeding cell $\tilde{c}_j$ is integrated in R_i. When the algorithm processes the seeding cell $\tilde{c}_j$, it will not give rise to a new cognitive region since it has already been assigned to the cognitive region R_i. The *semantic region growing* technique is implemented as shown in Algorithms 1 and 2.

As can be seen in the implementation of Algorithm 2, the size of the adjacent neighborhood of a cell can be adapted by using a customized implementation of $neighbors(c, C)$ to specify requirements such as larger or restricted neighborhoods. In any case, a larger

Algorithm 1 Identify cognitive regions.

Input

D = *a predicate to segment the map into cognitive regions*,
C_{seed} = *the set of seeding cells*,
C = *the set of cells*

Output

$\mathcal{R}$ = *cognitive regions*

procedure IdentifyCognitiveRegions
  $\mathcal{R} \leftarrow \emptyset$
  for all $\tilde{c} \in C_{seed}$ **do**
    if $(\forall R_i \in \mathcal{R} \colon \tilde{c} \notin R_i) \wedge f_{sim}(\tilde{c},\ D)$ **then**
      $R \leftarrow \{\tilde{c}\}$
      $\mathcal{R} \leftarrow \mathcal{R} \cup \{R\}$
      $RegionGrowing(D, \tilde{c}, R, \mathcal{R}, C)$
    end if
  end for
  return $\mathcal{R}$
end procedure

Algorithm 2 Region Growing.

Input

D = *a predicate to segment the map into cognitive regions*,
c = *a cell*,
R = *the current cognitive region that is extended*,
$\mathcal{R}$ = *the set of cognitive regions*,
C = *the set of cells*

Output

R = *the extended cognitive region*

procedure RegionGrowing
  $N \leftarrow neighbors(c, C)$
  for all $n \in N$ **do**
    if $(\forall R_i \in \mathcal{R} : n \notin R_i) \wedge f_{sim}(n,\ D)$ **then**
      $R \leftarrow R \cup \{n\}$
      $RegionGrowing(D, n, R, \mathcal{R}, C)$
    end if
  end for
  return R
end procedure

neighborhood can be used to ensure better coverage, or restrictions can be imposed to separate cognitive regions.

Example 3.3.1: Partitioning of City Maps and Cognitive Regions

Figure (a) shows an empty segmentation and Figure (b) shows a complete segmentation. Figure (c) shows a partitioning into cognitive regions of a certain type.

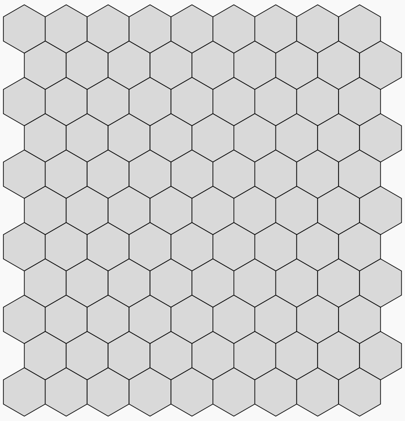

(a) Empty Segmentation

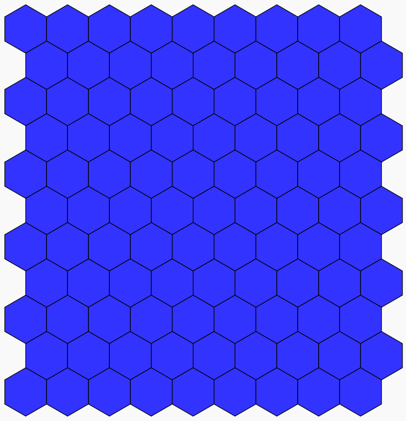

(b) Complete Segmentation

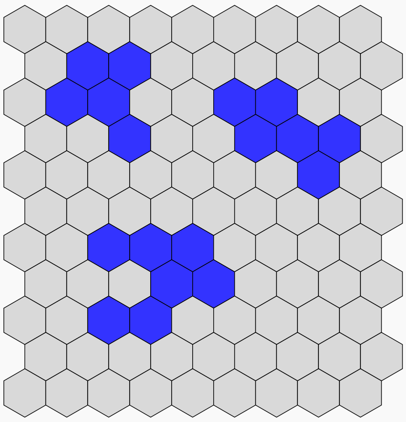

(c) Segments

3.4 Application Areas

In this section, two application areas of this hierarchical segmentation in functional regions are discussed.

3.4.1 Similarity in the Vector Space Model

The categorical types, i.e. services, can be modeled in the vector space or the bag-of-words model. The vocabulary is the set of all categorical tags occurring in the city, $\mathcal{T} := \{t_i \colon i = 1, \ldots, n\}$, where n denotes the size of the vocabulary. Specifically, each cell in the bag-of-words model is represented by the vector $(x_1, \ldots, x_n)$, where x_i is the frequency of the categorical tag t_i in this cell. The bag-of-words model yields a simplified 'semantic representation' of cognitive areas.

One particularly interesting question is how to compare different cognitive regions. Places as aggregates can be modeled by using the aggregated semantic attributes, which enables comparison and assessment of cognitive regions. Such a composed area offers not only single place affordances but rather encompasses a set of affordances which are seamlessly interconnected. In the case of OpenStreetMap, the aggregated categorical tags in cells offer a baseline for those affordances, i.e. services. However, before going into further analysis, tags that can not be mapped to a specific place affordance, e.g. a tag with the key "ref", have to be removed from the set of categorical tags.

In order to simulate human decision-making, artificial intelligent systems developed in this thesis can make use of distance functions to derive closest or best matches of different cognitive regions. So far, we have presented a reduction of cognitive regions to the vector space model. However, a priori there is no canonical way to define a metric on these vectors which exactly mirrors the cognitive reasoning process mentioned before. In mathematics, many different ways to define the "metric concept" have been proposed. In the following examples, we start with the notion of a metric space:

Definition 3.4.1: Metric Space

A **metric space** is an ordered pair (A, d) where A is a set and d is a metric on A, i.e. a function

$$d \colon A \times A \to \mathbb{R} \tag{3.1}$$

such that for any $r_1, r_2, r_3 \in A$, the following holds:

- $d(r_1, r_2) \geq 0$
- $d(r_1, r_2) = 0 \iff r_1 = r_2$
- $d(r_1, r_2) = d(r_2, r_1)$
- $d(r_1, r_3) \leq d(r_1, r_2) + d(r_2, r_3)$

The following is a list of distance function examples that adhere to the requirements of the metric space:

Definition 3.4.2: Discrete Metric

Discrete metric $(r_1, r_2 \in A)$:

$$d_{Discrete}(r_1, r_2) := \begin{cases} 0 & \text{, if } r_1 = r_2 \\ 1 & \text{, if } r_1 \neq r_2 \end{cases} \tag{3.2}$$

Definition 3.4.3: Euclidean Metric

Euclidean distance $(r_1 = (r_{1_1}, r_{1_2}, \ldots, r_{1_n}), r_2 = (r_{2_1}, r_{2_2}, \ldots, r_{2_n}) \in \mathbb{R}^n)$:

$$d(r_1, r_2) := \sqrt{\sum_{i=1}^{n} |r_{1_i} - r_{2_i}|^2} \tag{3.3}$$

Definition 3.4.4: Manhattan Metric

Manhattan distance $(r_1 = (r_{1_1}, r_{1_2}, \ldots, r_{1_n}), r_2 = (r_{2_1}, r_{2_2}, \ldots, r_{2_n}) \in \mathbb{R}^n)$:

$$d(r_1, r_2) := \sum_{i=1}^{n} |r_{1_i} - r_{2_i}| \tag{3.4}$$

Definition 3.4.5: Hamming Metric

Hamming distance $(r_1 = (r_{1_1}, r_{1_2}, \ldots, r_{1_n}), r_2 = (r_{2_1}, r_{2_2}, \ldots, r_{2_n}) \in \mathbb{R}^n)$:

$$d(r_1, r_2) := \sum_{i=1}^{n} d_{Discrete}(r_{1_i}, r_{2_i}) = |\{i \in \{1, \ldots, n\} \mid r_{1_i} \neq r_{2_i}\}| \tag{3.5}$$

We argue that urban search is directly comparable to traditional Information Retrieval, where cognitive regions correspond to documents containing the categorical tags as words.

In 1957, Luhn [1957] suggested a statistical approach to search for information with the similarity criterion:

> "The more two representations agreed in given elements and their distribution, the higher would be the probability of their representing similarity information."

Therefore, the degree of similarity between a representation of a fixed cognitive region and a collection of cognitive regions can be used to rank the search results. This can be interpreted as counting the number of elements that two cognitive regions share. For instance, let the fixed cognitive region's representation be a vector $\vec{r_1} = (r_{1_1}, r_{1_2}, \ldots r_{2_n})$ of which each component r_i $(1 \leq i \leq n)$ is associated with a tag and a cognitive region returned by a search is a similar vector $\vec{r_2} = (r_{2_1}, r_{2_2}, \ldots, r_{2_n})$, then a similarity measure can be based on the standard inner product of two vectors:

Definition 3.4.6: Standard Inner Product

$$\iota(r_1, r_2) := \sum_{i=1}^{n} r_{1_i} \cdot r_{2_i} \tag{3.6}$$

Example 3.4.1: Cosine Similarity

The cosine similarity measures the similarity between two vectors of an inner product space that measures the cosine of the angle between them. A schematic illustration of the similarity of two documents r_1 and r_2 is given below:

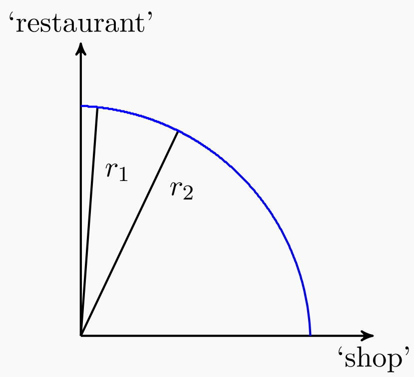

In 1983, Salton et al. [1983] proposed that comparisons can be made by the cosine-similarity function, which does not adhere to the mathematical concept of metric space (see Definition 2.5.1), and is defined as follows:

Definition 3.4.7: Cosine-Similarity

$$sim_{\mathbf{cos}}(r_1, r_2) := \cos(\sphericalangle(r_1, r_2)) = \frac{\sum_{i=1}^{n} r_{1_i} \cdot r_{2_i}}{\sqrt{\sum_{i=1}^{n} r_{1_i}^2} \cdot \sqrt{\sum_{i=1}^{n} r_{2_i}^2}} \tag{3.7}$$

In the scope of this thesis, these similarity measures are important building blocks of artificial intelligent systems used to answer more complex questions than what is possible today. For instance, today's geographic search engines cannot answer which areas are similar to shopping streets, or which areas are similar to a familiar area in my surroundings. By using similarity functions, these comparisons can be made by intelligent systems, and facilitate modern Geographic Information Retrieval.

3.4.2 Human Preference Models and Cognitive Regions

Human search depends on different aspects of personal search criteria, which can be incorporated in addition to the already introduced distance metrics.

Fine-grained or significant differences in place constellations can reveal how suitable the composition of a cognitive region is for someone's preferences, or they can be used to identify flaws in the naturally evolved or planned structure of a city. For instance, when people are required to travel by car due to an inefficient public transport system, or out of personal necessity, shopping areas with parking spots are certainly more attractive destinations. It follows that cities without dedicated parking spots in the vicinity of shopping areas must have an efficient public transport system.

Human search is characterized by specific needs, which are the activities required of or by the human being in question. To mimic human decision making, the preferences must be integrated into spatial search systems. This can be done by formalizing spatial needs as a list of requirements $n_1, n_2, \ldots, n_m$, where m is the total number of required place affordances. Then $\tau_{(R)}$ is denoted as the total number of cells in a given segment/cognitive region R. We set the absolute values $n_1, n_2, \ldots, n_m$, which we defined in the list above, in relation to the area of the cognitive region, which yields normalized density values:

$$a_i = \frac{n_i}{\tau_{(R)}} \; \forall i = 1, \ldots, m \tag{3.8}$$

Therefore, based on an ontology of specific place affordances and human preferences, a matching of normalized ranking criteria can be defined. A distance metric listed above can then be selected based on the requirements of the user.

3.5 Case Study

For an evaluation of our approach, we attempted to identify shopping areas in two cities. By using GeoTools[1], an open source library for geospatial data, we set up a fine-grained hexagonal grid, whereby the side length of the cells was set to 0.0005 degrees, and preprocessed the OpenStreetMap data by assigning the nodes and their tags to the enclosing cells.

3.5.1 Identification of Cognitive Regions

In this example, the rules for the description of the semantic region growing algorithm are based on the following assumptions:

A cell has to encompass at least two places where you can *shop* (i.e. shops of any type) OR a cell has to encompass at least two tags that relate to places where someone can get something to eat or drink a coffee (e.g. restaurants, fast food outlets, cafes).

These simple constraints were sufficient to find the commonly known *shopping areas* in Vienna, and many smaller clusters that can be interpreted as local shopping and leisure areas. The segmentation result for Vienna is shown in Figure 3.1. Using the same description, we employed the algorithm on the dataset of London, and obtained a comparable result (see Figure 3.2). By selecting a slightly different description, different cognitive regions are derived, which is coherent with results from different studies such as the experimental study to infer the location of downtown Santa Barbara [Montello et al., 2003]. The introduced approach allows us to derive human conceptualizations of cognitive regions. The derived concepts have clearly defined boundaries and have uniform and homogeneous characteristics.

Arguably, there is no *hard* method to evaluate the result, since the topics of interest are cognitive regions which do not really allow for a ground truth. Nevertheless, an estimation of feasibility is still possible, either by looking at descriptions found on the internet (e.g. tourism guides) or by comparing the results to expert knowledge (i.e. people familiar with the city). Indeed, Mariahilferstraße and Oxford Street are well-known shopping streets that have been correctly identified as part of shopping settings by our algorithm. Additionally, detailed explorations of some other clusters identified in the Vienna dataset consistently revealed that all larger regions found can be considered shopping areas.

3.5.2 Semantic Similarity of Cities

Consider the following scenario:

Alice grew up in London and she knows from experience that in the *urban setting* of Oxford Street there are plenty of places to withdraw money (i.e. ATMs and banks), that there is a wide selection of cafes and restaurants to have lunch or get something to drink, and also that there is a large diversity of shops and several tourism attractions that can

[1] `http://docs.geotools.org/`

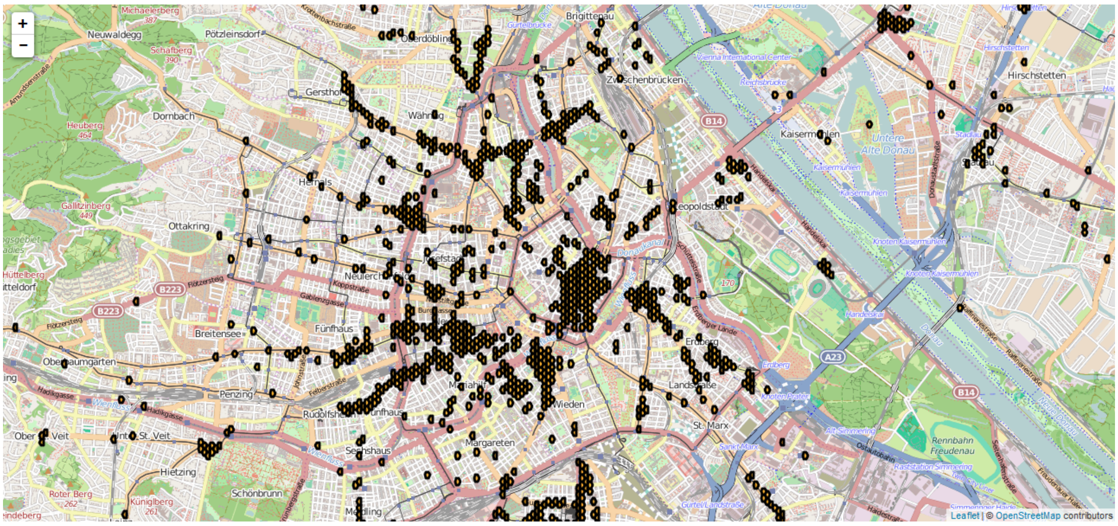

Figure 3.1: Visualization of the results identified by the semantic region growing algorithm in Vienna.

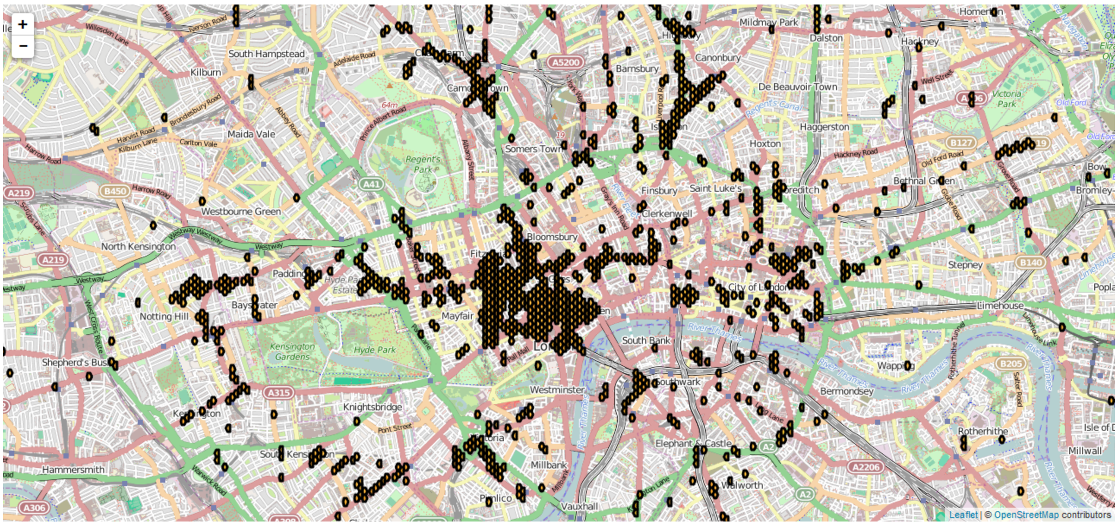

Figure 3.2: Visualization of the results identified by the semantic region growing algorithm in London.

be visited while moving from one shop to the next. Alice plans a trip to Vienna, and she would like to find, in advance, areas of Vienna that are similar to her idea of Oxford Street.

	B/A	Tourism	R/C/FF	Shop	Absolute
Inner City	0.31	1.28	2.60	0.81	5.02
Mariahilferstraße	0.10	0.62	6.39	2.28	9.41

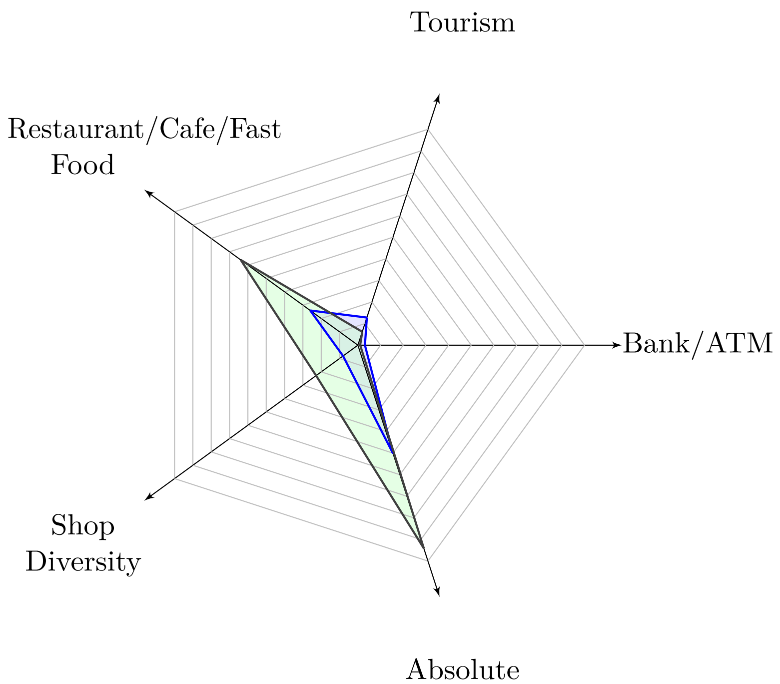

Figure 3.3: The deviations of the conceptual shopping area Oxford Street (London) to the conceptual Mariahilferstraße and Inner City (Vienna) in respect to the defined feature vector (values are multiplied by 10).

To model these preferences and action possibilities, we defined the following four features, which will later be used to define a similarity distance to other identified shopping areas (cognitive regions):

1. The number of tags in a setting of type bank or ATM n_1.
2. The number of tags in a setting of type restaurant or cafe or fast food n_2.
3. The number of tags in a setting of type tourism n_3.

4. The number of different shopping types (i.e. subcategories of shops) n_4.

We denote by $\tau_{(S)}$ the total number of cells in a given segment/setting S. We set the absolute values n_1, n_2, n_3, and n_4, which we defined in the list above, in relation to the area of the setting, which yields normalized density values (where m is the number of defined features):

$$r_i = \frac{n_i}{\tau_{(S)}} \; \forall i = 1, \ldots, m \tag{3.9}$$

To explore the similarity in respect to our defined feature vector, we are now considering the following distance measure

$$\sum_{i=1}^{m} |r_i^{(S_1)} - r_i^{(S_2)}| \tag{3.10}$$

Equation (3.10) formalizes the sum of the absolute values of the differences between corresponding features for two settings with normalized values $r_{(\cdot)}^{(S_1)}$ and $r_{(\cdot)}^{(S_2)}$, and as such corresponds to the Manhattan distance.

Based on the use case scenario outlined above, Alice wants to find shopping areas in Vienna that compare in similarity to London's Oxford Street. Therefore, we denote by $r_i^{(S_1)}$ the values of Oxford Street and make a comparison with the larger-sized extracted conceptual shopping settings of Vienna, since we normalized the data based on the size of the settings. According to the total deviation (see Eq. (3.10)) the best matching setting is the area found around the *Inner City*, and the second best is the cluster around the lower part of *Mariahilferstraße*, which is illustrated in Figure 3.4.

Figure 3.3 illustrates the deviations of the defined preferences between the areas of *Inner City* and *Oxford Street* (blue), as well as *Marahilferstraße* and *Oxford Street* (green). The total deviation, which is defined through the similarity distance given in Eq. (3.10), can be read off the *absolute deviation* axis. A lower deviation is indicated when the instance that is being compared, i.e. the line for *Inner City* or *Mariahilferstraße*, respectively, is nearer to the center. In this case, it can clearly be seen that the total deviation of *Inner City* is lower than the deviation of *Mariahilferstraße*, as compared to the original *Oxford Street* area. To enable a more fine-grained comparison, we plotted for each $i = 1, \ldots, 4$ the value of $|r_i^{(S_1)} - r_i^{(S_j)}|$, which is the single deviation on an independent axis. In the previous formula, the variable j stands for either 2 or 3, which correspond to *Inner City* and *Mariahilferstraße*, respectively.

In the following, we briefly elaborate on the individual feature differences according to Figure 3.3:

1. Regarding the density of banks and ATMs, the area found in the *Inner City* as well as the one around *Mariahilferstraße* are both relatively close to the area that contains *Oxford Street*.

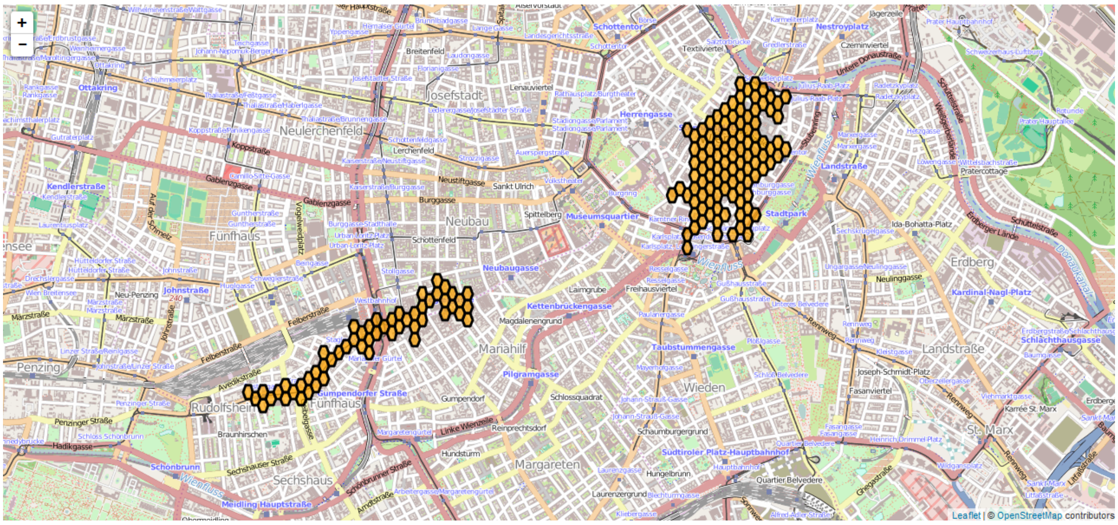

Figure 3.4: Visualization of the identified shopping areas in Vienna, which are most *similar* to the Oxford Street (London).

2. In terms of the density of tourism attractions, *Mariahilferstraße* is a bit closer to *Oxford Street* than the *Inner City*.
3. The higher deviations in terms of density of restaurants, cafes, and fast food outlets, and in terms of shop diversity of *Mariahilferstraße* indicate that the *Inner City* is more similar to *Oxford Street*.

3.6 Summary and Focus

In this chapter, a generalization procedure has been defined that subsequently allows for more semantic aspects to be considered than is possible with current spatial search engines. We denote homogeneous areas as cognitive regions or settings, which are simplified "semantic representations" and relate to human conceptualization of space.

In the beginning of this chapter, we argued that conceptualization of space is similar to image segmentation and proposed a method inspired by a technique called region growing. The proposed method allows us to partition a map into homogeneous and conceptually uniform areas, mimicking human decision making by segmenting a map into functional areas.

In order to facilitate the comparison of cities with respect to certain functional areas, distance and ranking mechanisms are discussed. Human requirements and preferences are modeled as spatial search criteria and integrated into the proposed computational model of cognitive regions.

In this chapter, two significant assumptions have been made: (i) first, seeding cells have to be known; and (ii) second, a description D for the regions is available. In the next chapter, we will try to mitigate these issues by utilizing textual place descriptions and Machine Learning to automatically derive conceptually uniform areas.

CHAPTER 4

Deriving the Geographic Footprint of Cognitive Regions

In the previous chapter, we discussed cognitive aspects of urban search. Cognitive regions are highlighted as the pillars of city generalization, where methods of traditional Information Retrieval serves as foundation in respect to similarity of regions. In this chapter, the basic idea of cognitive regions is integrated in an automated approach to extract conceptual areas by the use of Natural Language Processing and Machine Learning. This chapter is mainly drawn from Hobel et al. [2016].

4.1 Excursus

In this chapter, methods of Information Retrieval, Natural Language Processing, and Machine Learning are utilized to both identify seeding cells for cognitive regions from User Generated Content, and to advance the identification of homogeneous and uniform areas. This section, therefore, addresses two research directions: term extraction, and probabilistic classification.

Geographic Information Retrieval is a specialization of traditional Information Retrieval supported by geographic knowledge bases that enables the retrieval of geographic information and geotagged objects. The respective tools enable the identification and disambiguation of place names, the mapping of place names onto spatial features and vice versa, and the derivation of place semantics. Regarding the latter, the literature primarily focuses on the identification and classification of places [Tversky and Hemenway, 1983, Smith and Mark, 2001], and on the automatic generation of ontologies [Popescu et al., 2008].

To enhance the capabilities of the next generation of geographic search engines, different approaches are currently being pursued to facilitate the retrieval of geo-related content.

Applications range from the conceptualization of space into a metric space algebra [Adams and Raubal, 2009], to the contextualization of unstructured text [Adams et al., 2015, Adams and McKenzie, 2012] to relate concepts to places, to the development of content-rich knowledge bases and vocabularies [Ballatore et al., 2015], and to semantic similarity measures for geographic terms [Ballatore et al., 2013].

Interesting approaches of automatically mapping spatial content are pursued in different fields. Jones et al. [2008] focused on modeling vague regions by statistical density surfaces, and mining place descriptions in natural language to infer the approximate region. Grothe and Schaab [2009] exploited freely available georeferenced photographs to derive the geographic footprint of imprecise regions by using Kernel Density Estimation and Support Vector Machines. Cunha and Martins [2014] derived imprecise regions by exploiting Machine Learning for interpolating from a set of point locations. Lüscher and Weibel [2013] concentrated on using characteristics of topographical data to delineate regions.

The current focus on similarity measures for geographic terms [Ballatore et al., 2015, 2014] is further proof that there is an interest in the disambiguation of places and place descriptions. One of the goals is to prepare shared and universally accepted vocabularies to facilitate the interpretation and the resolution of spatial requests. For instance, if the task is to search for a place where "one can get something to eat", there are more possible matches than just restaurants. Coffee shops, pubs, or even supermarkets may also fulfill the requirements of the request.

The availability of mature Natural Language Processing tools [Manning et al., 2014] allows for advanced processing of textual spatial descriptions [Chang et al., 2015, 2014a,b, Coyne and Sproat, 2001] where tokenization and part-of-speech taggers are used to automatically break text into meaningful symbols – a selection of Part-of-Speech Tags (POST) is shown in Table 4.1. Two recent interesting approaches are presented in [Alazzawi et al., 2012] and [McKenzie et al., 2013]. The former builds upon current state-of-the-art Natural Language Processing to extract spatial activities from unstructured text; the latter presents a model to derive user similarity from spatial topics they discuss on social media.

Table 4.1: A Selection of Part-of-Speech tags (POST) [Santorini, 1990]

POST tag	Definition	POST tag	Definition
CC	coordinating conjunction	DT	determiner
IN	preposition or subordinating conjunction	WRB	adverb
JJ	adjective	WP	pronoun
NN	noun, singular or mass	TO	to
NNP	proper noun, singular	VB	verb

4.2 Framework

In the following, we outline a processing workflow (see Figure 4.1) to derive the geographic footprint of a given cognitive region from textual descriptions of that region. Details of the single steps involved are given in further subsections.

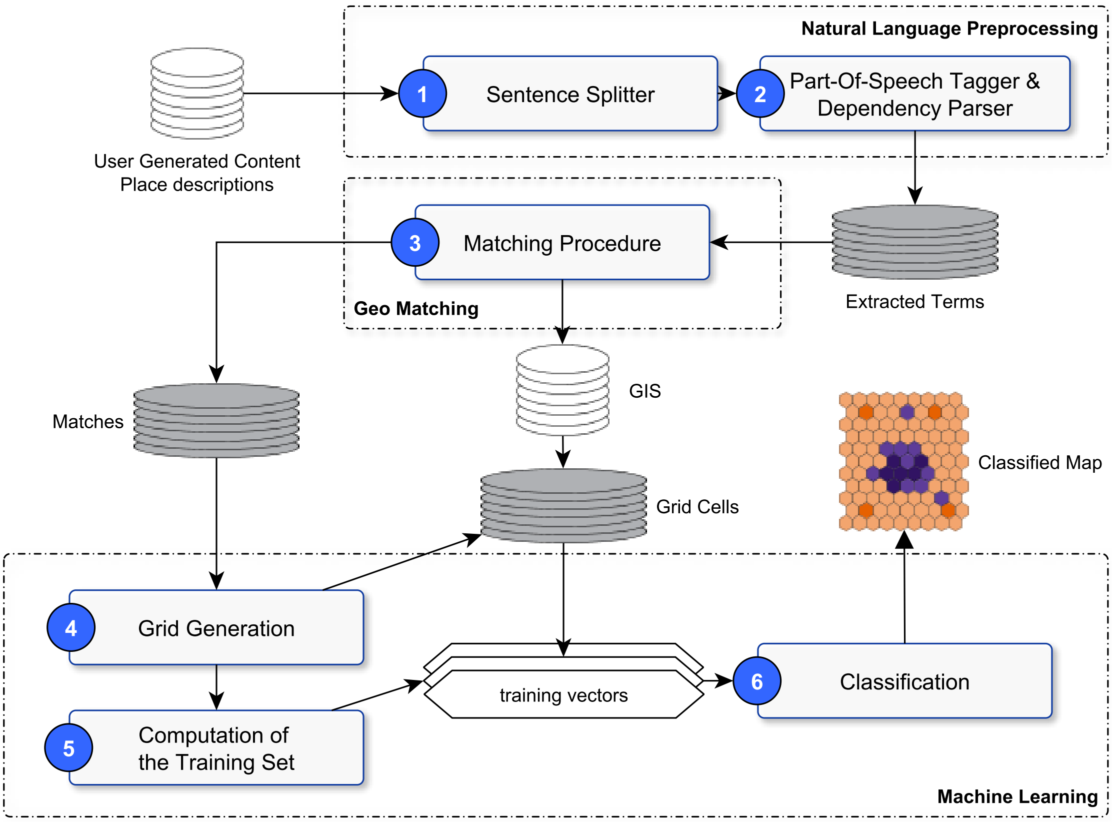

Figure 4.1: Schematic illustration of the proposed workflow to derive the geographic footprint of cognitive regions.

The proposed approach relies on two types of data sources (depicted as white databases in Figure 4.1): (*i*) a User Generated Content database containing textual descriptions of a given cognitive region, and (*ii*) a geographic database that might be either a Geographic Information System or a Volunteered Geographic Information system.

The workflow consists of three main stages labeled in Figure 4.1 as *Natural Language Preprocessing*, *Geo Matching*, and *Machine Learning*, respectively. First, the textual descriptions undergo a Natural Language Processing phase in order to extract from them a set of nouns referring to geographic features. In the next step, this set is compared to the geonames available in the spatial database in order to assign each a location on the map. Finally, a grid of regular cells is superimposed onto the map and the cells containing at least one of the geographic features mentioned in the textual descriptions are selected. These, together with a different set of cells selected randomly from the grid as counterexamples, are used as training samples for a Machine Learning model that categorizes all other cells according to the activities they allow. As a result, each cell is associated to either of the two training sets, unless too little information is known about it – in which case it is marked as "unclassified".

4.2.1 Natural Language Preprocessing

The Natural Language Processing stage relies on the Stanford CoreNLP Natural Language Processing Toolkit [Manning et al., 2014]. More specifically, it relies on three of the tools provided: the *sentence splitter*, the *part-of-speech tagger*, and the *dependency parser*.

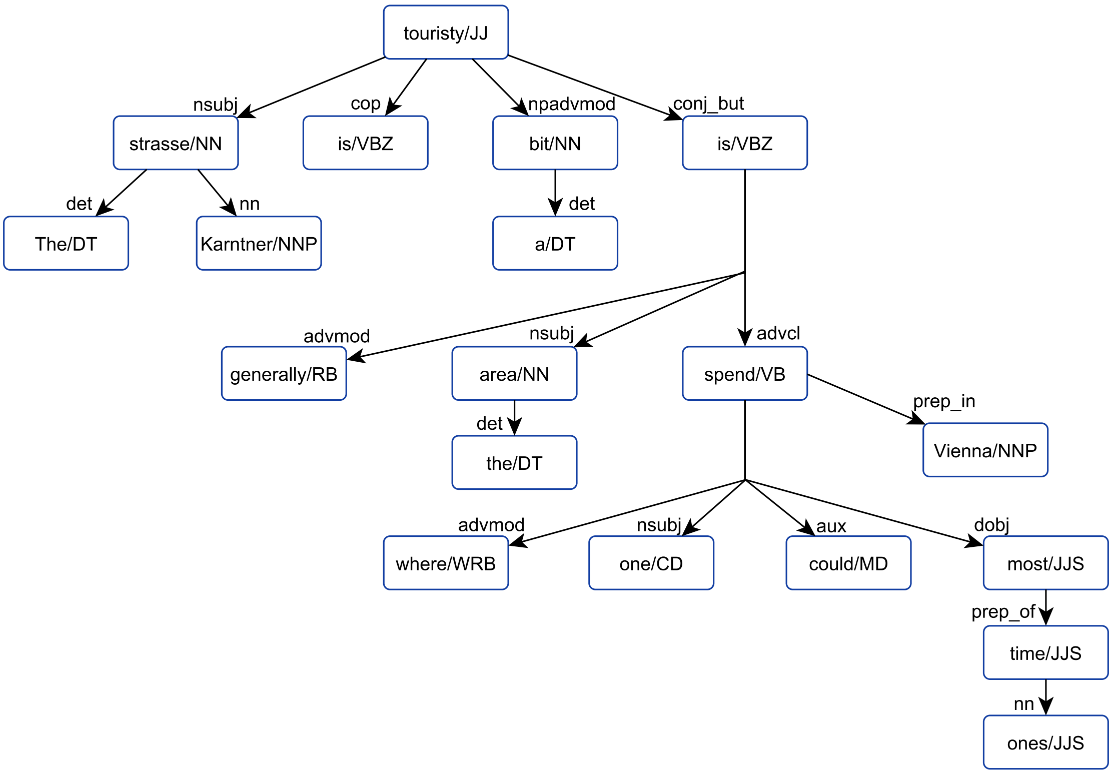

Figure 4.2: The dependency tree generated by the Stanford's part-of-speech tagger and dependency parser [Manning et al., 2014] for the sentence *"The Karntner Strasse is a bit touristy, but generally the area is where one could spend most of one's time in Vienna"*.

The sentence splitter tokenizes each User Generated Content description into sentences (step 1 in Figure 4.1) that are passed over to the Part-of-Speech tagger and the dependency parser (step 2 in Figure 4.1). The tagger classifies every word in a sentence according to its syntactical class, e.g. noun (NN), verb (VB), adjective (JJ) (see Table 4.1 for a more complete list of syntactical classes and tags). The parser generates a so-called dependency tree whose nodes denote the syntactical class of each word in a sentence, with edges representing the hierarchical structure of grammatical relations between the words. For example, given the sentence *"The Karntner Strasse*[1] *is a bit touristy, but generally the area is where one could spend most of one's time in Vienna."*, the Part-of-Speech tagger and the dependency parser produce the tree shown in Figure 4.2. Note that each term is also lemmatized, i.e. it is transformed into its base form.

Given a dependency tree, it is easy to extract from it the set $\mathcal{O}$ of common and proper nouns, tagged NN and NNP, respectively. This set then possibly contains a reference to geographic features in the textual description that we are interested in locating on the map. Since the name of a geographic feature may be a compound noun (e.g. Kärntner Straße, St. Stephen's Cathedral), we need to further process the set of nouns before trying to match them with geonames available in the geographic database.

Algorithm 3 Finding candidate compound geonames.

Input
$\mathcal{O}$ = *the set of nouns in UGC descriptions*,
x = *maximum number of words making up a compound geoname*
Output
$\mathcal{A} = \{\mathcal{A}_n : n \in \mathcal{O}\}$ = *the set of candidate geonames for each noun* n *in* $\mathcal{O}$

1: **procedure** CompoundGeonames
2: $\quad \mathcal{A} \leftarrow \emptyset$
3: $\quad$ **for all** $n \in \mathcal{O}$ **do**
4: $\qquad \mathcal{A}_n \leftarrow \emptyset$
5: $\qquad \mathcal{D} \leftarrow \{n\} \cup RetrieveDependencies(n, x)$
6: $\qquad$ **for all** $d \in 2^{\mathcal{D}}$ **do**
7: $\qquad\quad \mathcal{A}_n \leftarrow \mathcal{A}_n \cup PermutationsOf(d)$
8: $\qquad$ **end for**
9: $\qquad \mathcal{A} \leftarrow \mathcal{A} \cup \{\mathcal{A}_n\}$
10: $\quad$ **end for**
11: **end procedure**

We propose the procedure reported in Algorithm 3 that, given a noun $n \in \mathcal{O}$, produces a set $\mathcal{A}_n$ consisting of simple and compound nouns that we refer to as *candidate geonames* for n. For each noun $n \in \mathcal{O}$ we access the dependency tree to retrieve other nouns that, together with n, might make up a compound noun. This is done through the function *RetrieveDependencies(n, x)* (line 5) which, starting from the node corresponding to n, traverses the tree upwards (towards the root) and downwards (towards the leaves), and retrieves up to $x \in \mathbb{N}$ other nouns in both directions. These nouns, together with n, are stored in the set $\mathcal{D}$. The final set $\mathcal{A}_n$ of candidate compound nouns associated to n consists of all possible permutations of the elements of the powerset of $\mathcal{D}$ (lines 6-8). The set $\mathcal{A}$ is the collection of all candidate geonames associated to all nouns $n \in \mathcal{O}$, i.e. $\mathcal{A} = \{\mathcal{A}_n : n \in \mathcal{O}\}$ (line 9).

[1] The correct spelling in German language is Kärntner Straße. This comment has been retrieved from the Web and is purposely reproduced here with its original incorrect spelling.

Example 4.2.1: Natural Language Preprocessing

In our example, from the dependency tree in Figure 4.2 we derive:

$$\mathcal{O} = \{\textit{Strasse, Karntner, bit, area, time, Vienna}\}$$

And for the noun $n = \textit{Karntner}$ we have:

$$\mathcal{D}_{\textit{Karntner}} = \{\textit{Karntner, Strasse}\}$$
$$\mathcal{A}_{\textit{Karntner}} = \{\emptyset, \textit{Karntner, Strasse, Karntner Strasse, Strasse Karntner}\}$$

Note that in this case the number x of dependencies to be retrieved does not influence the sets of candidate compound names, as far as $x > 0$.

4.2.2 Geographic Matching

This stage does not rely on any external tool. The objective is trying to match every candidate geoname in the set $\mathcal{A}_n$ obtained in the previous stage against a unique feature in the geographic database according to name comparison (step 3 in Figure 4.1). The result is a set $\mathcal{G}$ that contains at most one geographic feature from the database for each element in $\mathcal{A}_n$: We choose the one (if it exists) whose name best matches against the candidate geonames in $\mathcal{A}_n$. This implies that we also discard nouns referring to categorical features (e.g. street, square), as our final goal is to pinpoint an initial area on the map that the textual descriptions refer to.

We propose the procedure reported in Algorithm 4 that works as follows. For each noun $n \in \mathcal{O}$ we retrieve (line 4) from the geographic database a set $\mathcal{P}$ of features whose names pattern-match (i.e. via regular expression) against n. In defining the regular expression (i.e. a search pattern), particular attention must be given to encode case-insensitivity and special characters (e.g. vowel mutations) to deal with spelling issues occurring when people write place names in a non-native language (e.g. the German word Straße vs. Strasse). Of all the retrieved features $\mathcal{P}$ we are only interested in selecting (lines 5-11) one whose name best matches against an element of the set $\mathcal{A}_n$ of candidate geonames associated to the noun n. Specifically, for each candidate geoname $c \in \mathcal{A}_n$ and for each feature $p \in \mathcal{P}$ we compute the Levenshtein distance[2] between c and the name of p (line 7).

[2] The Levenshtein distance is a string metric that measures similarity by the minimal number of required editing steps to transform one string into another string.

Algorithm 4 Geographic matching.

Input

$\mathcal{O}$ = *the set of nouns in UGC descriptions*,

$\mathcal{A} = \{\mathcal{A}_n : n \in \mathcal{O}\}$ = *the set of candidate geonames for each noun* n *in* $\mathcal{O}$,

ϵ = *threshold*

Output

$\mathcal{G}$ = *the set of matched geonames*

1: **procedure** GEOMATCHING
2: $\quad \mathcal{G} \leftarrow \emptyset$
3: $\quad$ **for all** $n \in \mathcal{O}$ **do**
4: $\qquad \mathcal{P} \leftarrow patternMatch(n)$
5: $\qquad (\underline{p}, \underline{d}) \leftarrow (nil, +\infty)$
6: $\qquad$ **for all** $(c, p) \in \mathcal{A}_n \times \mathcal{P}$ **do**
7: $\qquad\quad d \leftarrow Levenshtein(c, p.name)$
8: $\qquad\quad$ **if** $d \leq \epsilon \cdot WordsIn(c) \wedge d < \underline{d}$ **then**
9: $\qquad\qquad (\underline{p}, \underline{d}) \leftarrow (p, d)$
10: $\qquad\quad$ **end if**
11: $\qquad$ **end for**
12: $\qquad$ **if** $\underline{p} \neq nil \wedge \text{IsUnique}(\underline{p})$ **then**
13: $\qquad\quad \mathcal{G} \leftarrow \mathcal{G} \cup \{\underline{p}\}$
14: $\qquad$ **end if**
15: $\quad$ **end for**
16: **end procedure**

Example 4.2.2: Geographic Matching

Let us resume the example sentence introduced in Section 4.2.1 and whose dependency tree is shown in Figure 4.2. Assume that for the noun $n =$ *Karntner* and for the case-insensitive regular expression "k(a|ae|ä)rntner" (i.e. a search pattern that matches against *karntner*, *kaerntner*, and *kärntner*) the function *patternMatch* (line 4) returns only one feature named 'Kärntner Straße'. The following table then shows the resulting Levenshtein distance for each candidate geoname in $\mathcal{A}_n$ and the threshold (assuming $\epsilon = 3$ in our experiments) multiplied by the number of words in each noun:

$c \in \mathcal{A}_n$	Levenshtein Distance	$\epsilon \cdot WordsIn(c)$
$\emptyset$	15	0
'Karntner'	8	3
'Strasse'	11	3
'Karntner Strasse'	**3**	**6**
'Strasse Karntner'	12	6

It is easy to see that there is only one entry in this table whose distance is admissible and is minimum: the entry 'Karntner Strasse'.

To find possible matches (line 9) we enforce (line 8) that the distance not be larger than a given threshold ϵ. Since a candidate geoname might be a compound name, we multiply ϵ by the number of words making up the candidate geoname c. The best match is then the one with the smallest Levenshtein distance. At the end of the loop the variable $\underline{p}$ is either empty or it contains a geographic feature. In the first case no match has been found. Otherwise we must make sure that the feature is unique in the geographic database (line 12). This might not be the case for features like e.g. shops or restaurants that have several branches in the same city.

4.2.3 Machine Learning

This stage relies on a Machine Learning model called Multinomial Naïve Bayes: a probabilistic approach mainly used for text classification that learns from a given set of pre-classified samples (called training vectors) how to classify other, unclassified feature vectors according to their similarity with the given training vectors.

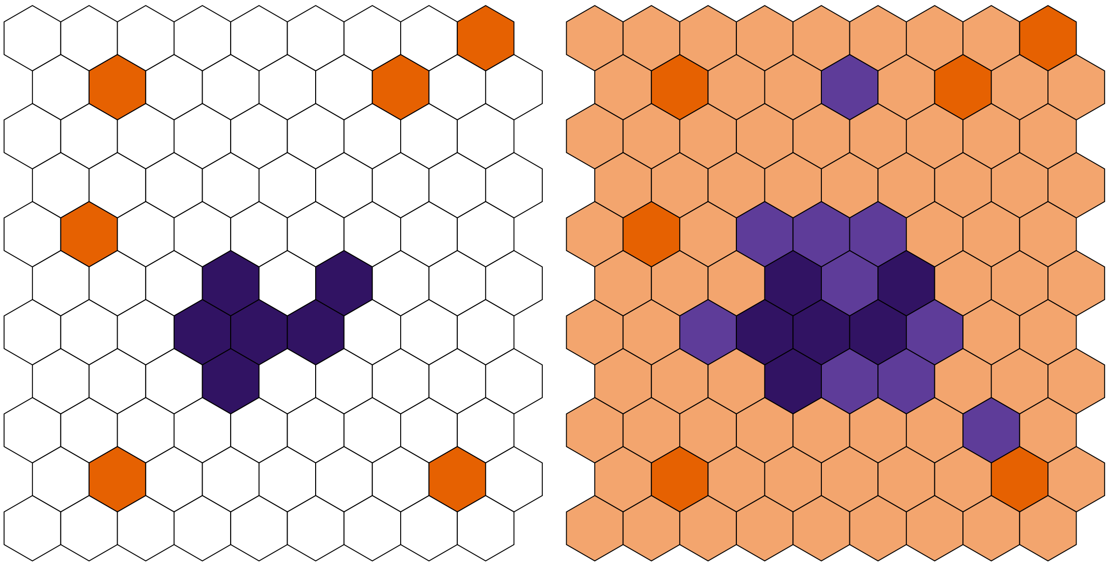

(a) Initial Configuration (b) Classified Area

Figure 4.3: A schematic representation of the classification process. Given training vectors for the cognitive region of interest (purple cells in Figure (a)) and for the counter-examples (orange cells in Figure (a)), the Machine Learning classifier associates each other cell to one of the two classes (light purple and light orange in Figure (b)).

We adapt Multinomial Naïve Bayes to classify geographic areas as either being part of the cognitive region of interest (class 1) or not being part of it (class 2). The training vectors are obtained by tessellating the map with a regular grid (step 4 in Figure 4.1) and retrieving the cells $\mathcal{I}_1$ containing at least one of the geographic features $\mathcal{G}$ derived in the previous stage. Such cells are the training vectors for the first class. The training vectors $\mathcal{I}_2$ for the second class consist of the same number of randomly selected cells that do not contain any of the geographic features in $\mathcal{G}$.

We adopt a bag-of-words model which is typically used for text classification to obtain a simplified semantical representation of the training cells by extracting certain categorical attributes from all the geographic features contained in each such cell (step 5 in Figure 4.1). Recall from Chapter 3 that $\mathcal{T} := \{t_i \colon i = 1, \ldots, n\}$ is a vocabulary containing all categorical attributes of interest from the whole map. Then, each cell is represented by a vector $(x_1, \ldots, x_n)$, where x_i is the frequency of the categorical attribute t_i in this cell. Since our focus is on cognitive regions conceptualized as homogeneous areas in terms of the activities they allow, we only select categorical attributes of geographic features that offer a service (e.g. bars, shops, restaurants, banks).

Given the two training sets $\mathcal{I}_1$ and $\mathcal{I}_2$ as described above, the Machine Learning procedure is capable of classifying all the remaining cells (step 6 in Figure 4.1), as graphically exemplified in Figure 4.3.

For convenience, the following two definitions are introduced:

Definition 4.2.1: A Priori Probability

The **a priori probability** $\hat{P}(k)$ for class k is the number of cells that are categorized as type k divided by the total number of cells:

$$\hat{P}(k) = \frac{\textit{number of cells of type } k}{\textit{total number of cells}} \tag{4.1}$$

Definition 4.2.2: Conditional Probability

The **conditional probability** $\hat{P}(t|k)$ of tag t in class k is the number of tags t occurring in class k plus a smoothing factor of 1 divided by the number of all tags in class k plus the size of the vocabulary $|V|$:

$$\hat{P}(t|k) = \frac{count(t,k) + 1}{count(k) + |V|} \tag{4.2}$$

Example 4.2.3: Multinomial Naïve Bayes

Given a simplified example for demonstration purposes as shown in the table below, the steps to classify a single test cell are illustrated.

	Cell	Tags	Class
Training	c_1	restaurant restaurant shop	$\mathcal{I}_1$
Training	c_2	restaurant restaurant monument	$\mathcal{I}_1$
Training	c_3	restaurant information	$\mathcal{I}_2$
Training	c_4	restaurant fountain bar	$\mathcal{I}_2$
Test	c_5	restaurant restaurant restaurant fountain bar	?

Then the **a priori probabilities** for the classes $\mathcal{I}_1$ and $\mathcal{I}_2$ are given as follows:

$$P(\mathcal{I}_1) = \frac{2}{4}, \quad P(\mathcal{I}_2) = \frac{2}{4}$$

The **conditional probabilities** for the relevant tag occurrences 'restaurant', 'fountain', and 'bar' for each class $\mathcal{I}_1$ and $\mathcal{I}_2$ are given as follows:

$$\begin{aligned} P(restaurant|\mathcal{I}_1) &= \frac{5}{12} \\ P(fountain|\mathcal{I}_1) &= \frac{1}{12} \\ P(bar|\mathcal{I}_1) &= \frac{1}{12} \\ P(restaurant|\mathcal{I}_2) &= \frac{3}{11} \\ P(fountain|\mathcal{I}_2) &= \frac{2}{11} \\ P(bar|\mathcal{I}_2) &= \frac{2}{11} \end{aligned}$$

Hence, for **choosing a class**, the following conditional probabilities are calculated:

$$\begin{aligned} P(\mathcal{I}_1|c_5) &\propto \frac{1}{2} * \left(\frac{5}{12}\right)^3 * \frac{1}{12} * \frac{1}{12} \\ P(\mathcal{I}_2|c_5) &\propto \frac{1}{2} * \left(\frac{3}{11}\right)^3 * \frac{2}{11} * \frac{2}{11} \end{aligned}$$

Since the probability $P(\mathcal{I}_2|c_5)$ is slightly higher than $P(\mathcal{I}_1|c_5)$, class $\mathcal{I}_2$ is chosen as class for the test example.

Definition 4.2.3: Classification

For **classification**, assuming that the cell c is identified with the vector $(x_1, \ldots, x_n)$ (e.g. in the example with (restaurant, restaurant, restaurant, fountain, bar)), the maximum a posteriori class k_{MAP} is defined as follows:

$$k_{MAP} = \max_{k \in K} P(c|k)P(k) = \max_{k \in K} P(x_1, \ldots, x_n|k)P(k) \tag{4.3}$$

Now, a simplifying independence assumption is made: the feature probabilities $P(x_i|k)$ are independent given the class k.

With this simplifying assumption, the problem can be efficiently solved in practice. Hence, the Multinomial Naïve Bayes classifier is defined as:

$$k_{NB} = \max_{k \in K} P(k) \prod_{i=1}^{n} P(x_i|k) \tag{4.4}$$

Besides the Naïve Bayes Classifier, many different approaches for *text* classification exist: decision trees, rule induction, neural networks, nearest neighbors, and support vector machines. However, due to its simplicity and effectiveness for small training samples, only the Naïve Bayes approach is utilized in this thesis.

4.3 Application Areas

In this section, two possible application areas of this approach are discussed.

4.3.1 Extraction of Cognitive Regions

In the scope of this thesis, we suppose that cognitive regions have homogeneous and uniform characteristics, e.g. a park consists of trees, green spaces, and children's playgrounds. Cognitive regions or settings are of substantial interest to the next generation of Geographic Information Systems and recommendation systems, considering that cognitive regions are conceptualizations of people's mental models, experiences, and individual knowledge. They also have the characteristic that they are generally understood by different people even when their boundaries are not clearly defined, which has been shown in an empirical study for inferring the location of downtown Santa Barbara [Montello et al., 2003].

Place is an ambiguous term and how people use place names in their everyday normal speech, describe places in natural language, and conceptualize places according to their individual experiences and given descriptions, are far beyond the scope of simple query to object matching. Therefore, the spatial footprint of cognitive regions can be expressed

Table 4.2: A Selection of the most common Tags within the Historic Center of Vienna

osm-places		osm-places	
tag	number	tag	number
natural=tree	1548	building=yes	1094
highway=footway	221	amenity=restaurant	191
landuse=grass	143	amenity=cafe	74
historic=memorial	60	shop=clothes	59
amenity=bicycle_parking	59	highway=bus_stop	40
amenity=bench	36	railway=tram	34
amenity=bank	34	tourism=hotel	34
service=parking_aisle	33	amenity=bar	33
religion=christian	31	amenity=telephone	29
amenity=fast_food	29	amenity=place_of_worship	29
tourism=museum	28	amenity=embassy	26
shop=books	24	amenity=fountain	23
shop=bakery	22	shop=shoes	21
amenity=post_box	21	building=church	20
amenity=pub	20	amenity=atm	18

in a bag-of-words model and Machine Learning approaches can be exploited. Table 4.2 shows the most common tags within the historic center of Vienna. Given that unique characteristics exist, these places can easily be identified by the proposed approach.

A different approach would be to identify *"activity clusters"* by searching for categorical tags and form clusters with the region growing algorithm, which was proposed in Chapter 3.

In this respect, the extraction of cognitive regions has several areas of applications. For example, one could be interested in the geometric extent of a given concept such as "downtown" that is not mapped as an explicit feature in Geographic Information Systems, and subsequently integrate it in spatial search engines. Another example would be the next generation of location-aware routing services, which could profit from recognizing a driver's unfamiliar regions and automatically switching to more detailed routing instructions. In conclusion, this concept has great potiential for several different areas of application and research.

4.3.2 Relocation Problem

Learning spatial knowledge is one of the main research areas of geographic information science. According to the Anchorpoint Theory [Golledge and Stimson, 1997], people learn spatial knowledge in a hierarchical structure where anchorpoints are initial starting points of knowledge generation. Such anchorpoints can be home, shopping, and work places, and according to these anchorpoints, a skeletal structure is created that is ordered hierarchically by primary, secondary, tertiary, and lower-order paths and places.

@Residential

↦ local supply supermarkets, small restaurants, kindergarden

@Work

↦ local supply supermarkets, restaurants

@Shopping

↦ supermarkets, restaurants, shops

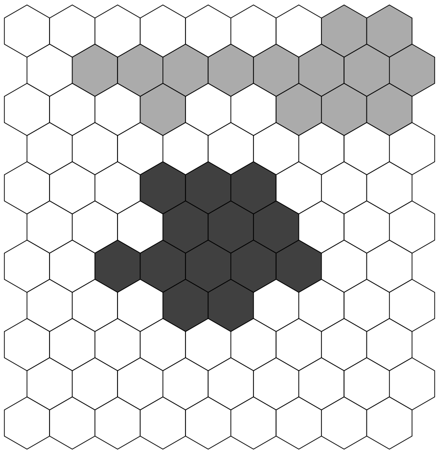

Figure 4.4: Concept Map – white cells designate residential, grey cells work, and black cells shopping settings of an individual.

A simple generalization of "activity clusters" can be made in respect to residential, work, and shopping areas. When people move from one city to another, the previous Machine Learning approach can be utilized to partition the environment in such conceptual classes. Figure 4.4 schematically illustrates that with given *feature vectors* of the original environment, and by utilizing Machine Learning, a new thematic map can be created.

4.4 Case Study

This section describes an implementation of the processing workflow, illustrating the feasibility of the described approach.

As data sources (see Figure 4.1) we selected two well-known User Generated Content and Volunteered Geographic Information projects: TripAdvisor[3] and OpenStreetMap[4]. By means of a customized crawler we retrieved English textual descriptions of the *historic center of Vienna* from a dedicated comment page on TripAdvisor. For the geographic database we used the OpenStreetMap extract of Vienna as provided by Mapzen Metro Extracts[5]. OpenStreetMap provides spatial data in the form of *points* (e.g. a park bench), *ways* (e.g. streets and buildings), and *relations* (e.g. spatial entities consisting of several parts). Semantic information such as name and categorical attributes are defined as 'tags', which are key-value pairs. For example, OpenStreetMap contains an entry for the "Hofburg Imperial Palace" that includes the name of the feature in several languages and is described by the following tags (among others): *(building, yes)*, *(historic, castle)*,

[3]http://www.tripadvisor.com/
[4]https://www.openstreetmap.org/
[5]https://mapzen.com/

Figure 4.5: Visualization of the OpenStreetMap dataset of Vienna as used in our experiments. The whole dataset consists of 290.586 nodes, 368.112 ways, and 810.145 relations, for a total of 1.468.843 features.

(castle_type, palace), *(tourism, attraction)*. The spatial dataset (see Figure 4.5) was stored in a dedicated database where the geometry of ways and relations was simplified by their centroid.

Finally, for the implementation of the Machine Learning stage (see Section 4.2.3) we resorted to a hexagonal grid with uniform cells with an edge-length of 0.0025°, and we used the MatLab implementation of the Multinomial Naïve Bayes[6] classifier.

4.4.1 Experimental Results of the Framework

We ran our workflow implementation on two experimental scenarios. Both scenarios use the same data sources with the following difference: in the first scenario (see Figure 4.6), all training vectors were kept in their integrity. In the second scenario (see Figure 4.7), we manually removed outlier cells from the training vector associated to the cognitive region, i.e. those cells that fall far away from the actual city center (compare the distribution of dark purple cells in Figures 4.6 and 4.7).

[6]http://de.mathworks.com/help/stats/naive-bayes-classification.html

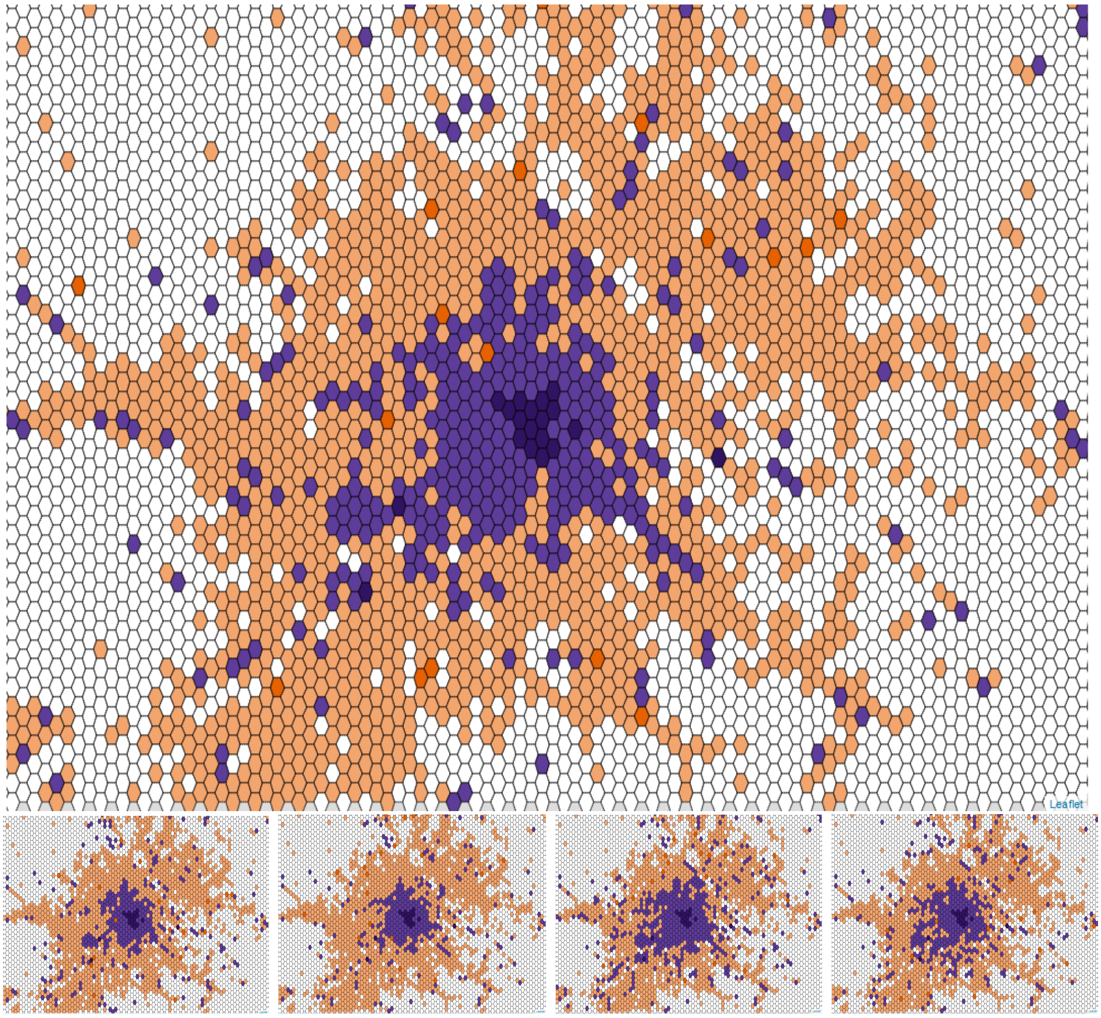

Figure 4.6: Visualization of classification results for the first scenario (several runs). Dark purple cells represent training vectors for the cognitive region *historic center of Vienna*; light purple cells are classified as *historic center of Vienna*; dark orange cells represent training vectors for the counter-example; light orange cells are classified as counter-example. White cells are unclassified.

For the image representations of the results we adopted the following color scheme: dark purple cells represent training vectors for the cognitive region *historic center of Vienna* as extracted from the textual descriptions; dark orange cells represent training vectors for the counter-example. Light purple and light orange cells show the areas classified as *historic center of Vienna*, and as the counter-example, respectively. White cells denote areas that have not been classified because of insufficient semantic information.

Since counter-examples are selected randomly from the grid, we decided to perform several runs for each scenario. Figure 4.6 shows the results obtained for five runs on the

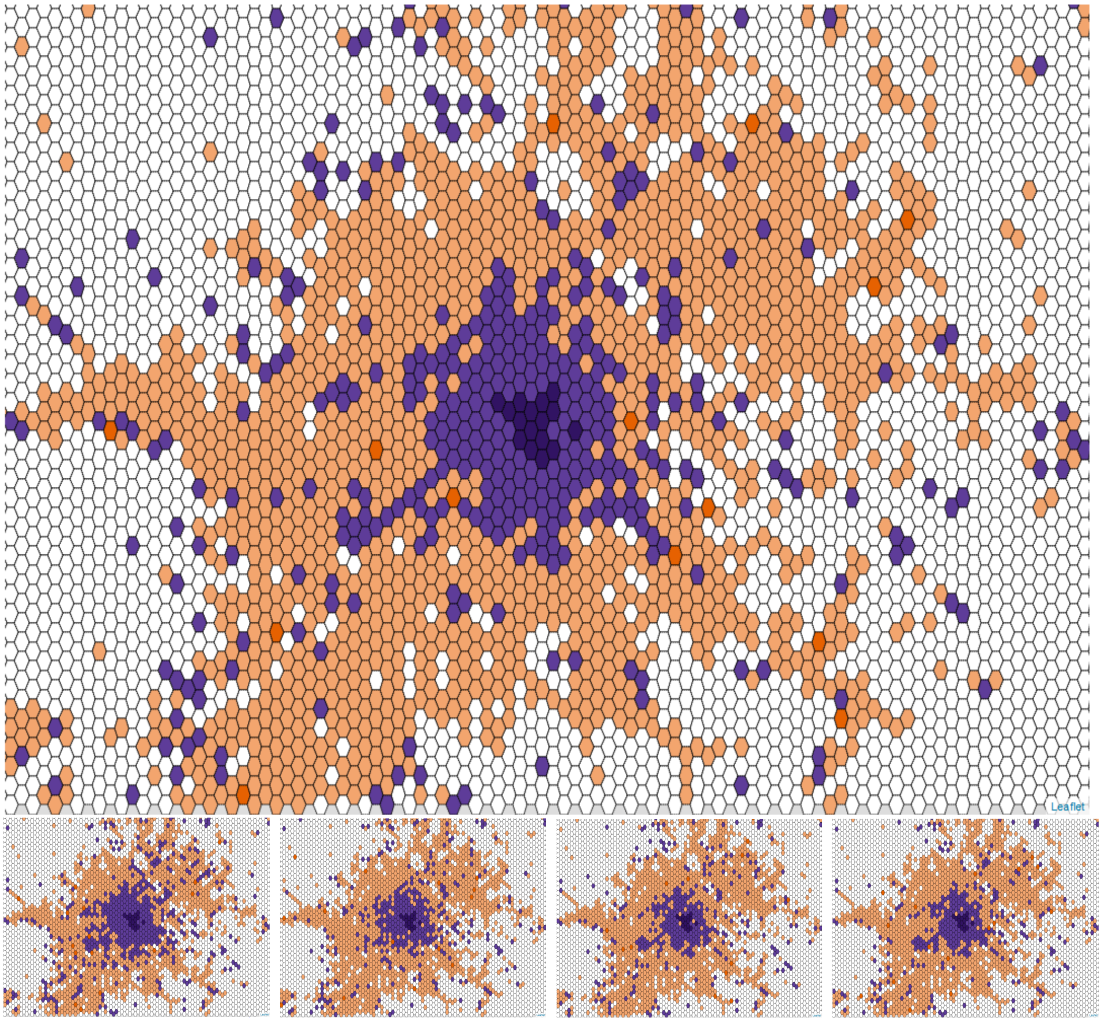

Figure 4.7: Visualization of classification results for the second scenario (several runs). Dark purple cells represent training vectors for the cognitive region *historic center of Vienna*; light purple cells are classified as *historic center of Vienna*; dark orange cells represent training vectors for the counter-example; light orange cells are classified as counter-example. White cells are unclassified.

first scenario. Figure 4.7 shows similar results for the second scenario, where outlier cells were removed from the training vector of the cognitive region. The results for the two scenarios mostly coincide, and the cells classified as similar to the cognitive region *historic center of Vienna* form a region approximately corresponding to the central district of the city and its immediate surroundings. Interestingly, the cells that were manually removed in the second scenario are associated to the class corresponding to the cognitive region anyway.

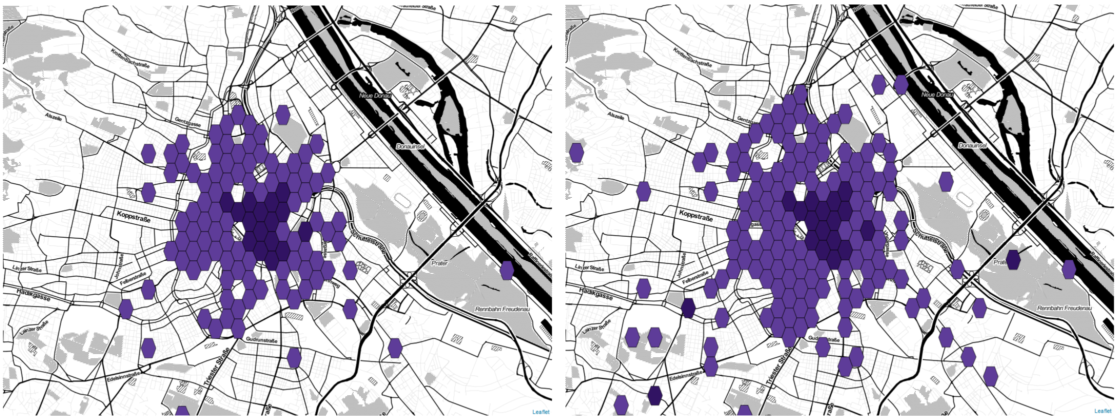

(a) Scenario 1 (b) Scenario 2

Figure 4.8: Visualization of robust results

To mitigate the effects of using randomly selected counter-examples, we performed ten runs for each scenario and intersected the results to obtain robust results: only cells classified as *historic center of Vienna* that occur in the result of each run form the robust results, as shown in Figure 4.8.

4.4.2 Qualitative Evaluation

We present here a preliminary qualitative evaluation of the outcomes by comparing the obtained robust results with a historical map of the city of Vienna that dates back to 1850. For that, we geographically overlaid the derived regions with the map, as shown in Figure 4.9 for the first scenario. It is easy to see how the shape and extent of the derived region correspond nicely to the city boundaries of 1850: the outer boundary of the main part of the classified area coincides with a physical separation which is now a major street of the city, while some the few outlier cells correspond to historical sites that are not reported in the historical map (e.g. the Schönbrunn Palace).

In summary, the approach, which relied solely on a knowledge base derived from Volunteered Geographic Information and crowd-sourced information sources, shows promising results.

4.4.3 The Relocation Problem

For the evaluation of the idea to classify cities according to their conceptual functions, four conceptual classes where defined:

- **Residential** (White)
- **Central Business Unit** (Red)

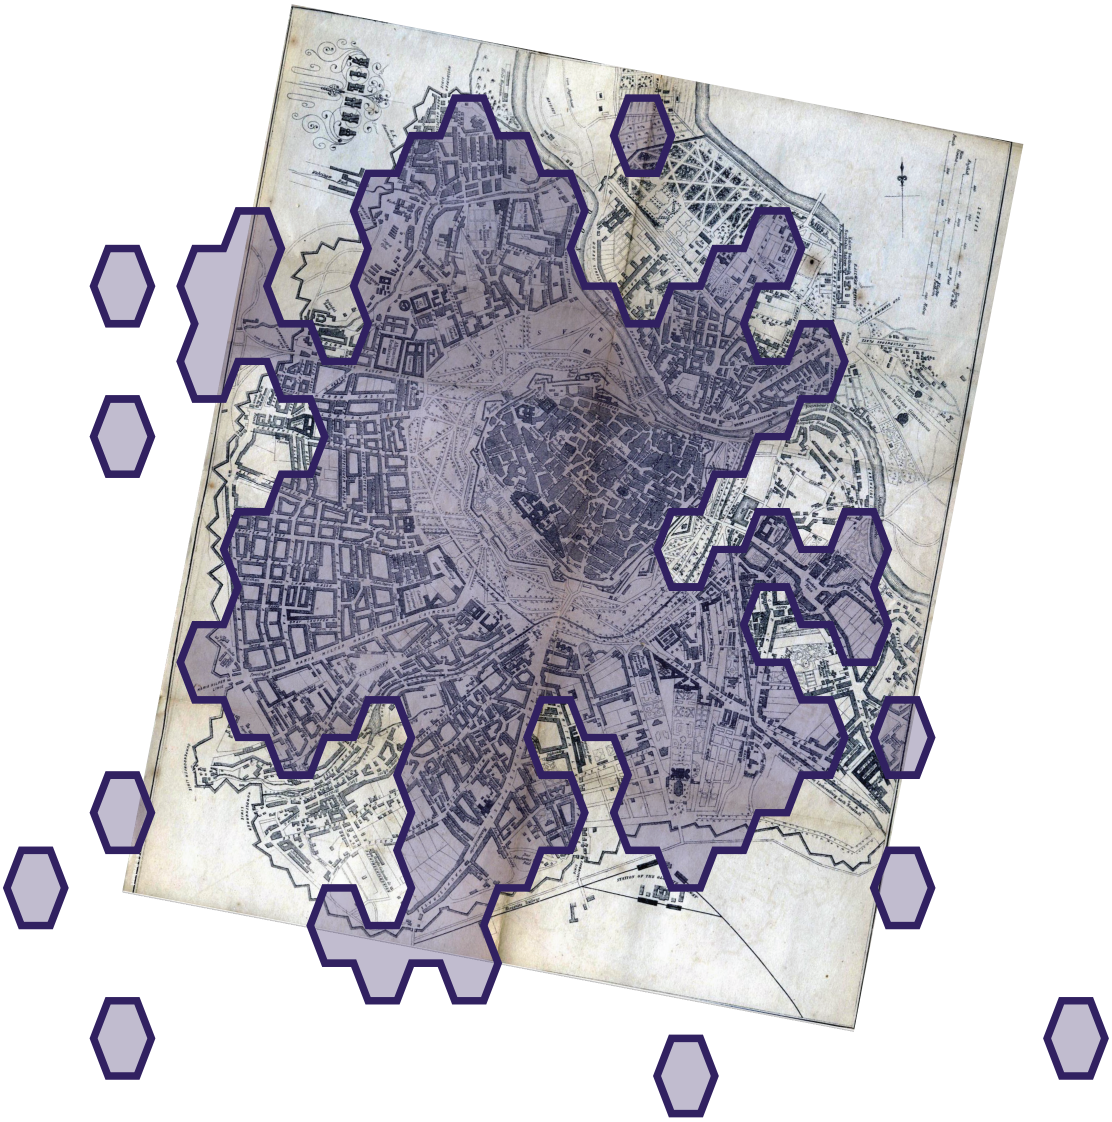

Figure 4.9: Approximate overlay of the robust result (Scenario 1) over a historic representation of Vienna (map retrieved from `www.valentina.net`).

- **Shopping** (Purple)
- **Amusement** (Blue)

Given one training sample for each of the conceptual classes – except Residential, which was provided with three samples – the remainder of Vienna is classified. Figure 4.10 shows the final classification. Grey indicates that the cells are not classified because insufficient information is available.

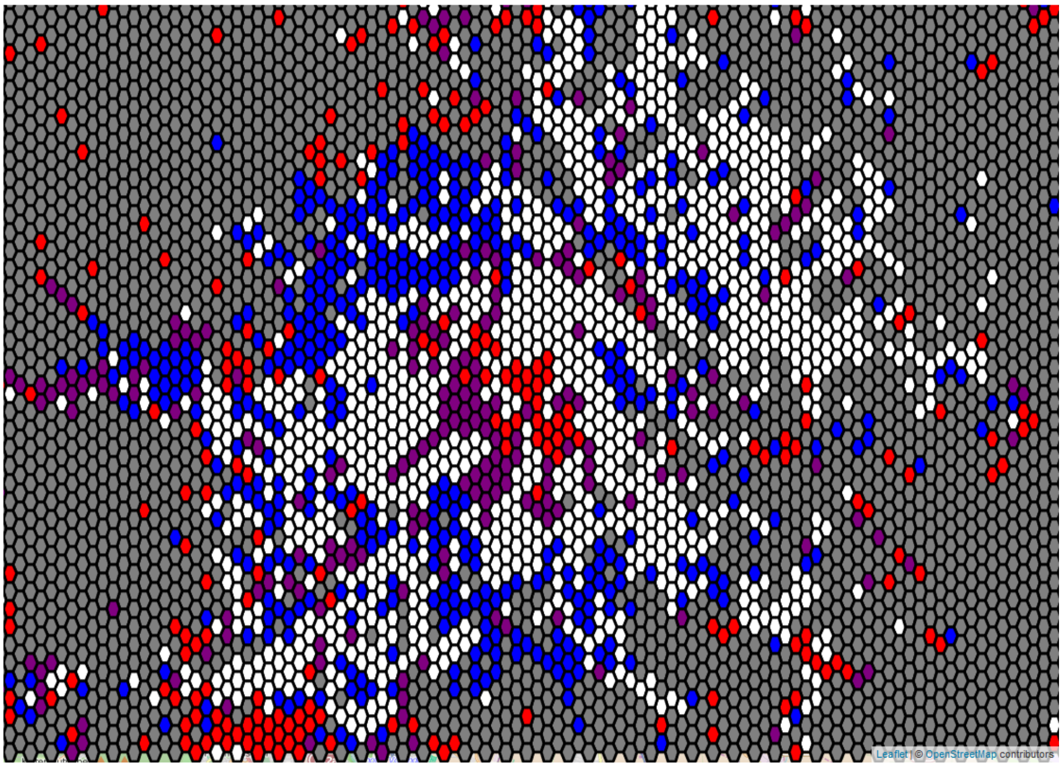

Figure 4.10: Visualization of classified areas (Vienna).

Again, all results were computed by using MatLab's Multinomial Naïve Bayes[7] classifier.

In the following, we briefly elaborate on the obtained results. As can be clearly seen in Figure 4.10, different functional areas give rise to different distinct regions obtained by the classifier. Notably we see that large extents of the "Donauinsel" and the north-western part of Vienna are correctly classified as "Amusement" areas. In contrast, the concept "Central Business Area" gives rise to a lot of small areas that are dedicated tourist attractions.

Consequently, we believe that the demonstrated result is strong enough to prove that tag distributions can indeed be used as basis to obtain a large-scale classification of an entire city.

4.5 Summary and Focus

This chapter shows a novel and automated approach to infer the geometric extent of "cognitive regions" by utilizing solely crowd-sourced geographic information as fundamental knowledge bases. Based on the Natural Language Preprocessing and combinatorial

[7] `http://de.mathworks.com/help/stats/naive-bayes-classification.html`

place matching procedure tailored to identify unique OpenStreetMap names, we have shown that the identification of seeding cells can be achieved by a simple Named Entity Recognition approach, yielding an improvement of the previously introduced approach (see Chapter 3).

Based on the previously identified seeding cells, an approach has been proposed which builds upon a "semantic representation" of OpenStreetMap service tags, allowing the automated clustering of cities into "cognitive regions" as a set of cells having clearly defined boundaries. We showed that the classification problem can be efficiently solved by utilizing the Multinomial Naïve Bayes model as classifier. Thereby, we have discussed a bi-classification approach operating on initial seeding cells identified by the combinatorial place matching procedure, and on counter-examples derived from a Monte Carlo approach.

Finally, two possible application areas have been discussed: first (i), the extraction of homogeneous and uniform cognitive regions with respect to urban search; and second (ii), the relocation problem by clustering urban areas in cognitive regions. Nevertheless, the proposed approach cannot answer more complex and nuanced questions about a cognitive region. For instance, which activities are *typically* referred to given places. The next chapter will therefore concern itself with the automated processing of User Generated Content to derive activities ascribed to geo-features.

CHAPTER 5

Extracting Semantics of Places from User Generated Content

The need for integrating and enabling human conceptualization, cognition, and common-sense reasoning in spatial search and Information Retrieval systems has been addressed in different specializations of geographic information science [Fogliaroni and Hobel, 2015, Zhang et al., 2008]. Modern Geographic Information Systems mostly support searches based on name and category matching, as well as on spatial relations among geographic features (e.g. the Opera House in Vienna, restaurants near Vienna). An important category of search that is not yet supported, however, concerns the retrieval of places based on the activities they afford. Agnew [1987] defined place as a combination of three elements: *location*, *locale* (the structure present in places), and *sense of place* (feelings and attachment to places). One major reason hindering place search is the lack of a cognitively plausible model for places that is capable of truly capturing the human understanding of such a fuzzy term. In the previous chapters it was argued that place affordances are essential for the development of recommendation systems based on place similarity and automated extraction of spatial footprints. This chapter describes modeling and extraction of place affordances by exploiting User Generated Content, which is mainly drawn from Hobel and Fogliaroni [2016].

5.1 Excursus

Drawing upon the taxonomy for geographic regions proposed by Montello [2003] and on affordance theory [Gibson, 1977, 1979], we regard a place as a cognitive region. More specifically, as a region of space conceptualized as a whole by people based on the activities it affords.

As a source for activity information we suggest exploiting User Generated Content like geo-logs, travel social media (e.g. TripAdvisor), and place review forums. These are naturally suitable sources for extracting the coveted information, as they convey human conceptualizations of places in the form of unstructured textual representations of cognitive regions.

In the spirit of the Semantic Web [Berners-Lee et al., 2001], this chapter introduces a vocabulary to model places in terms of the activities they afford, and describes a computational workflow to populate the model from User Generated Content. In addition, we also present an algorithmic approach that exploits Natural Language Processing tools to map unstructured text onto the proposed semantic model.

To enable a cognitive view in Information Retrieval, common-sense knowledge bases for Information Retrieval tasks were proposed, which are based on relationships among spatial objects [Zhang et al., 2008]. Based on the notion of Naïve Geography [Egenhofer and Mark, 1995], a novel framework was proposed in [Fogliaroni and Hobel, 2015], which enables qualitative spatial relation and configuration queries as a means to provide the casual user with more natural spatial search possibilities. Besides the spatial arrangement of objects, the question how people perceive places is tightly coupled with place affordances. The term affordance refers to Gibson's theory of visual perception and designates action potentials that are recognized by an agent in its environment [Gibson, 1977, 1979]. An object only affords an action if the agent's capabilities allow for the performing such an action.

Recently, much work has also focused on modeling, publishing, and consuming spatial data within the Semantic Web [Kuhn et al., 2014]. A striking example is the LinkedGeoData project [Stadler et al., 2012] that provides an encoding of OpenStreetMap data into the Resource Description Framework. The Linked Spatial Data trend is supported by different spatial (and temporal) extensions of basic Resource Description Framework and SPARQL. Some examples are the GeoSPARQL [Battle and Kolas, 2011] and the stSPARQL [Koubarakis and Kyzirakos, 2010] vocabularies and query extensions.

Natural Language Processing has also become a prominent research topic in geographic information science. The question how to model semantics of space for the retrieval from unstructured text is addressed, for example, by Bateman et al. [2010] where an ontology is presented for the processing of language concerned with space, actions in space, and spatial relationships. Khan et al. [2013] derived spatial triplets from unstructured text while Alazzawi et al. [2012] concentrated on pattern mining to derive language patterns for service identification.

Recent years have seen widespread research dealing with the extraction of structured information from unstructured text. This has been largely made possible by the availability of mature Natural Language Processing software like the Stanford CoreNLP toolkit [Manning et al., 2014]. This is a software suite that offers tools to parse and map unstructured text onto formal structures. One of the most interesting tools in this suite is the dependency parser, which generates a so–called dependency tree from a given sentence. The nodes of such a dependency tree denote the syntactical class of each word in a sentence (see Figure 5.1 for a partial list of such syntactical classes/tags). The

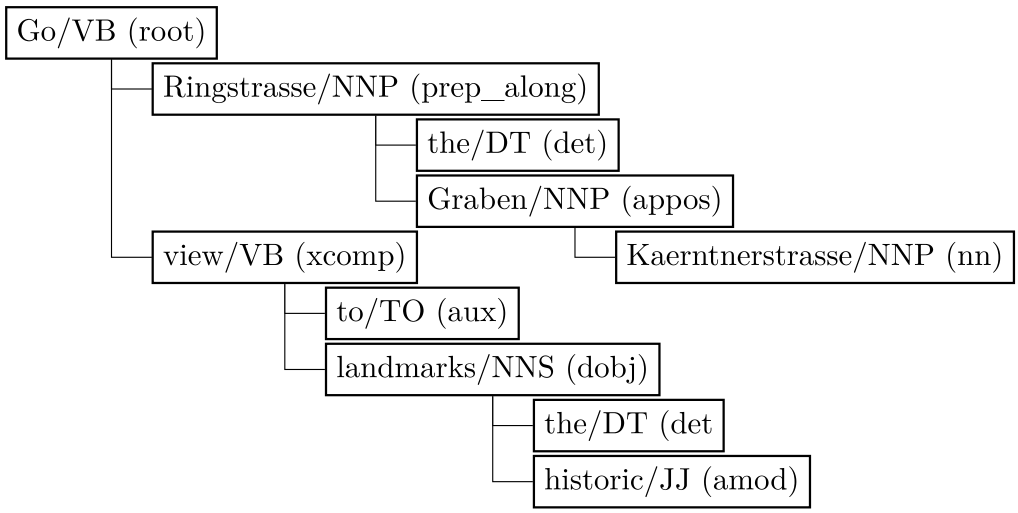

POST tag	Definition	POST tag	Definition
IN	preposition or subordinating conjunction	DT	determiner
JJ	adjective	WRB	adverb
NN	singular noun	WP	pronoun
NNS	mass noun	TO	to
NNP	proper noun, singular	VB	verb

Figure 5.1: Dependency tree obtained with the Stanford Dependency Parser [Manning et al., 2014] for the sentence "Go along the Ringstrasse, Kaerntnerstrasse, Graben to view the historic landmarks", and the partial list of syntactical classes and tags extracted from it [Santorini, 1990].

labeled edges represent the hierarchical structure of grammatical relations between the words. The root of a dependency tree always contains the verb of the independent clause of the sentence. The encoding of a sentence in a tree structure is done by extracting sequences of dependencies among words. Such dependencies are exactly the grammatical relations holding among different terms. If the parser is unable to narrow down the relation to a specific one, the edge is generally labeled rel. For example, the sentence “Go along the Ringstrasse, Kaerntnerstrasse, Graben to view the historic landmarks” is parsed onto the dependency tree shown in Figure 5.1, where syntactical classes and grammatical relations are both reported in the nodes.

5.2 Mapping Semantics onto Cognitive Regions

We propose the semantic model for places (specifically, cognitive regions) depicted in Figure 5.2. A place *pl* (root node in the diagram) recursively consists of other places or of an open number of geo-features that define the geographic footprint of *pl*. A geo-feature *gf* is a spatial entity (real or abstract) that possibly affords a number of different activities to be performed at or nearby the feature location. We regard an activity $act = (vb, ctx)$ as a pair consisting of a verb and a context. The verb expresses the type of activity (e.g. see, eat) while the context is any piece of ancillary information that narrows down or, more generally, modifies the semantics of the activity (e.g. see historical buildings, eat ice cream).

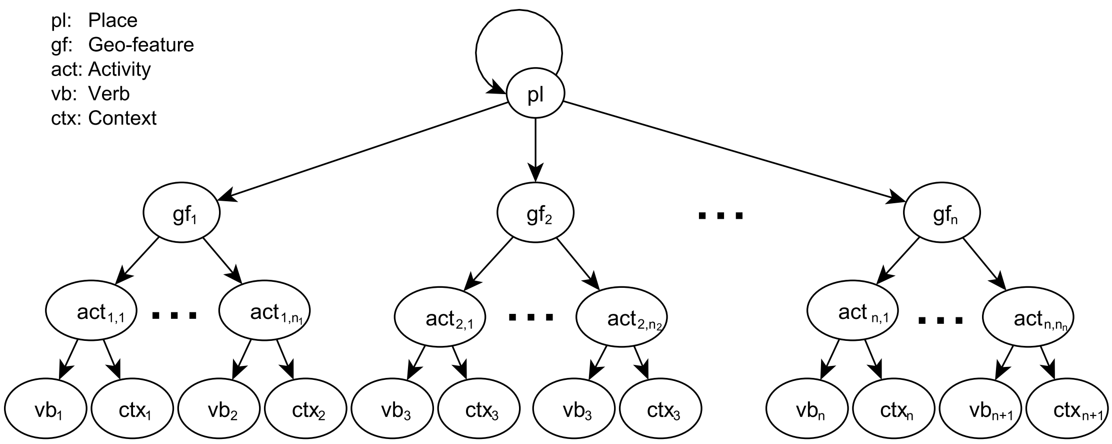

Figure 5.2: Abstract semantic model of a place. A place consists of one or more places or geo-features. Geo-features afford for activities. Activities consist of a verb in a given context.

Note that there is no restriction on the uniqueness of verbs, contexts, and verb-context pairs. In fact, a historical building (ctx_1), for example, can be seen (vb_1) but can also be photographed (vb_2). Moreover, historical buildings can be seen and photographed at different locations.

According to the proposed semantic model, an activity is also not unique – but geo-feature-activity pairs (gf_i, $act_{i,j}$) are. Consequently, a place is identified by its constituting geo-features and the activities they afford.

Embedding this abstract model within the Resource Description Framework [Consortium et al., 2014], we obtain a data model represented by the vocabulary in Figure 5.3. In this diagram, rounded nodes represent entities consisting of several attributes; rectangular nodes denote literals, whose types are defined in other vocabularies in the Semantic Web; edges report relations among nodes.

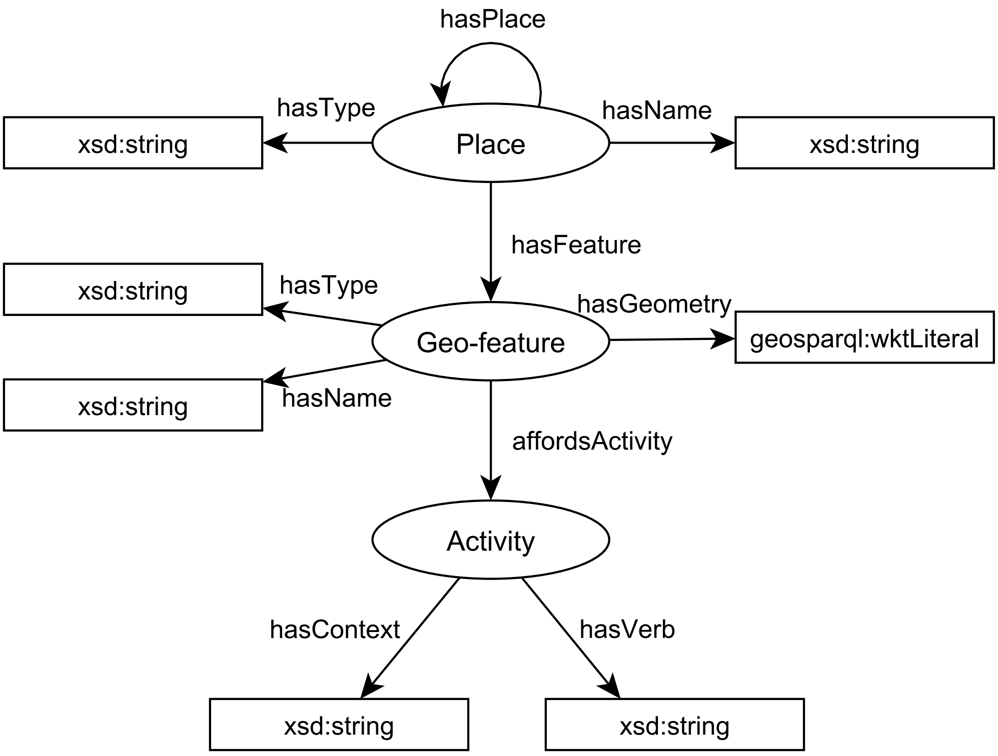

Figure 5.3: Vocabulary representing the data model for the abstract semantic model in Figure 5.2.

This less abstract model shows the composition of entities in more detail. Places and geo-features are typically referred to by one or more names, and may belong to a certain category or type (e.g. market, road). A geo-feature also has a geographic footprint that can be expressed, for example as well-known text (geosparql:wktLiteral). Activities afforded by a geo-feature are simply represented by a pair of strings, denoting the verb and the context of the activity.

5.3 A Workflow to Derive Semantic Place Representations from Unstructured Text

As of today, information about the central part of the model in Figure 5.3 (i.e. extent, type, and name of geographic features) can be easily retrieved from several open and private sources. What still remains an open question, however, is how to retrieve information about the rest of the model. Namely, we should find an answer to the following questions: "which geo-features constitute a place?", "which activities do these features afford?"

We present an approach that uses Natural Language Processing tools to extract this information from User Generated Content. More specifically, we present a workflow that uses the Stanford Dependency Parser [Manning et al., 2014] to process textual descriptions of places available on the Web (e.g. place reviews, tourist guides, travel logs, geo-blogs).

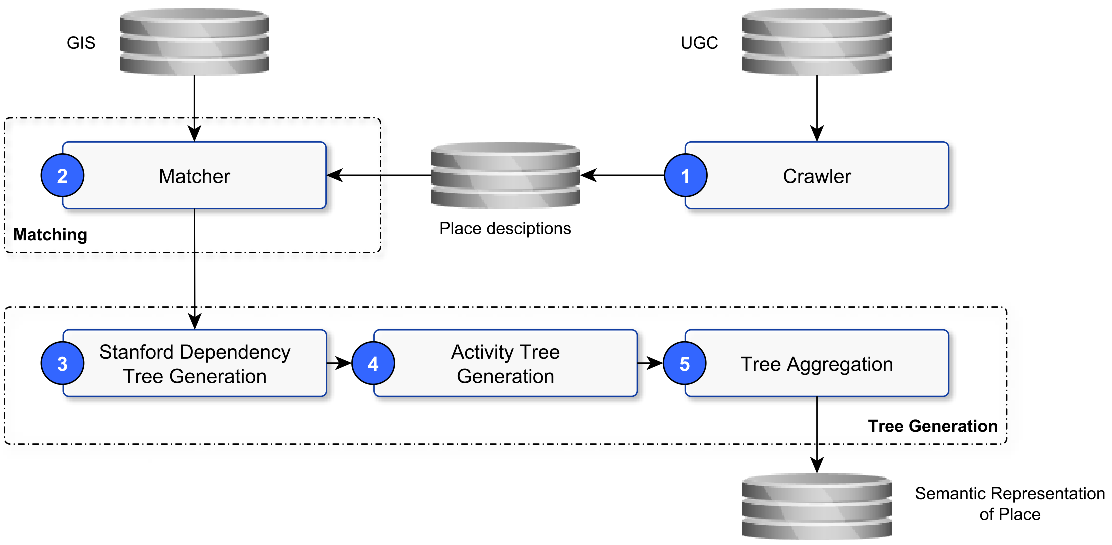

Figure 5.4: Processing workflow to derive semantic place representations from unstructured text.

The workflow (see Figure 5.4) consists of five main steps:

1. A web crawler automatically collects place descriptions in the form of unstructured natural language texts referring to a given cognitive region – e.g. the Historic Center of Vienna.

2. A spatial dataset of the area of interest is used as a knowledge base to match names and categories of geo-features against the textual descriptions. This step detects which geo-features constitute the place of interest.

3. Textual descriptions containing references to geo-features are processed with the Stanford Dependency Parser [Manning et al., 2014] to obtain a dependency tree – i.e.

syntactical classes of the terms occurring in the text and a structured representation of the grammatical relations among them.

4. The resulting dependency tree is parsed to detect verb-context pairs making up activities that can be performed at or in the proximity of a given location. This step ensures that the activity sub-trees are attached to a geo-feature node.

5. Activity sub-trees are aggregated by geo-feature to obtain the overall semantic representation of the place.

Step 1 (data crawling) can be performed separately and targets one of the two input data sources: the first is a set *TD* of textual descriptions of a given place *pl*. The second data source is the set *GF* of geo-features in the area of interest. Step 2 (geo-matching) is assumed to be realized by the function $s.refers(gf)$ appearing at line 3 of Algorithm 6 that detects whether a sentence s refers to a geo-feature *gf* (either by name or by category). An implementation of this function based on regular expressions is provided in Section 4.2.2.

In the remainder of this section we present an algorithmic realization of the core part of the workflow (steps 3, 4, and 5 above).

Algorithm 5 Given a place name *pl*, a set *TD* of textual descriptions referring to *pl*, and a geo dataset *GF*, the algorithm produces a semantic representation of the place *pl* according to the model in Figure 5.3.

Input

$pl =$ *a place/cognitive region*,
$GF = \{gf \mid gf \text{ is a geo-feature}\}$,
$TD = \{td \mid td \text{ is a textual description of } pl\}$

Output

$t_{pl} =$ *semantic tree description of pl*

function GeneratePlaceTree
  $t_{pl} \leftarrow initializePlaceTree(pl)$
  $S \leftarrow getSentences(TD)$
  for all $(gf, s) \in GF \times S$ **do**
    $t_{gf} \leftarrow$ GenerateGeoSubtree(gf, s)
    if $t_{gf} \neq \emptyset \wedge t_{gf}.hasActivities()$ **then**
      $t_{pl}.append(t_{gf})$
    end if
  end for
  return t_{pl}
end function

The main function is reported in Algorithm 5. Given a place *pl*, a set *TD* of textual descriptions of *pl*, and a spatial dataset *GF*, the function GeneratePlaceTree produces a semantic representation t_{pl} of pl according to the model described in the previous section. As an example, assume that pl is the "Historic Center of Vienna", GF is the OpenStreetMap dataset of the city of Vienna, and that TD comprises the text "I even enjoyed walking down the beautiful Kärntnerstrasse admiring many nice, original shops". The dependency tree of the latter is depicted in Figure 5.5, while the partial semantic representation derived from it is shown in Figure 5.6. As a first step (line 2) the algorithm creates the root node of the place tree (see Figure 5.6).

Subsequently (line 3), the textual descriptions are split into single sentences. This is done to avoid associating, later on, activities found in a sentence to the geo-feature(s) referred to in another sentence of the description. For each geo-feature-sentence pair (gf, s) the algorithm calls (line 5) the function GenerateGeoSubtree that is in charge of producing a so-called geo-sub-tree. A geo-sub-tree displays the part of the semantic representation starting at the geo-feature node and proceeding all the way down to the verbs and contexts that make up the activities (cf. Figure 5.3). If such a sub-tree is not empty and it includes at least one activity, it is appended to the place node t_{pl} (line 7).

Algorithm 6 Given a geo-feature gf and a sentence s produces a semantic representation t_{gf} of gf according to the model in Figure 5.3.

Input
  $gf =$ *a geo-feature*,
  $s =$ *a sentence*
Output
  $t_{gf} =$ *semantic subtree description of a geo-feature*

1: **function** GenerateGeoSubtree
2: $\quad t_{gf} \leftarrow \emptyset$
3: $\quad$ **if** $s.refers(gf)$ **then**
4: $\qquad \mathcal{D} \leftarrow getDependencyTrees(s)$
5: $\qquad$ **for all** $d \in \mathcal{D}$ **do**
6: $\qquad\quad T_{act} \leftarrow$ GenerateActivitySubtrees(d)
7: $\qquad\quad$ **if** $T_{act} \neq \emptyset$ **then**
8: $\qquad\qquad t_{gf} \leftarrow initializeGeoTree(gf)$
9: $\qquad\qquad t_{gf}.appendAll(T_{act})$
10: $\qquad\quad$ **end if**
11: $\qquad$ **end for**
12: $\quad$ **end if**
13: $\quad$ **return** t_{gf}
14: **end function**

GenerateGeoSubtree is described in Algorithm 6. An empty geo-sub-tree t_{gf} is initialized (line 2) which is actually built (line 8) only if the following two conditions hold true: (i) The input sentence s refers to the input geo-feature gf (line 3). (ii) It is possible to associate at least one activity to gf (line 7). If the first condition is satisfied, a set $\mathcal{D}$ of dependency trees is generated from the given sentence (line 6). The function $getDependencyTrees(s)$ utilizes the Stanford Dependency Parser [Manning et al., 2014] to generate a dependency tree for each independent clause of the sentence s (step 3 of the workflow in Figure 5.4).

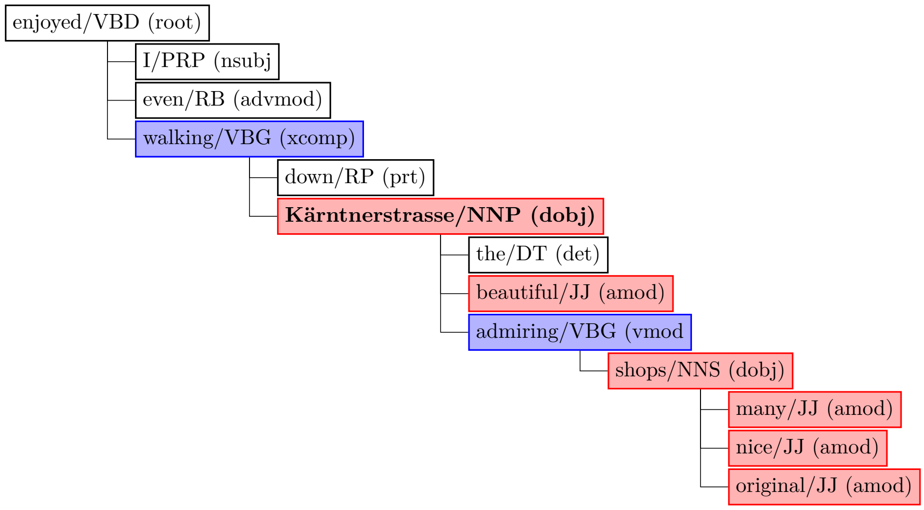

Figure 5.5: Dependency tree obtained through the Stanford Dependency Parser [Manning et al., 2014] for the sentence "I even enjoyed walking down the beautiful Kärntnerstrasse admiring many nice, original shops". The node labeled in bold text is the only geo-feature. Blue and red nodes indicate verbs and contexts of activities, respectively.

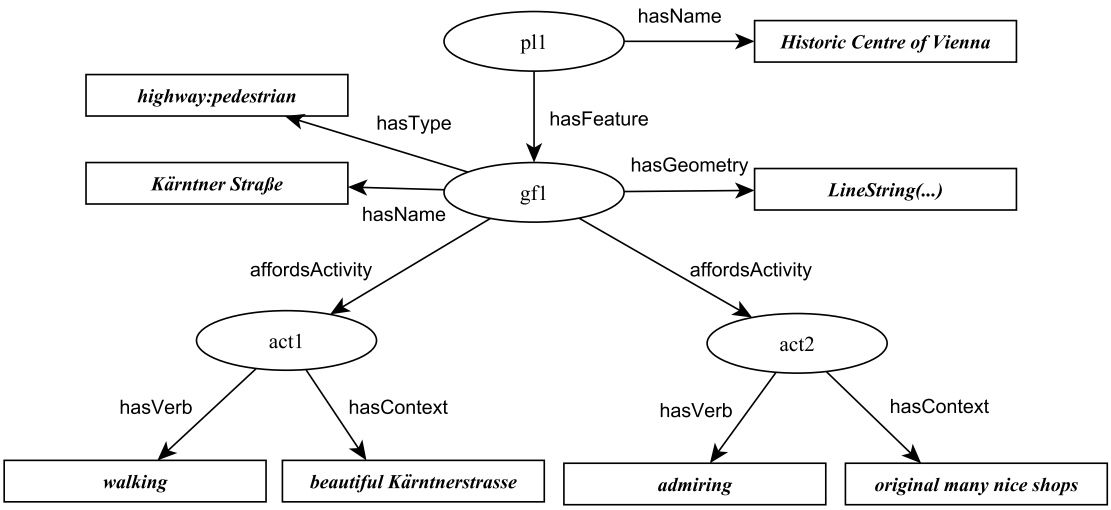

Figure 5.6: (Partial) semantic representation (according to the model in Figure 5.3) of the Historic Center of Vienna as derived by the description of Figure 5.5.

For the running example we only consider one sentence and one independent clause, so only the tree in Figure 5.5 is generated. The node labeled in bold text indicates the only geo-feature for which a name match was found in the spatial dataset *GF*. Each such dependency tree *d* is given as input to the function GenerateActivitySubtree that is in charge of mapping from *d* onto the part of the semantic representation rooted at the activity node (cf. Figure 5.3). This function returns a (possibly empty) set T_{act} of activity-sub-trees: one for each verb-context pair found in the dependency tree. If the returned set is not empty, the geo-sub-tree is finally initialized (line 8) and the activity-sub-trees are appended to the geo-feature node t_{gf}. Note that the given algorithmic realization assumes data to be represented as Resource Description Framework triples, and persisted in a triple store. This means that a geo-feature *gf* corresponds to a uniquely identified node in the triple store. The function $initializeGeoTree(gf)$ creates a new node only if *gf* is not already stored, otherwise it retrieves it from the triple store. This implements step 5 (aggregation) of the workflow in Figure 5.4.

Algorithm 7 Given a dependency tree *d*, the algorithm produces a set T_{act} of activity sub-trees to be attached to a geo-feature sub-tree.

Input
$d = a\ dependency\ tree$
Output
$T_{act} = \{t_{act} \mid t_{act}\ is\ the\ semantic\ description\ of\ an\ activity\ act\ \}$

1: **function** GenerateActivitySubtree
2: $T_{act} \leftarrow \emptyset$
3: $O \leftarrow d.getDObjs()$
4: **for all** $o \in O$ **do**
5: $t_{act} \leftarrow initializeActivityTree()$
6: $vb, ctx \leftarrow getVerb(o), getContext(o)$
7: $t_{act}.appendVerb(vb)$
8: $t_{act}.appendContext(ctx)$
9: $T_{act} \leftarrow T_{act} \cup t_{act}$
10: **end for**
11: **return** T_{act}
12: **end function**

GenerateActivitySubtree is described in Algorithm 7. The dependency tree under consideration possibly contains several verbs (e.g. enjoyed, walking, admiring), but we are only interested in those that also have a refining context. To achieve this we suggest using the grammatical relation *dobj* (direct object) that leads to the part of the sentence recognized as the (accusative) object of a verbal predicate [De Marneffe and Manning, 2008]. The dependency tree in Figure 5.5 contains two such objects: Kärntnerstrasse and shops. For each such direct object *o* an activity-sub-tree t_{act} is initialized (line 5). The functions $getVerb(o)$ and $getContext(o)$ (line 6) start from node *o* of the dependency tree and traverse it to retrieve the verb *vb* and the context *ctx*. By construction, the verb vb is always the nearest verbal (i.e. tagged VB) ancestor of o in the dependency tree. In our example, we get walking for the object Kärntnerstrasse and admiring for the object shops. We assume that the context consists of the object o itself plus all related adjectives (i.e. tagged JJ). By construction, these are located in the tree branch rooted at vb and going trough o. In the example, we have beautiful Kärntnerstrasse and many nice original shops. Verb and context are appended (lines 7-8) to the root t_{act} of the activity-sub-tree, which is finally added to the return set T_{act} (line 9).

5.4 Application Areas

In this section, two possible application areas of the outlined approach are discussed.

5.4.1 Place Search Queries expressed in Natural Language

While today's spatial search services mostly support searches based on name and category matching as well as on spatial relations among geographic features (e.g. the Opera House in Vienna, restaurants nearby Vienna), more intuitive questions about the environment are not possible. Arguments are frequently made that in order to sufficiently cover the information needs of a user's daily task the information sources have to be enriched. By utilizing the approach presented in this chapter, however, human impressions (e.g. which places are assumed to be *near* another place) and the activities that can be performed at places can be transformed to a machine-processable model.

One possible application for the presented model is to make natural language query answering possible. In fact, a spatial question posed in natural language can be interpreted into the same (tree) semantic structure onto which we encoded unstructured place descriptions. Answering the query then simply consists of matching two graphs. While the presented approach is part of ongoing research, we believe it provides a solid foundation for future improvements.

Two major simplifications were the result of neglecting (i) spatial relations, and (ii) negations. The former can strongly modify the spatial location where an activity can be performed (e.g. many inexpensive shops can be found outside the city center) – although this is supposedly rare in reviews that describe a place, as we assume that people typically describe what one can do at a location, rather than away from it. One approach in tackling this particular challenge is to utilize ontologies of spatial relations, as provided, for example, by qualitative spatial calculi [Bateman et al., 2010]. The case of negations is more concerned with Natural Language Processing, and can be addressed with techniques used in sentiment analysis.

Another aspect that is disregarded in the current model concerns the semantic similarity of verbs, contexts, and activities. For example, the activity of "seeing historical landmarks" is obviously similar (if not outright equivalent) to "admiring old monuments". To address this aspect, the semantic similarity of terms should be accounted for. One approach is to resort to a synonym structure as provided e.g. by WordNet [Miller, 1995].

5.4.2 Activity Recognition for Place Recommendation and Machine Learning

In Chapter 3, we introduced a model for partitioning cities into functional regions, which enables the recommendation and ordering of places based on semantic similarity using distance metrics. In that preliminary model one important assumption is made: that place affordances are modeled based on the services an area offers. By using the approach presented in this chapter, the activities can be derived from textual descriptions of places, and, consequently, place recommendations based on semantic similarity can be facilitated.

In Chapter 4, we outlined an approach that utilizes Natural Language Processing and Machine Learning to derive the geographic footprint of cognitive regions. One essential part of this model relies on manually selecting the services excluded for the vocabulary, for instance tags with the key “ref”. This manual selection can be replaced by using the approach introduced in this chapter to infer the services people associate to activities.

In conclusion, the approach presented in this chapter complements the methods presented in the previous chapters. In combination, the framework then allows us to address and answer place-related questions.

5.5 Case Study

For the following evaluation, we use real-world and publicly available datasets to construct the semantic representation. Before giving details of experimental results, we introduce how the datasets were retrieved and the quantitative and qualitative characteristics of the constructed semantic representation.

5.5.1 Data Sources and Collection

Two data sources are utilized for the proposed approach: (i) place descriptions of the Historic Center of Vienna originating from Tripadvisor[1]; and (ii) OpenStreetMap as the geographic knowledge base for Points of Interest.

Our first data source, TripAdvisor[2], was utilized as the source for place descriptions. TripAdvisor provides around 225 million crowd-sourced reviews of different attractions, hotels, and (most interestingly and left unspecified on the website) “places”. A customized crawler starts at the main page of a given place and dives into the subpages, collecting all of the place’s full-text reviews. A set of exemplary place descriptions referring to the spatial object “Kärntner Straße” is shown in Table 5.1.

Our second data source, OpenStreetMap, is a major Volunteered Geographic Information platform and serves as the knowledge base for the physical features, e.g. roads or buildings, or generally accepted boundaries (e.g. administrative boundaries). The data set was

[1] http://www.tripadvisor.com/
[2] TripAdvisor.at

Table 5.1: Examples for the retrieved Place Descriptions

We carried on an easy stroll along Karntner Strasse where good restaurants, bars and some stores are located.
The cafe is called Bistro 59, just at the junction of Karntner Strasse and Borsendorforstrasse.
The Karntnerstrasse is a bit touristy but generally the area is where one could spend most of ones time in Vienna.
It's great fun to wander down Karntnerstrasse and check out the shop windows. By historic center, I assume Tripadvisor means Karntnerstrasse, which is the main pedestrian drag in Vienna and its tourist center.
I even enjoyed walking down the Karntnerstrasse admiring the shops.

Table 5.3: OSM Dataset

Name of Dataset	Features	Points	Ways	Relations
Vienna Bratislava, Austria	1.468.843	290.586	368.112	810.145

retrieved from Mapzen Metro Extracts[3] and provides spatial data in the form of Points (e.g. a park bench), Ways (e.g. streets), and Relations (e.g. administrative boundaries). The spatial entries (see Table 5.3) were simplified by their centroid and stored in a dedicated database.

5.5.2 Quantitative Evaluation

For the dataset, we used the concept "Historic Center of Vienna[4]". We collected 1235 reviews in the English language. Subsequently, we ran our proposed approach on the dataset retrieved from TripAdvisor by utilizing the OpenStreetMap dataset to map the gained information to spatial features. In total, we obtained 33060 geo-features that are enriched with activity information. We only considered places with names longer than two characters in order to exclude, among others, categories of buses or trains. This elimination process then resulted in 2135 single places with a total of 3810 activities. A first overview based on a visualization of the results revealed that most of the extracted information is consistent with the aims of the work, and can thus be used as a knowledge

[3] https://mapzen.com/

[4] https://www.tripadvisor.at/Attraction_Review-g190454-d1534524-Reviews-Historic_Center_of_Vienna-Vienna.html

base for spatial requests. A few outliers such as *xpedit* or *fuel* could be identified which are not places and occurred based on the naive regex-matching.

5.5.3 Qualitative Evaluation

For a brief qualitative evaluation, the processed sentences and the obtained results for the geo-feature *Kärnter Straße* are shown and discussed.

Table 5.4 shows a selection of sentences and inferred activities that are referred to the geo-feature *Kärnter Straße*, which is a well-known shopping street of Vienna. The derived results support the character of the street as shopping and sightseeing opportunity.

Table 5.4: Examples for the activities mapped as semantic graph nodes for the place *Kärnter Straße*.

Sentence	**Verb**	**Contextual Information**
The Karntnerstrasse is a bit touristy but generally the area is where one could spend most of ones time in Vienna.	spend	most ones time
It's great fun to wander down Karntnerstrasse and check out the shop windows.	wander; check	Karntnerstrasse; shop windows
I would suggest Schwedenplatz as the starting point down to Stephansplatz and while your enjoying the street crowd and shops of Karntner Strasse, you'll reach the Opera house or Wiener Staatsoper.	suggest; enjoying; enjoying; reach; reach	Schwedenplatz; crowd; shops; Opera; Wiener Staatsoper
be sure to visit St. Stephen's Cathedral, walk along the delightful pedestrian Kärntnerstrasse, stop to enjoy a piece of decadent cherry strudel in a Viennese coffee house.	visit; walk; enjoy	St. Stephen Cathedral; pedestrian Kärntnerstraße; decadent piece
There is something for each weather and taste - enjoy a walk in the generous pedestrian areas (Kärntner Strasse, Graben) and do some window shopping (or more if you can afford!)	enjoy; do	walk; window shopping

The results obtained by the proposed workflow support the previously introduced approaches of city generalization and place recommendation as well as Machine Learning by extracting the activities as well as the geo-features to which the activities are associated. What remains for future work is generalizing activities and clustering them into meaningful classes, which would go beyond the scope of this thesis. However, the presented approach does show great potential in facilitating spatial search, since place is a socially constructed concept, and the current trend of research is trying to provide models for place representation and querying. We focus on a theory that regards place as emergent of possible activities. In this respect, the next generation of geographic information services can allow more intuitive queries such as “shopping in Vienna” or “sightseeing for historic landmarks”.

5.6 Summary

In this chapter, a framework for extracting activities from User Generated Content is discussed. A semantic model was introduced, modeling place as emergent of possible activities. The framework extends the previously proposed concepts by populating an ontology with people’s perception.

In the beginning of this chapter, a preliminary model of place is introduced that complements the previously introduced approaches by modeling the activities and geo-features to which activities are asserted.

For the first approach, city generalization and place recommendation based on semantic similarity, one important assumption is made: that place affordances are modeled based on the services an area offers. By using the approach presented in this chapter, the activities can be derived from textual descriptions of places, and, consequently, place recommendations based on semantic similarity can be facilitated.

The second approach, Machine Learning, depends on manually selecting the services an area offers. This manual selection can be replaced by using the approach introduced in this section to automatically infer the services people associate to activities.

Finally, an algorithmic solution is outlined that automatically extracts activities from User Generated Content by exploiting mature Natural Language Processing tools. It is argued that besides the applicability in the previous approaches, the approach can allow for place search where queries are expressed in natural language. By processing questions such as “where can I view historic landmarks in a place that offers shopping opportunities?” with the approach presented in this chapter, the problem is essentially reduced to simple graph matching. What remains for future work is generalizing activities and clustering them into meaningful classes, which would go beyond the scope of this thesis.

CHAPTER 6

Discussion

In this section, the previously posited research questions are discussed.

6.1 Extraction of Cognitive Regions

Montello et al. [2014] explained the characteristics of cognitive regions as follows:

> "Cognitive regions are regions in the mind, reflecting informal ways individuals and cultural groups organize their understanding of earth landscapes. Cognitive region boundaries are typically substantially vague and their membership functions are substantially variable – the transition from outside to inside the region is imprecise or vague, and different places within the region are not equally strong or clear as exemplars of the region."

In this thesis, it is argued that cognitive regions have relatively uniform and homogeneous characteristics, allowing the extraction of such regions by two different methods discussed in Chapter 3 and Chapter 4. Therefore, the question

> "How to model cities as regions of functional areas?",

can be answered by taking the categorical tags of aggregated places as the simplified 'semantic' representation (i.e. bag-of-words model) of an area. In Chapter 3 a segmentation procedure is proposed that is inspired by a technique called region growing used in image segmentation [Adams and Bischof, 1994]. The essential idea is to categorize areas according to the conceptualization of human beings. This idea is further extended in Chapter 4, where exemplary cognitive regions are automatically derived by exploiting User Generated Content, and by utilizing the bag-of-words model of categorical attributes the remaining parts of a city can be classified. Additionally, it is shown that the Machine Learning approach can be further used to classify parts of a city into functional regions based on an ontological description of the environment.

6.2 Semantic Similarity

By taking the categorical tags as simplified 'semantic' representation (i.e. bag-of-words model) of an area, semantic similarity can be addressed in a way analogous to traditional Information Retrieval (see Chapter 3). Places as aggregated concepts can, therefore, be exploited, and the question

> "How can semantic similarity be integrated for spatial search engines?",

can be answered by utilizing distance metrics defined on pairs of abstract representations in order to derive or rank the similarity of places. The feasibility of this approach was illustrated in an example comparing the conceptual place Oxford Street (London) with shopping areas in Vienna. It is highlighted that when combining this approach with human preference models, spatial search can be improved. The human preference model is closely related to *place affordances* and can be intuitively integrated into spatial search models.

6.3 Integration of Machine Learning

In Chapter 4, a probabilistic model for cognitive regions was discussed. Since in the scope of this thesis it is argued that cognitive regions have uniform and homogeneous characteristics, the question

> "How to incorporate statistics and Machine Learning in spatial Information Retrieval?",

can be answered by using a bi-classification approach taking seeding cells and Monte-Carlo-random counterexamples, and exploiting the bag-of-words model of categorical tags. Furthermore, the relocation problem of identifying similar functional regions in different cities shows promising results.

6.4 Natural Language Processing

In the spatial domain, User Generated Content has the potential to bridge the gap from conceptual models to successful expert systems by providing the necessary data for novel knowledge bases. To make sense out of today's huge amount of available information sources, mature Natural Language Processing methods can be utilized, and, therefore, the question

> "Can Natural Language Processing be used in Geographic Information Retrieval?"

is answered in the affirmative in Chapter 4 and Chapter 5. In conclusion, Natural Language Processing allows us to derive the geometric extent of cognitive regions by matching geographic content against User Generated Content as shown in Chapter 4, and to create semantically enriched search graphs as shown in Chapter 5.

6.5 Knowledge Bases

Spatial search is characterized by human activity. To answer the question

> "How to build the next generation of knowledge bases for spatial search?"

one must consider places in the context of spatial search. Throughout this thesis, essential information is produced to aid spatial search engines with decisions about activities. In Chapter 3, the main objective was to model cognitive regions according to the conceptualization of human beings. That in turn allows for comparisons of aggregated places. Chapter 4 deals with the automatic extraction of cognitive regions aided by Natural Language Processing and Machine Learning. In Chapter 5 a context model was proposed that is populated by User Generated Content. The main objective of the semantic graph structure is to aid humans when they search for implicit semantics of places. In all three chapters mentioned, knowledge produced by humans can be more intuitively consumed in spatial search tasks. In conclusion, the proposed techniques are aimed at supporting spatial searches according to the needs of the users.

6.6 Conclusion

In today's spatial search engines, users mostly depend on 'name' or categorical tags, or on already concrete address or location data to derive the information they are looking for. Throughout this thesis, it is argued that spatial search should be intuitive, and that naïve geography should be supported by the next generation of Geographic Information Systems.

The main objective of models proposed in Chapter 3 and Chapter 4 is the extraction of cognitive regions by exploiting both crowd-sourced geographic information and User Generated Content. Utilizing the distribution of categorical tags of OpenStreetMap in the form of a bag-of-words model, the spatial footprint of cognitive regions can be derived. This means that even when the area the user is looking for is not directly mapped in Geographic Information Systems, a uniform and homogeneous area can be derived solely by measuring the tag distribution in a city. The approach proposed in Chapter 3 is based on the intuitive assumption that space can be classified according to the conceptualization of human beings. In Chapter 4, the previous approach is extended to an automated approach which is inspired by text classification used in Information Retrieval. Based on the hypotheses that humans classify space according to the activities they can perform there and that cognitive regions have uniform and homogeneous characteristics, the Machine Learning approach has the potential to aid humans in spatial search tasks. By applying the proposed techniques, more intuitive questions can be answered, such as "where is downtown?" or "which shopping areas are more recreational with respect to different preference models?". While semantic similarity is a difficult question addressed in different research areas and not easy to solve, the results presented in this thesis can be seen as evidence that similarity can be tackled in ways analogous to traditional Information Retrieval. Furthermore, a knowledge base was proposed in Chapter 5 that incorporates activities and context information for spatial analysis. By using the proposed semantic model, knowledge produced by humans can be more intuitively consumed for spatial search tasks, as is shown in two different use cases. Especially when implicit semantics are involved, such as activities one can perform at a given place, the proposed approach shows promising results in enhancing spatial search engines.

In conclusion, the results presented in this thesis provide clear evidence that the integration of methods adapted from different fields has several advantages. Consequently, the central research question of this thesis can be answered with a resounding Yes:

> "Can cognitive regions be formalized and processed with the synergistic interplay of methods arising in different fields?"

CHAPTER 7

Summary and Future Work

This chapter concludes this thesis with a summary of the results achieved. It outlines possible future research directions in the field of bridging the gap between effective spatial search engines and cognitive science.

7.1 Summary

The rise of personal, mobile, and in general ubiquitous computing and applications making use of location data, has led to the creation of geo-data encoded as annotated points, lines and primitive geometric shapes (in the sense of GPS-coordinates, possibly with a tag such as river, bridge or restaurant). Using these application, it quickly became clear that not all questions could be modeled and answered satisfactorily with the available, mostly distance-based, spatial search engines. On a semantic level, spatial search can be more than a straightforward address or exact location matching. In this thesis, we concentrate on socially constructed reality (i.e. cognitive regions), which represents how humans perceive their spatial surroundings. While sophisticated methods exist to compute complex routes or plan trips, spatial search is inadequately supported for answering nuanced and fuzzy questions such as searching for 'shopping areas' and 'recreational regions' within a city. Throughout this thesis, it is argued that a generalization of areas bridges the gap between cognitive models of human understanding of space and the capabilities of today's spatial search engines. To consume spatial information more intuitively and in order to advance spatial search, several techniques are investigated in this thesis. The central argument is that spatial search is characterized by human activity. In general, the idea of cognitive regions is discussed, and two techniques are investigated to derive these regions in the same manner a human would conceptualize them. The semantic similarity of places is investigated, and distance metrics are applied which allow a judging of semantic similarity for aggregations of places. To process spatial information produced by humans a semantic model is introduced, and a technique is presented to transform text expressed in natural language into a machine-processable form.

The phenomena of cognitive regions are based on the fact that humans can conceptualize and generalize space according to the activities they can carry out at a given place, leading to uniform and homogeneous areas. For instance, people can, according to the task ‘shopping’, think of suitable regions to fulfil their list of subtasks. Therefore, in Chapter 3, we discussed the generalization and segmentation of cities into functional regions, which is inspired by the method region growing used in image segmentation. The results of this segmentation are simplified “semantic representations” which correspond to human conceptualization of space. It is argued that the mathematical abstraction in terms of feature vectors encapsulates sufficient knowledge to arrive at a meaningful computational representation of cognitive regions. The employed mathematical abstraction lends itself naturally as a basis for measuring the semantic similarity in terms of distance metrics. In a second step, the integration of expressed human requirements and preferences can be achieved by selecting services only according to the needs of the user. We evaluated the approach in a use case scenario, where conceptual places of London and Vienna are extracted and their similarity is judged based on the offered *place affordances*. We noted that the approach discussed in Chapter 3 relied on two premises, namely that the initial seeding cells are known and the partitioning ruleset is modeled according to a semantic-region-specific formula, the existence of which is assumed. Since these two premises are not automatically fulfilled, this issue naturally links to the next chapter.

In order to overcome the problems of initializing a training set for a cognitive region and asses multiple regions based on their similarity, we developed an automated approach which is detailed in Chapter 4. Our proposed approach is twofold: firstly, we process an additional data source of textual place descriptions and exploit this data to derive the initial area to which the descriptions refer to; and secondly, with the help of Machine Learning we are able to classify the remaining parts of a city. For the recognition of unique geo-features, we developed a Named Entity Recognition approach tailored to the OpenStreetMap knowledge base. The Machine Learning model utilizes the simplified semantic representation of regions (i.e. bag-of-words model of offered services) to measure and decide about semantic similarity of regions. Using this automated approach, large-scale analysis of different regions is made possible, which enables researchers and urban planners to explore cognitive regions from their respective perspectives in an automated manner. The results also confirm the applicability and strength of the proposed bag-of-words model as the underlying fundamental mathematical modeling methodology. However, in this model it is assumed that there exists a collection of services that are recognised by people. The next chapter, therefore, deals with a more sophisticated Natural Language Processing approach to populate a model of place as emergent of possible activities.

Our efforts so far to successively replace any a priori manual steps in all of the proposed methodologies with automated procedures now culminate in the introduction of the proposed semantic model. For the questions discussed in this thesis, the inception of a semantic model has important implications. It allows us to automate the derivation of activities from textual descriptions, strengthening the approach discussed in Chapter 3. The semantic model also allows us to automatically infer services which are associated to activities by people, strengthening the Machine Learning approach discussed in Chapter 4. Independent of their beneficial impact to the previous approaches, their most striking property is that also they can serve as the fundamental data structure to model semantic regions for spatial search engines. As has been shown in Chapter 5, the notion of the semantic model offers an interpretation as a spatial knowledge base that is able to incorporate the theory of cognitive regions. The canonical way to interpret queries in this model is to reduce the problem to a graph-matching problem. Therefore, the reasoning methodology in semantic graphs is different to and independent from the reasoning approaches based on the bag-of-words model. Although semantic graphs do possess many desirable features, we do not consider them a replacement for the bag-of-words model.

In conclusion, this thesis demonstrates how the presented computational models can be integrated in a holistic way to enable cognitively enhanced Geographic Information Systems. Synergies of different fields are utilized to bridge the gap between cognitive science and successful expert systems. The examples shown in this thesis serve as proof-of-concepts and, in combination, serve as proof that the cognitive understanding of place can be modeled with the proposed techniques in spatial knowledge bases and search engines.

7.2 Future Research Directions

The work presented in this thesis can be extended in several ways. In the following, seven possible future research directions are outlined.

7.2.1 Individual Learning

In the near future, location-aware routing systems are going to consider the cognitive interpretation of the surroundings of the user in their reasoning process. For instance, they could switch the level of detail when the driver is about to enter a familiar area, or they could give instructions based on cognitive regions from the view of the driver. For enhanced usability, the next generation of Geographic Information Systems should incorporate the option to derive cognitive regions from instructions expressed in natural language. Future research directions could, consequently, deal with individual learning instead of deriving a cognitive region based on the collection of reports submitted by multiple users.

7.2.2 Spatial Clustering

In Section 3.2, a grid-based clustering approach was introduced for spatial region clustering. Analyzing other clustering paradigms such as clusters created by k-nearest neighbors or R*-trees could be further investigated. Of special interest is the development of a system that detects co-occurring activities in spatial data sets [Ballatore, 2014, Hobel et al., 2015]. This could be another approach to answering nuanced questions about the surroundings. Different shopping areas, for example, afford different activities: those near industrial estates at the periphery of a city are mostly reachable by car only, and offer long-term activities such as “buying furniture” [Golledge and Stimson, 1997].

7.2.3 Spatial Vocabularies and Reasoning

In Section 3.3, an algorithm was introduced that ‘grows’ cognitive regions based on the distribution of categorical tags. This model was then extended in Chapter 4, where spatial entities are retrieved from unstructured text to determine the seeding cells. However, place descriptions expressed in natural language can also refer to spatial arrangements such as topological information. To incorporate this information into the next generation of Geographic Information Systems, approaches using Natural Language Processing can make use of qualitative spatial reasoning to derive the spatial arrangements of given entities [Chang et al., 2014a]. As a starting point, spatial vocabularies could be created to derive geographic arrangements of given concepts. The next research direction would be to develop a novel reasoning approach, which is more “intelligent” and can incorporate spatial arrangements when clustering cognitive regions.

7.2.4 Context Graph Refinement from a Natural Language Processing Perspective

The semantic representation of place proposed in Chapter 5 is just a starting point of more complex adjustments to capture individual impressions of the surroundings. Contextual information is stored in a single leaf node of an activity branch. Breaking the contextual information into a more fine-grained data format could reveal more about the actual semantics of a sentence. However, in order to refine the proposed approach a deeper understanding of the current state of Natural Language Processing would be required, which, consequently, could be a promising topic for future work.

7.2.5 Optimizing Database Management Systems (DBMSs) for Knowledge Bases

One aspect, which is not part of this thesis, is the performance of the utilized graph database. In real-life applications, the graph database has to handle not only search queries but also update and optimization measures. Moreover, the semantic graph database has to be extended to deal with parallelization. Therefore, possible further investigation directions could concern the optimization of DBMSs and an investigation of different types of graph databases.

7.2.6 Future Research in the Field of Artificial Intelligence

The used data sources inherit the natural properties of User Generated Content, i.e. incorrectness, incompleteness, and irregularity. To deal with these properties, future investigations could be concerned with methods of Artificial Intelligence to deal with these issues. For instance, in Section 4.2.2 combinatorial place matching was introduced, where one possible starting point could be the investigation of spheres of neighboring words that can be corrected. This would be similar to a prominent problem of coding theory – error correction.

7.2.7 Semantic Similarity and Anti-Unification

To quantify the cognitive notion of semantic similarity, we applied a specialized distance metric based on the vector space model in Section 3.4.1. Techniques from anti-unification could be analyzed regarding their suitability for operating on semantic models of place in order to asses their semantic similarity. This way of classification would enable a different way of reasoning based on primitives developed in theoretical computer science.

List of Figures

List of Tables

List of Algorithms

Index

Glossary

Artificial Intelligence The intelligence that software incorporates.. 77

bag-of-words model The bag-of-words model is a simplifying representation used in Natural Language Processing and Information Retrieval.. 8, 13, 23, 41, 44, 69–71, 74, 75

Geographic Information Retrieval A specialization of traditional information retrieval focused on geographically tagged content.. 2, 3, 5, 26, 33, 71

Geographic Information System A system that allows to store and process geographical data.. xi, 1–3, 5, 7, 14, 35, 43, 44, 53, 71, 75, 76

image segmentation Image segmentation is the process of partitioning an image into multiple segments.. 3, 5, 8, 11, 14, 19, 20, 31, 69, 74

Information Retrieval The process of retrieving information according to the needs of the user.. 2, 3, 5, 11, 14, 17, 24, 33, 53, 54, 70, 72

Levenshtein distance The Levenshtein distance is a metric for measuring string similarity by the number of used editing steps, required to change on text into the other.. 38–40

Machine Learning A field of computer science that enables artificial intelligent systems by pattern recognition and learning approaches.. 2, 3, 5, 8–10, 12, 14, 15, 32–36, 40, 41, 44–46, 64, 66, 67, 70–72, 74, 75, 79

Natural Language Processing The methods used to understand and process information expressed in natural language.. ix, xi, 2, 3, 5, 10, 13, 33–36, 54, 58, 63, 64, 67, 70, 71, 74, 76

OpenStreetMap A collaborative project to capture and provide geographic information for public use.. 1, 4, 5, 13, 17, 18, 23, 27, 45, 46, 52, 54, 60, 64, 65, 71, 74, 79

Part-of-Speech tagger In corpus linguistics, part-of-speech tagging is the procedure of marking the words in a text as corresponding to a particular part of speech, based on both its definition as well as its context.. 36, 37

Points of Interest The term used for geographic information that is of interest according to the needs of the user.. 2, 7, 9, 18, 64

region growing A simple image segmentation procedure.. 19, 31, 69, 74

regular expression A regular expression is a sequence of characters that define a string search pattern.. 38, 39

Resource Description Framework A framework that is used to conceptually represent information.. 4, 13, 54, 57, 62

User Generated Content Any form of content produced by human beings and often made publicly available.. ix, xi, 2–4, 9–11, 13, 15, 33, 35, 36, 45, 52–54, 58, 67, 69–71, 77

Volunteered Geographic Information The information captured by human sensors and provided fo open use.. ix, xi, 1, 3, 17, 18, 35, 45, 49, 64

Bibliography

Amin Abdalla and Andrew U Frank. Combining trip and task planning: How to get from a to passport. In *Geographic information science*, pages 1–14. Springer, 2012.

Benjamin Adams and Grant McKenzie. Frankenplace: An application for similarity-based place search. In *ICWSM*, 2012.

Benjamin Adams and Martin Raubal. A metric conceptual space algebra. In *Spatial information theory*, pages 51–68. Springer, 2009.

Benjamin Adams, Grant McKenzie, and Mark Gahegan. Frankenplace: Interactive thematic mapping for ad hoc exploratory search. In *Proceedings of the 24th International Conference on World Wide Web*, pages 12–22. International World Wide Web Conferences Steering Committee, 2015.

R. Adams and L. Bischof. Seeded region growing. *IEEE Trans. Pattern Anal. Mach. Intell.*, 16(6):641–647, #jun# 1994. ISSN 0162-8828. doi: 10.1109/34.295913. URL `http://dx.doi.org/10.1109/34.295913`.

J.A. Agnew. *Place and Politics: The Geographical Mediation of State and Society.* Allen & Unwin, 1987. ISBN 9780043201770. URL `https://books.google.at/books?id=9EUVAAAAIAAJ`.

Ahmed N. Alazzawi, Alia I. Abdelmoty, and Christopher B. Jones. What can i do there? towards the automatic discovery of place-related services and activities. *International Journal of Geographical Information Science*, 26(2):345–364, 2012. doi: 10.1080/13658816.2011.595954. URL `http://dx.doi.org/10.1080/13658816.2011.595954`.

Andrea Ballatore. The search for places as emergent aggregates. 2014.

Andrea Ballatore, David C. Wilson, and Michela Bertolotto. Computing the semantic similarity of geographic terms using volunteered lexical definitions. *International Journal of Geographical Information Science*, 27(10):2099–2118, 2013. doi: 10.1080/13658816.2013.790548. URL `http://dx.doi.org/10.1080/13658816.2013.790548`.

Andrea Ballatore, Michela Bertolotto, and David C. Wilson. An evaluative baseline for geo-semantic relatedness and similarity. *GeoInformatica*, 18(4):747–767, 2014. doi: 10.1007/s10707-013-0197-8. URL `http://dx.doi.org/10.1007/s10707-013-0197-8`.

Andrea Ballatore, Michela Bertolotto, and David C Wilson. A structural-lexical measure of semantic similarity for geo-knowledge graphs. *ISPRS International Journal of Geo-Information*, 4(2):471–492, 2015.

John A. Bateman, Joana Hois, Robert Ross, and Thora Tenbrink. A linguistic ontology of space for natural language processing. *Artificial Intelligence*, 174(14):1027 – 1071, 2010. ISSN 0004-3702. doi: http://dx.doi.org/10.1016/j.artint.2010.05.008. URL `http://www.sciencedirect.com/science/article/pii/S0004370210000858`.

Robert Battle and Dave Kolas. Geosparql: enabling a geospatial semantic web. *Semantic Web Journal*, 3(4):355–370, 2011.

Tim Berners-Lee, James Hendler, Ora Lassila, et al. The semantic web. *Scientific american*, 284(5):28–37, 2001.

Angel Chang, Will Monroe, Manolis Savva, Christopher Potts, and Christopher D Manning. Text to 3d scene generation with rich lexical grounding. *arXiv preprint arXiv:1505.06289*, 2015.

Angel X Chang, Manolis Savva, and Christopher D Manning. Interactive learning of spatial knowledge for text to 3d scene generation. *Sponsor: Idibon*, page 14, 2014a.

Angel X Chang, Manolis Savva, and Christopher D Manning. Learning spatial knowledge for text to 3d scene generation. EMNLP, 2014b.

World Wide Web Consortium et al. Rdf 1.1 concepts and abstract syntax. 2014.

Helen Couclelis. Location, place, region, and space. *Geography's inner worlds*, 2:15–233, 1992.

Helen Couclelis and Nathan Gale. Space and spaces. *Geografiska Annaler. Series B, Human Geography*, 68(1):pp. 1–12, 1986. ISSN 04353684. URL `http://www.jstor.org/stable/490912`.

Bob Coyne and Richard Sproat. Wordseye: an automatic text-to-scene conversion system. In *Proceedings of the 28th annual conference on Computer graphics and interactive techniques*, pages 487–496. ACM, 2001.

Eduardo Cunha and Bruno Martins. Using one-class classifiers and multiple kernel learning for defining imprecise geographic regions. *International Journal of Geographical Information Science*, 28(11):2220–2241, 2014. doi: 10.1080/13658816.2014.916040. URL `http://dx.doi.org/10.1080/13658816.2014.916040`.

Marie-Catherine De Marneffe and Christopher D Manning. Stanford typed dependencies manual. Technical report, Technical report, Stanford University, 2008.

Fausto D'Antonio, Paolo Fogliaroni, and Tom Kauppinen. Vgi edit history reveals data trustworthiness and user reputation. 2014.

MaxJ. Egenhofer and DavidM. Mark. Naive geography. In AndrewU. Frank and Werner Kuhn, editors, *Spatial Information Theory A Theoretical Basis for GIS*, volume 988 of *Lecture Notes in Computer Science*, pages 1–15. Springer Berlin Heidelberg, 1995. ISBN 978-3-540-60392-4. doi: 10.1007/3-540-60392-1_1. URL http://dx.doi.org/10.1007/3-540-60392-1_1.

Paolo Fogliaroni and Heidelinde Hobel. Implementing naive geography via qualitative spatial relation queries. 2015.

Scott M Freundschuh and Max J Egenhofer. Human conceptions of spaces: implications for gis. *Transactions in GIS*, 2(4):361–375, 1997.

Song Gao, Krzysztof Janowicz, Grant McKenzie, and Linna Li. Towards platial joins and buffers in place-based gis. In *Proceedings of the 1st ACM SIGSPATIAL international workshop on computational models of place (COMP'2013)*, pages 1–8, 2013.

J.J. Gibson. The theory of affordances. In Robert Shaw and John Bransford, editors, *Perceiving, Acting, and Knowing: Toward and Ecological Psychology*, pages 62–82. Erlbaum, Hillsdale, NJ, 1977.

J.J. Gibson. *The ecological approach to visual perception*. Houghton Mifflin, Boston, 1979.

R.G. Golledge and R.J. Stimson. *Spatial Behavior: A Geographic Perspective*. Guilford Press, 1997. ISBN 9781572300507. URL https://books.google.at/books?id=2JPMvpMbLrMC.

Michael F Goodchild. Citizens as sensors: the world of volunteered geography. *GeoJournal*, 69(4):211–221, 2007.

Michael F Goodchild. Formalizing place in geographic information systems. In *Communities, Neighborhoods, and Health*, pages 21–33. Springer, 2011.

Christian Grothe and Jochen Schaab. Automated footprint generation from geotags with kernel density estimation and support vector machines. *Spatial Cognition & Computation*, 9(3):195–211, 2009. doi: 10.1080/13875860903118307. URL http://dx.doi.org/10.1080/13875860903118307.

Heidelinde Hobel and Paolo Fogliaroni. Extracting semantics of places from user generated content. 2016.

Heidelinde Hobel and Andrew U. Frank. Exploiting linked spatial data and granularity transformations. In *Proceedings of the Workshop on Geographic Information Observatories 2014 co-located with the 8th International Conference on Geographic Information Science (GIScience 2014), Vienna, Austria, September 23, 2014.*, pages 15–22, 2014. URL http://ceur-ws.org/Vol-1273/paper2.pdf.

Heidelinde Hobel, Lisa Madlberger, Andreas Thöni, and Stefan Fenz. Visualisation of user-generated event information: Towards geospatial situation awareness using hierarchical granularity levels. In *Joint Proceedings of the 1th Workshop on Semantic Sentiment Analysis (SSA2014), and the Workshop on Social Media and Linked Data for Emergency Response (SMILE 2014) co-located with 11th European Semantic Web Conference (ESWC 2014), Crete, Greece, May 25th, 2014.*, pages 43–54, 2014. URL http://ceur-ws.org/Vol-1329/papersmile_3.pdf.

Heidelinde Hobel, Amin Abdalla, Paolo Fogliaroni, and Andrew U Frank. A semantic region growing algorithm: Extraction of urban settings. In *AGILE 2015*, pages 19–33. Springer, 2015.

Heidelinde Hobel, Paolo Fogliaroni, and Andrew U Frank. Deriving the geographic footprint of cognitive regions. In *AGILE 2016*. Springer, 2016.

Mohammed Imaduddin Humayun and Angela Schwering. Representing vague places: Determining a suitable method. In *Proceedings of the international workshop on place-related knowledge acquisition research (P-KAR 2012), Monastery Seeon, Germany*, volume 881, pages 19–25. Citeseer, 2012.

C. B. Jones, R. S. Purves, P. D. Clough, and H. Joho. Modelling vague places with knowledge from the web. *International Journal of Geographical Information Science*, 22(10):1045–1065, 2008. doi: 10.1080/13658810701850547. URL http://dx.doi.org/10.1080/13658810701850547.

Christopher B Jones, Harith Alani, and Douglas Tudhope. Geographical information retrieval with ontologies of place. In *Spatial information theory*, pages 322–335. Springer, 2001.

Troy Jordan, Martin Raubal, Bryce Gartrell, and M Egenhofer. An affordance-based model of place in gis. In *8th Int. Symposium on Spatial Data Handling, SDH*, volume 98, pages 98–109, 1998.

M. Karwan and A.U. Frank. *Cognitive and Linguistic Aspects of Geographic Space*. Nato Science Series D:. Springer Netherlands, 2012. ISBN 9789401126069. URL https://books.google.at/books?id=l1gPCQAAQBAJ.

Carsten Keßler and René Theodore Anton de Groot. Trust as a proxy measure for the quality of volunteered geographic information in the case of openstreetmap. In *Geographic information science at the heart of Europe*, pages 21–37. Springer, 2013.

Arbaz Khan, Maria Vasardani, and Stephan Winter. Extracting spatial information from place descriptions. In *Proceedings of the First ACM SIGSPATIAL International Workshop on Computational Models of Place*, pages 62–69, 2013.

Manolis Koubarakis and Kostis Kyzirakos. Modeling and querying metadata in the semantic sensor web: The model strdf and the query language stsparql. In *The semantic web: research and applications*, pages 425–439. Springer, 2010.

Werner Kuhn. Ontologies in support of activities in geographical space. *International Journal of Geographical Information Science*, 15(7):613–631, 2001. doi: 10.1080/13658810110061180. URL http://dx.doi.org/10.1080/13658810110061180.

Werner Kuhn, Tomi Kauppinen, and Krzysztof Janowicz. Linked data-a paradigm shift for geographic information science. In *Geographic Information Science*, pages 173–186. Springer, 2014.

Hans Peter Luhn. A statistical approach to mechanized encoding and searching of literary information, 1957.

K. Lynch. *The Image of the City*. Harvard-MIT Joint Center for Urban Studies Series. M.I.T. Press, 1960. ISBN 9780262620017. URL https://books.google.at/books?id=_phRPWsSpAgC.

Patrick Lüscher and Robert Weibel. Exploiting empirical knowledge for automatic delineation of city centres from large-scale topographic databases. *Computers, Environment and Urban Systems*, 37:18 – 34, 2013. ISSN 0198-9715. doi: http://dx.doi.org/10.1016/j.compenvurbsys.2012.07.001. URL http://www.sciencedirect.com/science/article/pii/S0198971512000609.

Christopher D Manning, Mihai Surdeanu, John Bauer, Jenny Finkel, Steven J Bethard, and David McClosky. The stanford corenlp natural language processing toolkit. In *Proceedings of 52nd Annual Meeting of the Association for Computational Linguistics: System Demonstrations*, pages 55–60, 2014.

Grant McKenzie, Benjamin Adams, and Krzysztof Janowicz. A thematic approach to user similarity built on geosocial check-ins. In *Geographic Information Science at the Heart of Europe*, pages 39–53. Springer, 2013.

Richard Meegan and Alison Mitchell. 'it's not community round here, it's neighbourhood': Neighbourhood change and cohesion in urban regeneration policies. *Urban Studies*, 38(12):2167–2194, 2001. URL http://EconPapers.repec.org/RePEc:sae:urbstu:v:38:y:2001:i:12:p:2167-2194.

George A. Miller. Wordnet: A lexical database for english. *Commun. ACM*, 38(11):39–41, #nov# 1995. ISSN 0001-0782. doi: 10.1145/219717.219748. URL http://doi.acm.org/10.1145/219717.219748.

Daniel R Montello. Scale and multiple psychologies of space. In *Spatial information theory a theoretical basis for gis*, pages 312–321. Springer, 1993.

Daniel R Montello. Regions in geography: Process and content. *Foundations of geographic information science*, pages 173–189, 2003.

Daniel R Montello, Michael F Goodchild, Jonathon Gottsegen, and Peter Fohl. Where's downtown?: Behavioral methods for determining referents of vague spatial queries. *Spatial Cognition & Computation*, 3(2-3):185–204, 2003.

Daniel R Montello, Alinda Friedman, and Daniel W Phillips. Vague cognitive regions in geography and geographic information science. *International Journal of Geographical Information Science*, 28(9):1802–1820, 2014.

Peter Mooney and Padraig Corcoran. The annotation process in openstreetmap. *Transactions in GIS*, 16(4):561–579, 2012. ISSN 1467-9671. doi: 10.1111/j.1467-9671.2012.01306.x. URL http://dx.doi.org/10.1111/j.1467-9671.2012.01306.x.

J.C. Pinoli. *Mathematical Foundations of Image Processing and Analysis*. ISTE. Wiley, 2014. ISBN 9781118649121. URL https://books.google.at/books?id=370BBAAAQBAJ.

Adrian Popescu, Gregory Grefenstette, and Pierre Alain Moëllic. Gazetiki: automatic creation of a geographical gazetteer. In *Proceedings of the 8th ACM/IEEE-CS joint conference on Digital libraries*, pages 85–93. ACM, 2008.

Martin Raubal. Human wayfinding in unfamiliar buildings: a simulation with a cognizing agent. *Cognitive Processing*, 2(3):363–388, 2001.

Daniela Richter, Maria Vasardani, Lesley Stirlng, Kai-Florian Richter, and Stephan Winter. Zooming in–zooming out hierarchies in place descriptions. In Jukka M. Krisp, editor, *Progress in Location-Based Services*, Lecture Notes in Geoinformation and Cartography, pages 339–355. Springer Berlin Heidelberg, 2013. ISBN 978-3-642-34202-8. doi: 10.1007/978-3-642-34203-5_19. URL http://dx.doi.org/10.1007/978-3-642-34203-5_19.

Gerard Salton, Edward A. Fox, and Harry Wu. Extended boolean information retrieval. *Commun. ACM*, 26(11):1022–1036, #nov# 1983. ISSN 0001-0782. doi: 10.1145/182.358466. URL http://doi.acm.org/10.1145/182.358466.

Beatrice Santorini. Part-of-speech tagging guidelines for the penn treebank project (3rd revision). 1990.

Theodore R. Schatzki. Spatial ontology and explanation. *Annals of the Association of American Geographers*, 81(4):pp. 650–670, 1991. ISSN 00045608. URL http://www.jstor.org/stable/2563428.

Simon Scheider and Krzysztof Janowicz. Place reference systems: A constructive activity model of reference to places. *Appl. Ontology*, 9(2):97–127, #apr# 2014. ISSN 1570-5838. URL http://dl.acm.org/citation.cfm?id=2766273.2766274.

Barry Smith. On drawing lines on a map. In *COSIT*, pages 475–484, 1995. URL http://dblp.uni-trier.de/db/conf/cosit/cosit95.html#Smith95.

Barry Smith and David M. Mark. Geographical categories: an ontological investigation. *International Journal of Geographical Information Science*, 15(7):591–612, 2001. URL http://dblp.uni-trier.de/db/journals/gis/gis15.html#SmithM01.

Claus Stadler, Jens Lehmann, Konrad Höffner, and Sören Auer. Linkedgeodata: A core for a web of spatial open data. *Semantic Web*, 3(4):333–354, 2012.

Yi-Fu Tuan. *Space and place: humanistic perspective*. Springer, 1979.

Barbara Tversky and Kathleen Hemenway. Categories of environmental scenes. *Cognitive Psychology*, 15(1):121 – 149, 1983. ISSN 0010-0285. doi: http://dx.doi.org/10.1016/0010-0285(83)90006-3. URL http://www.sciencedirect.com/science/article/pii/0010028583900063.

Maria Vasardani, Stephan Winter, and Kai-Florian Richter. Locating place names from place descriptions. *International Journal of Geographical Information Science*, 27(12):2509–2532, 2013. doi: 10.1080/13658816.2013.785550. URL http://dx.doi.org/10.1080/13658816.2013.785550.

Stephan Winter and Marie Truelove. Talking about place where it matters. In Martin Raubal, David M. Mark, and Andrew U. Frank, editors, *Cognitive and Linguistic Aspects of Geographic Space*, Lecture Notes in Geoinformation and Cartography, pages 121–139. Springer Berlin Heidelberg, 2013. ISBN 978-3-642-34358-2. doi: 10.1007/978-3-642-34359-9_7. URL http://dx.doi.org/10.1007/978-3-642-34359-9_7.

Stephan Winter, Werner Kuhn, and Antonio Krüger. Guest editorial: Does place have a place in geographic information science? *Spatial Cognition & Computation*, 9(3):171–173, 2009. doi: 10.1080/13875860903144675. URL http://dx.doi.org/10.1080/13875860903144675.

Yi Zhang, Yong Gao, LuLu Xue, Si Shen, and KaiChen Chen. A common sense geographic knowledge base for gir. *Science in China Series E: Technological Sciences*, 51(1):26–37, 2008. ISSN 1006-9321. doi: 10.1007/s11431-008-5003-8. URL http://dx.doi.org/10.1007/s11431-008-5003-8.

Dennis Zielstra and Alexander Zipf. A comparative study of proprietary geodata and volunteered geographic information for germany. In *13th AGILE international conference on geographic information science*, volume 2010, 2010.