Ontology-driven Enrichment of Spatial Databases

Patrick Lüscher, Dirk Burghardt, and Robert Weibel
Department of Geography, University of Zurich, CH-8057 Zurich, Switzerland
E-mail: {luescher, burg, weibel}@geo.unizh.ch

Keywords: Spatial data enrichment, pattern recognition, ontologies, urban structures

Abstract

Generalization is an abstraction process by which characteristics of spatial patterns should be preserved and highlighted. This requires the patterns to be detected beforehand. Additionally, automated enrichment of spatial data is of growing importance for many mapping agencies in order to respond to varying user needs. In this paper we present a framework for pattern recognition in urban environments that complements current algorithm-centered approaches by first formalizing spatial patterns in ontologies, and then deductively triggering appropriate low-level pattern recognition techniques. We start our paper by giving an introduction to the terminology of ontologies. Existing work on pattern recognition using semantic models is reviewed. We then outline our general framework and exemplify an ontological model of an urban structure for a case study we are currently working on. Finally, we discuss issues, benefits and challenges of the approach.

1. Introduction

Patterns play an important role during the generalization process: Since their characteristics need to be preserved, they provide a basis for an appropriate selection and parameterization of generalization algorithms. However, most of the spatial databases that exist today have been designed to serve multiple purposes and hence concentrate on the 'least common denominator'. Data models are usually simple in the sense that they define basic features such as buildings and roads. Therefore, existing databases have to be enriched with patterns that have to be extracted by means of automated pattern recognition techniques (Brassel & Weibel 1988; Ruas & Plazanet 1996).

For mapping agencies, automated enrichment of existing spatial databases with specific higher level concepts allows responding better to customer needs and is therefore useful for many applications. Some concrete examples for the urban domain might be the derivation of the construction period of particular buildings to infer the typical copper concentration per building, a more advanced application might be to connect patterns with urban evolution processes (Camacho-Hübner & Golay 2007), or improved adaptation in mobile services such as navigation by considering spatial contexts specified in the database (Winter 2002).

In the urban context, many specialized pattern recognition algorithms have been employed for detection of structures (Regnauld 1996; Barnsley & Barr 1997; Anders et al. 1999; Boffet 2001; Heinzle et al. 2005; Steiniger 2006a). In the main, these are 'bottom-up' in the sense that they first specify a (often visual) pattern to recognise, derive its (geometrical) properties, and use some elaborated detection algorithm (figure 1, left branch).

Then again, it has been argued that for better adaptation to varying applications, approaches that model the concepts to be derived are needed. For example, it has been pointed out by Mackaness (2006) that abstraction of large-scale databases to very general concepts requires the roles of the individual features and patterns they form to be understood and modeled explicitly. Dutton & Edwardes (2006), Kulik (2005) and Redbrake & Raubal (2004) show the importance of semantic modeling of geographic features in maps to guide user adaptation during generalization.

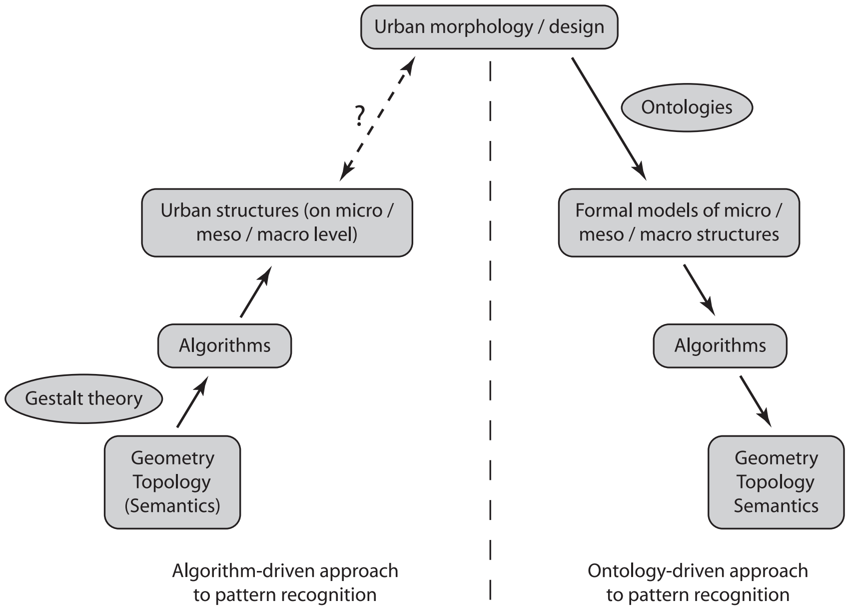

Figure 1. Bottom-up vs. top-down approaches to pattern recognition in urban areas

In our research project we aim at developing methods for the integration of rich semantic concepts into existing spatial databases of the urban domain. The approach we pursue is 'top-down' as shown in figure 1, right branch: We study the literature on urban morphology and urban design in order to identify specific urban patterns. The next step is to formalize these patterns, their context and hierarchical composition using ontologies. The formal definitions of patterns are then used to deductively trigger appropriate 'low-level' pattern recognition techniques in order to detect them in real databases. We hope that this way we can overcome some important drawbacks of the methods employed nowadays:

Firstly, current pattern recognition methods have often been developed and parameterised for specific databases. However, urban patterns are highly dependent on the cultural background and topographic conditions. For example, the German national atlas (*Nationalatlas Bundesrepublik Deutschland*, Friedrich et al. 2002) describes specific settlement forms (*Angerdorf*, *Hufendorf*, *Gutsdorf*) that cannot be found in other countries such as the UK, which in turn has its own very specific settlement patterns. Therefore, in an ideal approach a domain expert would model important patterns in a formalized language and then have tools available that convert the models automatically to pattern recognition processes.

Secondly, existing pattern recognition algorithms are often not flexible enough to include additional information, such as topography, which may be important to describe the genesis of certain urban forms. Ontologies are a promising means to achieve this integrative role (see Klien & Lutz 2005 for an application example).

Finally, more explanatory power will be contained in the final classifications, since a natural language description of the model can be generated upon request. The network of interlinked concepts can be used for versatile abstraction processes.

The structure of this paper is as follows: After an introduction to the terminology of ontologies (§ 2), we will give an overview of related research in pattern recognition using ontologies (§ 3). We will then state the methodology of our approach and the research issues connected to it (§ 4). Finally, we draw some conclusions of our preliminary work and report on our current and future work on this topic (§ 5).

2. Ontological Modeling

Since ontologies are used in many different contexts, we want to first clarify our understanding of the term. The roots of ontologies lie in philosophy, where the term Ontology is understood as "the science or study of being". It is a specification of "what constitutes reality" in the form of taxonomies (Agarwal 2005). It is independent of epistemology, and since there can be only one reality, there is also only one Ontology, hence the big 'O' and the singular use of the term.

In the last decade, ontologies have attracted large interest in the artificial intelligence community. In AI, an ontology is understood as an explicit specification of a conceptualization (Gruber 1993). A conceptualization is an abstract, simplified view of the world that we want to represent for some reason. Each concept has a concept name (e.g., *ResidentialHouse*), some properties ('number of floors', 'area'), and a set of relations (Rodríguez & Egenhofer 2004).

While this definition reveals some similarities to classic object-oriented modeling, there are some significant differences: Firstly, ontologies are linked hierarchically to higher-level ontologies such that the semantics of concepts is globally clearly defined (section 2.1). Secondly, concepts in ontologies are rich in semantically defined relations to other concepts (section 2.2). Thirdly, ontologies can be specified in machine-interpretable languages that allow automatic inference (section 2.3). Therefore, while object-oriented models define relations on *data*, ontologies define terms with which to represent *knowledge* (Gruber 1993).

2.1 Levels of ontologies

There exists no universally accepted classification of ontologies. For our purposes, we distinguish between three types according to the specialisation of the represented concepts that is similar to the one defined in Guarino (1998) and Fonseca et al. (2002):

- Top-level ontologies: They define very general concepts such as space, time, matter, object, event, action, etc. which are independent of a specific domain or problem. One example of top-level ontology is the SNAP/SPAN ontology by Grenon & Smith (2004) that generally distinguishes between two types of entities. On the one hand objects have a continuous existence through time. On the other hand processes, events, and activities are bound in time – they exist only in their successive temporal parts or phases (Grenon & Smith 2004).
- Domain ontologies: They describe the terminology of a certain domain (such as medicine), or of a general task. We will describe necessary domain ontologies for urban pattern recognition in section 4.
- Application ontologies: They describe the terms that are on the one hand dependent on a domain, and on the other hand on a very specific task.

The key point is that every level builds on the terms that have been defined in a higher-level ontology. In our framework, basic terms that are needed to trigger the recognition of higher-level concepts would be described as domain ontology. These basic terms comprise single features such as a residential house, and the necessary spatial relations (connected, adjacent, etc.).

2.2 Types of relations

Thus, an ontology is essentially a set of concepts. Concepts can be associated with each other through relations. When modeling entities with ontologies, we can distinguish three types of relations (Rodríguez & Egenhofer 2004 and Fonseca et al. 2002):

- Taxonomic relations: These define sub-concepts and thus create a hierarchy of concepts. For instance, a single family home is a sub-concept of *ResidentialHouse*, which is again a sub-concept of the general concept *Building*.
- Roles: They allow adapting ontologies to specific user views by dynamically assigning concepts to each other. For example, the role *spatialFootprint* for a *Building* can be either played by a polygon, or by a point.
- Partonomic relations: With partonomic relations, aggregate concepts can be defined from a set of basic concepts. Thus, a R*esidentialNeighbourhood* is composed mainly of instances of the concept *ResidentialBuilding*.

Spatial patterns are aggregate concepts that are characterized by the spatial arrangement of the individual parts. For their description, spatial relations have to be defined additionally. For example, “a floodplain is a meadow that is *adjacent* to a river” (Klien & Lutz 2005). Topological relations like *contains* or *touches* are a special class of spatial relations, but also the statement that several houses *are aligned* can be conceptualized as a spatial relation.

When using ontologies for the classification of real data, one wants to find out whether a specific set of objects satisfies all requirements to be classified as an instance of a specific concept. Hence, spatial relations form predicates that have to be evaluated by mapping them to geospatial processing operations (Peachavanish & Karimi 2007). For example, the topological relations mentioned above can be evaluated by the 9-intersection model (Egenhofer & Herring 1991).

One of the main problems is that spatial relations are often fuzzy and hence, the same semantic relation can have different implementations or parameterisations, depending on the context it is used in. For the above mentioned example of floodplains, *adjacent* actually denotes all areas low enough in order to be flooded by the nearby river. If *adjacent* is implemented as a buffer operation, how large should the buffer width be chosen?

2.3 Reasoning with Description Logics (DL)

Ontologies can be specified in a Description Logics (DL) language. In description logics, generally two types of knowledge are represented (Neumann & Möller 2004): A set of axioms (describing a concept) is referred to as terminological box or as *TBox*; factual (assertional) knowledge about the world is called an *ABox*. Let’s clarify the difference between TBoxes and ABoxes with two examples:

- The definition of a floodplain as “a meadow that is *adjacent* to a river” can be formalized in a DL language and states a concept of the TBox. We can tag all areas in a spatial database that satisfy the definition with “Floodplain”. Hence, these areas are part of the ABox.

- "A football stadium is a sports facility which is used for playing football" (Rodríguez & Egenhofer 2004) defines football stadium as a sub-concept of sports facilities in a TBox. The ABox of a London database comprises Highbury Stadium, Matchroom Stadium, Griffin Park, etc.

DL reasoners allow various types of inferences, of which the following might prove to be of importance to our project (from Neumann & Möller 2004):

- whether a concept is subsumed by another concept
- whether an ABox is consistent w.r.t. a TBox;
- whether an individual is an instance of a concept;
- what are the most-specific atomic concepts of which an individual is an instance;
- what are the instances of a concept;
- what are the individuals filling a role for a specified individual;
- what pairs of individuals are related by a specified role; and
- general queries for tuples of individuals mentioned in ABoxes that satisfy certain predicates (so-called conjunctive queries).

Formalizing urban patterns as ontologies reveals some exciting possibilities: As we hope, reasoners can be used to automatically associate instances with concepts; on the other hand, having an ontology-enriched database (enriched manually, or by another system), we can test whether and to which extent it is consistent with our own description.

3. Related work

We will summarize in this section previous and ongoing work that uses explicit semantic models for recognition of spatial patterns.

For computer vision, Neumann & Möller (2004) present an approach to using a DL for high-level scene interpretation. They point out that there has been a gap between low-level vision, which involves techniques for image segmentation and object recognition, and high-level vision, where interpretation tasks may be highly context dependent and knowledge-intensive. They show how specific configurations of objects constrained by temporal and spatial relations such as a table-laying scene for breakfast can be represented by a Description Logic ALCF(D) and sketch a method for using reasoning services as components for the interpretations.

Notable work on semantics-driven interpretation of spatial data has been done in remote sensing for automatic classification of aerial photographs. De Gunst & Vosselmann (1997) present a model-driven approach for the detection of roads using semantic networks. For instance, a two-lane road can be described by three white lines, where the middle line is dashed. Sester (2000) and Anders & Sester (1997) build semantic models for the automatic interpretation of large-scale databases, i.e. they extract different types of houses, streets, parcels and built-up areas from polygon data. The inductive machine learning algorithm ID3 is used to discover relevant spatial properties and relations in manually tagged data. An approach for combining DL with spatial reasoning to formalize spatial arrangements is presented by Haarslev et al. (1994). They propose to combine the reasoning mechanism with a spatial index in order to speed up calculations.

Many spatial concepts are inherently vague. Santos et al. (2005) use supervaluation semantics to integrate vagueness into logical reasoning. They show a prototype implementation which classifies water bodies according to an inland water feature ontology. The inference process is carried out in Prolog.

Ontologies are a means to achieve semantic interoperability in a distributed environment. In this context, Klien & Lutz (2005) discuss the automatic annotation of existing datasets with concepts defined in an ontology. Their approach emphasises spatial relations between features rather than individual feature properties.

Tina Thomson's work aims at building land use maps from OS MasterMap data. Therefore, she intends to use ontologies to model land use categories according to the specific spatial configurations, compositions, relations and other special characteristics (Thomson 2006).

A project of the Ordnance Survey aimed at identifying fields such as farming land or pasture in OS MasterMap data. They used ontologies in order to describe relevant field properties (Kovacs & Zhou 2007).

4. Ontology-driven pattern recognition

4.1 General approach

In this section we will outline our methodology for investigating the role of ontologies in pattern recognition and the benefit of ontology-enriched spatial databases.

Figure 2 shows the general framework. A domain expert (cartographer or urbanist) models the urban structures he/she wants to recognize. The model includes geometrical and semantic components which are needed for their automatic detection and hierarchical composition of patterns, e.g., the pattern might usually be part of an inner city area, which could be either used to restrict the search area for the pattern given inner city areas, or to gain hints for the detection of inner city areas. The model can also include contextual information such as a geographical region for which the pattern is defined, e.g., specific for UK or Israel, and the functional role it plays in a specific context, such as the connection to an urban development process ontology, and thus allow the abstraction to application specific representations.

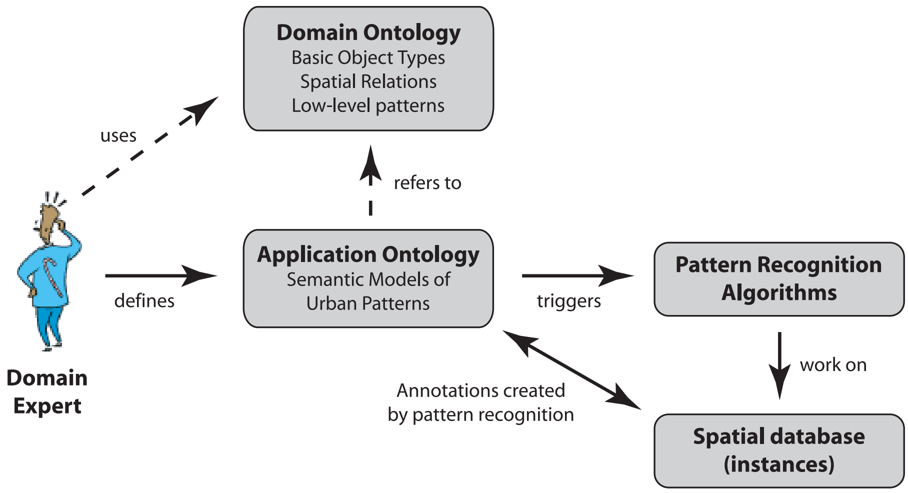

Figure 2. Workflow of the enrichment process using semantic models of urban patterns.

These specific models of patterns which we termed 'high-level patterns' constitute application ontologies. We will provide an example for a high-level pattern in the next subsection. In order to be able to define them, a basic vocabulary is needed which is provided as a set of domain ontologies. The 'GIS/cartography' ontology provides concepts for space representation (point, polygon, etc.) and spatial relations (adjacent, within, etc.). There exists also a set of 'low-level patterns' such as alignments and ring structures (buildings), grid patterns and star-shapes (roads), or southern slopes (topography) that are adopted when describing high-level patterns. Another domain ontology is therefore constituted by these low-level patterns.

Ontologies describe a set of concepts and relations between concepts. In order to do the actual data enrichment, a pattern recognition system has to interpret the models and transfer them to a series of spatial processing operations that can be carried out in a GIS environment. To this end, we directly link low-level concepts to spatial algorithms: The pattern recognition system knows how to handle concepts that describe spatial predicates and properties for spatial measures; furthermore, the low-level patterns mentioned above are identified using traditional pattern recognition algorithms. High-level patterns should then be detected automatically by triggering appropriate procedures for measurement of geometrical properties and detection of low-level patterns. Finally, the existing spatial database is annotated with detected low-level and high-level patterns, i.e. links between database objects and concepts are created.

4.2 Formalizing perimeter block developments

In a case study, we are currently working on the formalization of the high-level pattern 'perimeter block developments'. They were a dominant architectural style in Europe from 1880 to 1920 and, as the name implies, perimeter block developments are constituted by buildings that are aligned at the frontage around a rectangular courtyard. Some of the courtyards were originally occupied by workshops, but they were often removed later. Figure 3 shows an extract of a typical perimeter block development area in the City of Zurich.

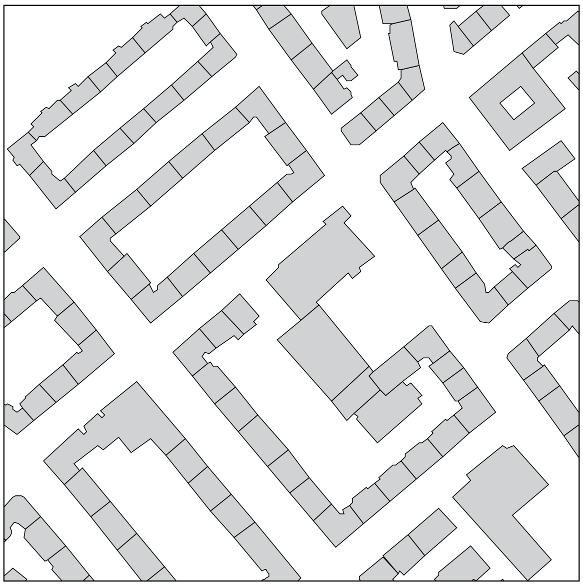

Figure 3. Typical perimeter block development in the City of Zurich. Source: General plan of Zurich 1:2500.

Figure 4 and 5 show extracts of an ontology that might be built for the urban concept PerimeterBlockDevelopment. We can see that the GIS/Cartography domain ontology also specifies a concept 'Scale', which is important because characteristics of urban structures may depend largely on the scale for which they are defined. For GIS processing functions, it has been proposed that the OGC Simple Feature Specification could be used as a basic domain ontology (Peachavanish & Karimi 2007). The urban morphology defines basic concepts such as urban block or inner city area, which are defined as sub-concepts of Micro- and MesoStructures, respectively. The arrows denote semantic relations of the concept PerimeterBlockDevelopment to its geographical and architectural context. This may be used for example to extract all areas that are instances of inner city concepts in Europe. Thus, through these links, abstraction processes can be formally defined.

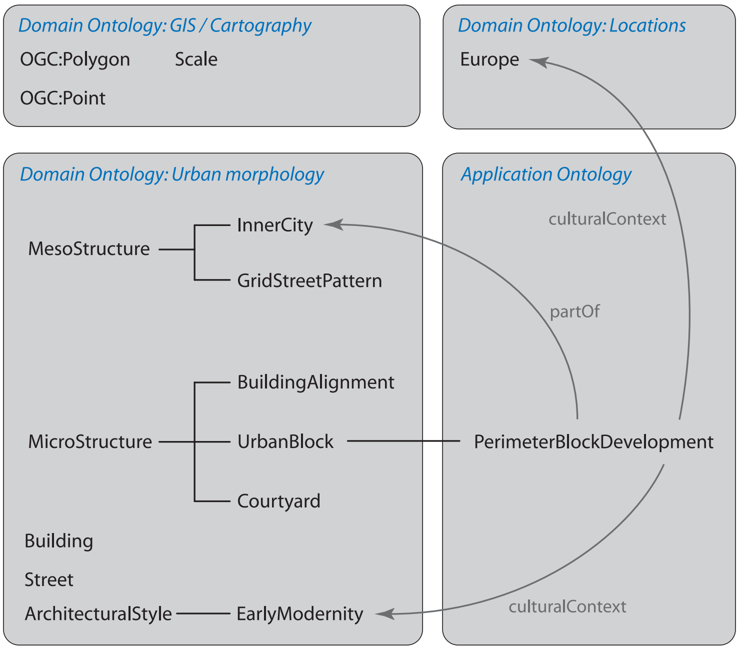

Figure 4. Connection of the concept PerimeterBlockDevelopment with its cultural context.

Contextual links allow to flexibly abstract and browse spatial information contained in the database. In order to actually enrich databases with defined concepts, their spatial and functional characteristics have to be encoded in the ontology. Spatial characteristics may include the compositional structures that may be formed from low-level patterns, as well as geometric measures such as typical building sizes. Figure 5 shows a preliminary attempt at linking PerimeterBlockDevelopment to lower-level patterns. Since perimeter block developments typically constitute a grid street pattern, there exists a containment relationship between these concepts. Furthermore, perimeter block developments consist of building alignments, which is also formalized as a containment relationship. A topological relationship between building alignment and street states that the alignments have to be arranged along streets.

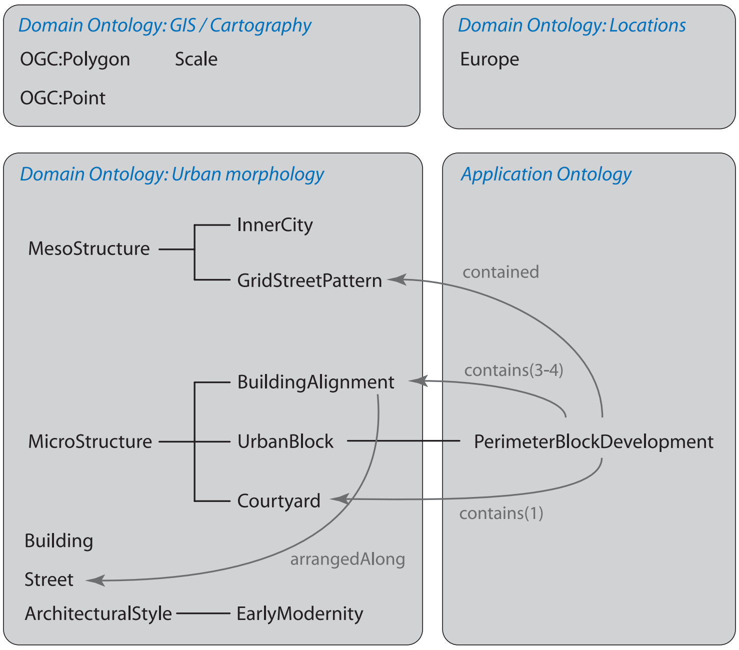

Figure 5. Attempt at linking PerimeterBlockDevelopment to its spatial characteristics.

4.3 Research issues

During the first part of our project, the emphasis is on identification and formalization of specific urban concepts. Later, we will have to look at issues concerning the design of the pattern recognition system. Generally, we pursue the following objectives:

1) *Identification and formalization of relevant urban concepts and their spatial properties*. This issue has mainly been addressed by a review of the relevant literature about urban forms and architecture. The formalization of the pattern knowledge is carried out using Protégé (Protégé 2007).

2) *Transformation from ontologies to algorithms that allow their automatic detection in existing spatial databases*. As stated before, we investigate the deployment of automatic reasoning techniques for triggering low-level recognition procedures from ontological descriptions. Commercial reasoners are available off-the-shelf, but they possess no spatial processing capabilities. Reasoners allow to import external functionalities as predicates and functions, so that they can be connected to a GIS environment such as JUMP/JTS (Vivid Solutions 2007).

3) *Actual enrichment of databases with the previously established ontological concepts*. This includes finding an appropriate data model for the connections between ontological concepts and the set of data base objects which instantiate the concepts. Since the concepts (the TBox model) are to be permanently connected to real data (the ABox) which naturally reside in a spatial database, data models have to be found which allow efficient traversal and machine interpretation. It may also be advantageous to store the classification history: If an object is changed during an update, it may affect the patterns it is related to (Haarslev et al. 1994). Another motivation might be that users can retrieve not only patterns, but also the reasons why a concept has been instantiated as such (for example as a textual explanation).

4) *Design of intuitive human-computer interaction methods with the pattern recognition system*: Protégé may be too complex for domain experts. Therefore, we investigate a specific user interface for creating spatial patterns and verify results of detected instances.

4.4 Benefits and challenges of the approach

Compared to the conventional method of building specific algorithms for pattern recognition, our approach has several benefits:

- Properties of patterns are explicitly stated instead of hidden in algorithms. Hence, we will have more explanatory power in the final classifications.
- Pattern recognition will be adaptable to different cultures or contexts by adapting pattern specifications, without actually having to alter the recognition engine.
- Knowledge discovery, representation, and exploitation are integrated within one global framework.
- As already mentioned in section 2.3, different ways of utilizing the system can be envisioned: On the one hand, it can be used to verify whether a concept is formalized consistently with regard to a certain reality. On the other hand, machine learning techniques can be used for exploring spatial relations that characterize concepts, and hence help domain experts to formalize patterns.

On the other hand, we can identify some issues that may cause difficulties or imply significant drawbacks:

- The semantics of natural language terms denoting spatial relations has been addressed within qualitative spatial reasoning research (Frank 1996). The same term may have different meanings within different contexts (ambiguity of terms), and they are often inherently vague. There is still a lack of knowledge regarding the roles of spatial relations terms in cognitive science research, which may hinder the translation of natural language descriptions into processing chains.
- Similarly, there is also ambiguity and vagueness of concepts. While formalisms to represent ambiguity in ontologies do exist, vagueness has not been profoundly treated so far. The method proposed in Santos et al. (2005) is simplistic since it relies on fixed thresholds. A more natural way to deal with vagueness would be to determine a value of certainty to which a set of objects is trusted to constitute a concept.
- Compared to conventional algorithms, the efficiency of the (spatial) reasoning process may be poor and hence prove to be a significant bottleneck.
- Klien & Lutz (2005) mention that it may not possible to find a fully automated process. In this respect, it is sensible to build a user interface that guides the domain expert through the recognition process and asks for help, where no automatic recognition is possible.

5. Conclusions

In this paper, we investigated the application of ontologies for describing spatial patterns. We believe this would be a sound basis for reasoning about which features and relations are important and hence have to be preserved in automated generalization. In this respect, ontologies are a means to make spatial databases more intelligent. Therefore, methods are needed to connect real data with ontological concepts.

In section 2, we have introduced the terminology and presented three different levels of ontologies. One conclusion is that application ontologies can be utilized to formalize urban structures.

Section 3 comprises a review about relevant research on spatial pattern recognition using semantic models. As it is pointed out, there has been some work on the conceptual level, but the feasibility for complex real-world problems needs to be proven.

In section 4, we have presented a methodology for semantic enrichment. The approach is to model high-level concepts in an ontology, whereas low-level pattern recognition procedures are automatically triggered.

The next steps in our work will be to complete the pilot study concerning the perimeter block developments, i.e. to enhance the ontological model and to build a processing chain for their actual detection in spatial databases. Furthermore, we also intend to build a taxonomy of salient urban patterns.

Acknowledgments

The research reported in this paper is part of the PhD project of the first author. Funding by the Swiss State Secretariat for Education and Research (SER) through COST Action C21 (project ORUS, grant no. C05.0081) is gratefully acknowledged.

References

Agarwal, P. (2005): Ontological considerations in GIScience. In: *International Journal of Geographical Information Science*, **19** (5), 501–536.

Anders, K.-H. and Sester, M. (1997): Methods of Data Base Interpretation – Applied to Model Generalization from Large to Medium Scale. In: Förstner, W. and Plümer, L. (Eds.) (1997): *Semantic Modeling for the Acquisition of Topographic Information from Images and Maps. SMATI 97*, Bonn Bad Godesberg, Germany, May 21–23, 1997, 89–103.

Anders, K.-H., Sester, M. and Fritsch, D. (1999): Analysis of Settlement Structures by Graph-Based Clustering. *SMATI'99 - Semantic Modelling for the Acquisition of Topographic Information from Images and Maps*, Munich, Germany, September 7.

Barnsley, M. J. and Barr, S. L. (1997): Distinguishing urban land-use categories in fine spatial resolution land-cover data using a graph-based, structural pattern recognition system. In: *Computers, Environment and Urban Systems*, **21** (3–4), 209–225.

Boffet, A. (2001): *Méthode de création d'informations multi-niveaux pour la généralisation cartographique de l'urbain*. Ph.D. thesis, Université de Marne-la-Vallée.

Brassel, K. E. and Weibel, R. (1988): A review and conceptual framework of automated map generalization. In: *International Journal of Geographical Information Systems*, **2** (3), 229–244.

Camacho-Hübner, E. and Golay, F. (2007): Preliminary insights on continuity and evolution of concepts for the development of an urban morphological process ontology. In: Teller, J., Lee, J. R. and Roussey, C. (Eds.) (2007): *Ontologies for Urban Development*. Studies in Computational Intelligence, Vol. 61, 95–108.

de Gunst, M. and Vosselman, G. (1997): A Semantic Road Model for Aerial Image Interpretation. In: Förstner, W. and Plümer, L. (Eds.) (1997): *Semantic Modeling for the Acquisition of Topographic Information from Images and Maps. SMATI 97*, Bonn Bad Godesberg, Germany, May 21–23, 1997, 107–122.

Dutton, G. and Edwardes, A. (2006): Ontological Modeling of Geographical Relations for Map Generalization. *Proceedings of the 9th ICA Workshop on Generalisation and Multiple Representation,* Portland, USA, June 25th, 2006.

Egenhofer, M. J. and Herring, J. R. (1991): Categorizing Binary Topological Relations Between Regions, Lines, and Points in Geographic Databases. Technical Report, Department of Surveying Engineering, University of Maine, 28 pages.

Fonseca, F. T. et al. (2002): Using Ontologies for Integrated Geographic Information Systems. In: *Transactions in GIS*, **6** (3), 231–257.

Frank, A. U. (1996): Qualitative spatial reasoning: cardinal directions as an example. In: *International Journal of Geographical Information Systems,* **10** (3), 269–290.

Friedrich, K., Hahn B. and Popp, H. (2002): *Nationalatlas Bundesrepublik Deutschland*. Bd. 5. Dörfer und Städte. Spektrum Akademischer Verlag, Berlin, 150 pages.

Grenon, P. and Smith, B. (2004): SNAP and SPAN: Towards Dynamic Spatial Ontology. In: *Spatial Cognition and Computation*, **4** (1), 69–104.

Gruber, T. R. (1993): A translation approach to portable ontology specifications. In: *Knowledge Acquisition*, **5** (2), 199–220.

Guarino, N. (1998): Formal Ontology and Information Systems. In: Guarino, N. (Ed.) (1998): *Formal Ontology in Information Systems. Proceedings of FOIS'98*, Trento, Italy, June 6–8, 1998, 3–15.

Haarslev, V., Möller, R. and Schröder, C. (1994): Combining Spatial and Terminological Reasoning. In: *KI-94: Advances in Artificial Intelligence: 18th German Annual Conference on Artificial Intelligence*, Saarbrücken, Germany, September 8–23, 1994, Lecture Notes in Artificial Intelligence 861, 142–153.

Heinzle, F. et al. (2005): Graph Based Approaches for Recognition of Patterns and Implicit Information in Road Networks. *Proceedings of the 22nd International Cartographic Conference*, A Coruña, Spain, July 11–16, 2005.

Klien, E. and Lutz, M. (2005): The Role of Spatial Relations in Automating the Semantic Annotation of Geodata. In: Cohn, A. G. and Mark, D. M. (Eds.) (2005): *Spatial Information Theory, International Conference, COSIT 2005*, Ellicottville, NY, USA, September 14–18, 2005, LNCS 3693, 133–148.

Kovacs, K. and Zhou, S. (2007): Key challenges in expressing and utilising geospatial semantics at Ordnance Survey. Presentation held at the *European Geoinformatics Workshop*, Edinburgh, UK, March 7–9, 2007. http://www.nesc.ac.uk/action/esi/download.cfm?index=3411. Accessed 22.03.2007.

Kulik, L. et al. (2005): Ontology-Driven Map Generalization. In: *Journal of Visual Languages and Computing*, **16** (3), 245–267.

Mackaness, W. (2002): The Importance of Modelling Pattern and Structure in Automated Map Generalisation. *Proceedings of the Joint ISPRS/ICA Workshop on Multi-Scale Representations of Spatial Data*, Ottawa, Canada, July 7–8, 2002.

Mackaness, W. (2006): Automated Cartography in a Bush of Ghosts. In: *Cartography and Geographic Information Science,* **33** (4), 245–256.

Marshall, S. (2005): *Streets & patterns.* Spon Press, London and New York, 150 pages.

Neumann, B. and Möller, R. (2004): *On Scene Interpretation with Description Logics*. Technical Report, Universität Hamburg, FBI-B-257/04, 30 pages.

OWL (2007): OWL Web Ontology Language Reference. http://www.w3.org/TR/owl-ref/. Accessed 22.03.2007.

Peachavanish, R. and Karimi, H. A. (2007): Ontological Engineering for Interpreting Geospatial Queries. In: *Transactions in GIS*, **11** (1), 115–130.

Protégé (2007): Protégé website. http://protege.stanford.edu/. Accessed 22.03.2007.

Redbrake, D. and Raubal, M. (2004): Ontology-Driven Wrappers for Navigation Services. *Proceedings of the 7th AGILE Conference on GIScience*, Heraklion, Crete/Greece, April 29 - May 1, 2004.

Regnauld, N. (1996): Recognition of Building Clusters for Generalization. In: *Advances in GIS Research II: Proceedings 7th International Symposium on Spatial Data Handling*, Delft, Netherlands, August 12–16, 1996, 4B.1–4B.14.

Rodríguez, M. A. and Egenhofer, M. J. (2004): Comparing geospatial entity classes: an asymmetric and context-dependent similarity measure. In: *International Journal of Geographic Information Science*, **18** (3), 229–256.

Ruas, A., and Plazanet, C. (1996): Strategies for automated generalization. In: *Advances in GIS Research II: Proceedings 7th International Symposium on Spatial Data Handling*, Delft, Netherlands, August 12–16, 1996, 6.1–6.17.

Santos, P. et al. (2005): Supervaluation Semantics for an Inland Water Feature Ontology. *Proceedings of the Nineteenth International Joint Conference on Artificial Intelligence*, Edinburgh, Scotland, July 30 - August 5, 2005.

Sester, M. (2000): Knowledge acquisition for the automatic interpretation of spatial data. In: *International Journal of Geographical Information Science*, **14** (1), 1–24.

Steiniger, S. (2006a): Classifying urban structures for mapping purposes using discriminant analysis. In: Priestnall, G. and Aplin, P. (Eds.) (2006): *Proceedings of GISRUK Conference 2006*, Nottingham, UK, April 5-7, 2006, 107–111.

Steiniger, S. et al. (2006b): Recognition of Island Structures for Map Generalization. In: *Proceedings of ACM-GIS'06*, Arlington, Virginia, November 10-11, 2006, 67–74.

Thomson, T. (2006): *Cartographic Data Analysis and Enhancement. Objective 1: Analysis of Functional Information – Needs & Opportunity Analysis, User Requirements, and Use Case.* Technical Report, Department of Geomatic Engineering, University College London.

Torres, M. et al. (2005): Ontology-Driven Description of Spatial Data for Their Semantic Processing. In: Rodríguez, M. A. et al. (Eds.) (2005): *GeoSpatial Semantics: First International Conference, GeoS 2005*, Mexico City, Mexico, November 29–30, LNCS 3799, 242–249.

Vivid Solutions (2007): JUMP/JTS website. http://www.vividsolutions.com/. Accessed 22.03.2007.

Winter, S. (2002): Ontologisches Modellieren von Routen für mobile Navigationsdienste. In: Kelnhofer, F. and Lechthaler, M. (Eds.) (2002): *TeleKartographie und Location-Based Services*. Schriftenreihe der Studienrichtung Vermessungswesen und Geoinformation, Technische Universität Wien, 111–124.