Visual Spatial Query Languages: A Semantics Using Description Logic

Volker Haarslev Ralf Möller Michael Wessel
University of Hamburg, Computer Science Department,
Vogt-Kölln-Str. 30, 22527 Hamburg, Germany
`http://kogs-www.informatik.uni-hamburg.de/~<name>/`

Abstract

We present a first treatment dealing with the semantics of visual *spatial* query languages for geographic information systems using a suitable description logic. This decidable space logic is described and its usefulness for geographic information systems is exemplified. The logic supports the specification of a semantics, reasoning about query subsumption and about applying default knowledge.

1   Introduction

For accessing spatial databases or geographic information systems (GIS), different query specification techniques have been proposed. For instance, the visual spatial query system VISCO developed in our group [9, 13] can be used to query a spatial database (GIS) in a *visual* way. In contrast to conventional textual query systems the user is not required to learn a complicated textual query language in order to effectively use an information system. Users can query the database by drawing diagrammatic representations of what is to be retrieved from the spatial information system. However, experiences with the current VISCO system indicate that in the context of VISCO (and query systems in general), the specification of queries in a GIS still could be made easier by advances in research areas combining spatial and terminological reasoning with visual language theory.

In this paper we discuss the application of a new logic-based formalism to specifying the semantics of visual spatial queries. To the best of our knowledge this is the first proposal utilizing an expressive and decidable spatial logic for this task. The formalism can be used to define the semantics of

visual spatial queries, to reason about query subsumption, and to deal with multiple worlds or query completion with the help of default reasoning. Examples for these kinds of reasoning are discussed in this paper. Our formalism is based on the description logic $\mathcal{ALCRP(D)}$ [6, 7] offering mechanisms for integrating so-called *concrete domains* and on a recent extension for default reasoning [11].

We like to emphasize that the work on VL theory presented in this paper truly extends our previous research as summarized in [3], where we used a logic that is more expressive than $\mathcal{ALCRP(D)}$ since it allows *qualified number restrictions* but also less expressive than $\mathcal{ALCRP(D)}$ since it has no *defined roles*. In [4] we made the first proposal for using $\mathcal{ALCRP(D)}$ for reasoning about visual representations. This proposal was extended in [5] by considering the semantics of visual spatial query languages. In [8] we integrated default reasoning. This paper summarizes our previous work and extends it by using so-called ABox patterns for describing *n*-ary queries.

2   The Description Logic $\mathcal{ALCRP(D)}$

This section gives a brief introduction to the description logic $\mathcal{ALCRP(D)}$ and to description logics (DLs) in general summarizing the notions important for this paper. We refer to [1, 14] for more information about description logics.

Many DLs can be viewed as subsets of first-order predicate logic. However it is important to note that particular DLs are only considered as practical if they are based on *sound, complete* and *terminating* reasoning algorithms, i.e. the *decidability* of a DL is of utmost importance. Of course, this is a major distinction to reasoning with general first-order predicate logic.

DLs are based on the ideas of structured inheritance networks [2]. In a DL a factual world consists of named individuals and their relationships that are asserted through binary relations. Hierarchical descriptions about sets of individuals form the terminological knowledge. Descriptions (or terms) about sets of individuals are called *concepts* and binary relations are called *roles*. Descriptions consist of identifiers denoting concepts, roles, and individuals, and of description constructors. For any individual x the set $\{y|r(x,y)\}$ is called the set of *fillers* of the role r. A role which may have at most one filler is referred to as a *feature*.

For instance, consider the following description used in our GIS scenario with the intended meaning "a cottage that is enclosed by a forest" that contains concept names (e.g. cottage), role names (e.g. is_g_inside), and constructors (e.g. $\sqcap$ and $\exists$).
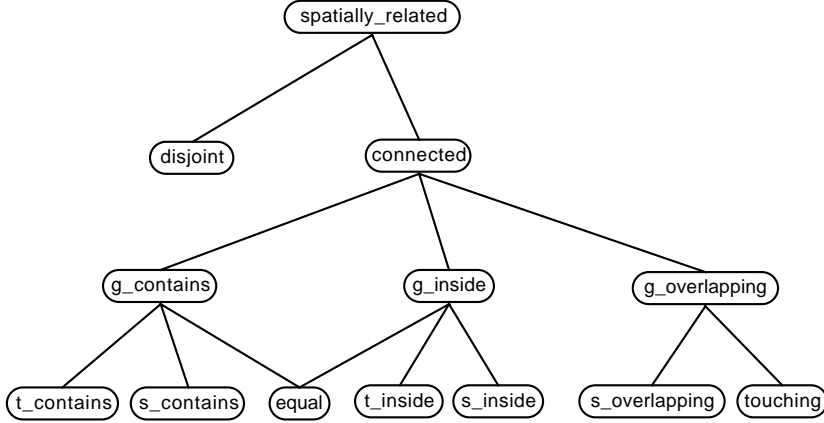
Figure 1: Subsumption hierarchy of spatial predicates.

| disjoint | touching | s_overlapping | t_contains | s_contains |

Figure 2: Spatial relations between A and B. The inverses of t_contains and s_contains as well as the relation equal have been omitted.

cottage_in_forest $\doteq$ cottage $\sqcap$ $\exists$ is_g_inside . forest

2.1 Terminologies

In this section, the language (syntax and semantics) for defining concepts and roles in $\mathcal{ALCRP}(\mathcal{D})$ is presented. $\mathcal{ALCRP}(\mathcal{D})$ is parameterized with a concrete domain which consists of a set of concrete objects and a set of predicates.

Concrete Domains   A *concrete* domain $\mathcal{D}$ is a pair $(\Delta_{\mathcal{D}}, \Phi_{\mathcal{D}})$, where $\Delta_{\mathcal{D}}$ is a set called the domain, and $\Phi_{\mathcal{D}}$ is a set of predicate names. Each predicate name P from $\Phi_{\mathcal{D}}$ is associated with an arity n, and an n-ary predicate $\mathsf{P}^{\mathcal{D}} \subseteq \Delta_{\mathcal{D}}^n$. A concrete domain $\mathcal{D}$ is called *admissible* iff (1) the set of its predicate names is closed under negation (i.e. for any $\mathsf{P} \in \Phi_{\mathcal{D}}$ there exists a $\overline{\mathsf{P}} \in \Phi_{\mathcal{D}}$ denoting the negation of P) and contains a name $\top_{\mathcal{D}}$ for $\Delta_{\mathcal{D}}$ and (2) the satisfiability problem for finite conjunctions of predicates is decidable.

A concrete domain can be understood as a device providing a bridge between conceptual reasoning with abstract entities and (qualitative) constraint reasoning with concrete or symbolic data. In this paper we use the admissible

concrete domain $\mathcal{RS}_2$. It is the union of the domains $\mathcal{R}$ (over the set $\mathbb{R}$ of all real numbers with predicates built by first order means from (in)equalities between integer polynomials in several indeterminates, see [12]) and $\mathcal{S}_2$ (over the set of all two-dimensional polygons with topological relations from Figure 1 and 2 as predicates, see [7]). The name 'concrete domain' is in some sense misleading since it suggests that a concrete domain realizes reasoning about 'concrete' (e.g. numeric) data. This kind of reasoning is sometimes supported (e.g. in the domain $\mathcal{R}$) but in our application we mainly use concrete domains for reasoning about the satisfiability of finite conjunctions of predicates. For instance, the domain $\mathcal{S}_2$ qualitatively decides the satisfiability of conjunctions such as $\mathsf{touching}(I_1, I_2) \wedge \mathsf{contains}(I_2, I_3) \wedge \mathsf{touching}(I_1, I_3)$ without any notion for 'concrete' polygons. This is a well-known example for a constraint satisfaction problem.

Without loss of generality we introduce a λ-like notation for anonymous predicates for the domain $\mathcal{R}$. Formally, each anonymous predicate and its negation could be replaced by unique names for the λ-term and its negated counterpart and, moreover, the negation sign in front of a λ-term can be safely moved inside this term.

Role Terms Let R and F be disjoint sets of role and feature names, respectively. Any element of R and any element of F is an *atomic role term*. The elements of F are also called *features*. A composition of features (written $\mathsf{f_1 f_2 \cdots}$) is called a feature chain. A feature chain of length one is also a feature chain. If $\mathsf{P} \in \Phi_{\mathcal{D}}$ is a predicate name with arity $n + m$ and $\mathsf{u_1}$, ..., $\mathsf{u_n}$ as well as $\mathsf{v_1}$, ..., $\mathsf{v_m}$ are $n + m$ feature chains, then the expression $\exists (\mathsf{u_1}, \ldots, \mathsf{u_n})(\mathsf{v_1}, \ldots, \mathsf{v_m}) . \mathsf{P}$ (*role-forming predicate restriction*) is a *complex role term*. Let S be a role name and let T be a role term. Then $\mathsf{S} \doteq \mathsf{T}$ is a terminological axiom.

An example for using a role-forming predicate operator is the definition of a role $\mathsf{is\_g\_inside}$ for a corresponding topological predicate $\mathsf{g\_inside}$ (see Section 2.1 for an explanation of the semantics). Intuitively speaking, this role holds for any pair of individuals (I_1, I_2) iff the spatial area (associated via the feature $\mathsf{has\_area}$) of I_1 is g_inside of the area of I_2 (see also Figure 3 where this is illustrated for the predicate $\mathsf{s\_inside}$).

is_g_inside $\doteq$ $\exists (\mathsf{has\_area})(\mathsf{has\_area}) . \mathsf{g\_inside}$

Concept Terms Let C be a set of concept names which is disjoint from R and F. Any element of C is an *atomic concept term*. If C and D are concept terms, R is an arbitrary role term or a feature, $\mathsf{P} \in \Phi_{\mathcal{D}}$ is a predicate name

with arity n, and $\mathsf{u}_1, \ldots, \mathsf{u}_n$ are feature chains, then the following expressions are also concept terms: $\mathsf{C} \sqcap \mathsf{D}$ (*conjunction*), $\mathsf{C} \sqcup \mathsf{D}$ (*disjunction*), $\neg\mathsf{C}$ (*negation*), $\exists\,\mathsf{R}\,.\,\mathsf{C}$ (*concept exists restriction*), $\forall\,\mathsf{R}\,.\,\mathsf{C}$ (*concept value restriction*), and $\exists\,\mathsf{u}_1, \ldots, \mathsf{u}_n\,.\,P$ (*predicate exists restriction*). Concept terms may also be written in parentheses.

We illustrate the notion of concept and role terms by extending the cottage example mentioned above.

$$\mathbf{normal\_cottage\_in\_forest} \doteq \mathsf{cottage\_in\_forest} \sqcap$$
$$\exists\,\mathsf{has\_space}\,.\,\lambda_{\mathcal{R}}x\,.\,(x \geq 30 \wedge x < 70)$$

The definition of $\mathsf{normal\_cottage\_in\_forest}$ roughly has the intended meaning "something is a standard cottage in a forest if and only if it is a cottage located in a forest with 30-70 square meters of total space for the cottage." This definition also gives an example for a predicate exists restriction for the domain $\mathcal{R}$ using a feature $\mathsf{has\_space}$.

In order to ensure the decidability, we had to restrict possible combinations of concepts terms w.r.t. defined roles (e.g. a nested concept term with defined roles such as $\forall\,\mathsf{is\_touching}\,.\,\exists\,\mathsf{is\_g\_inside}\,.\,\mathsf{cottage}$ is not allowed). Note that all examples in this paper are restricted. However, this restrictedness criterion is beyond the scope of this paper and is explained elsewhere [7].

Terminology Let A be a concept name and D be a concept term. Then $\mathsf{A} \doteq \mathsf{D}$ and $\mathsf{A} \sqsubseteq \mathsf{D}$ are terminological axioms as well. A finite set of terminological axioms $\mathcal{T}$ is called a *terminology* or *TBox* if no concept or role name in $\mathcal{T}$ appears more than once on the left-hand side of a definition and, furthermore, if no cyclic definitions are present.

The previous examples already informally introduced concept axioms for defined concepts using the $\doteq$ operator. For convenience, we also allow the $\sqsubseteq$ operator for the definition of primitive concepts, i.e. their definition consists only of necessary conditions. The concept $\mathsf{cottage}$ is a good candidate for a primitive definition documenting that we omitted in our terminology other conditions that are not relevant for this modeling task. For instance, a cottage has to be at least a building: $\mathbf{cottage} \sqsubseteq \mathsf{building}$.

Of course, there exist other description logics that allow more than one axiom for a particular concept name or even support generalized concept inclusions (implications) with arbitrary concept terms on the left- and right-hand side of terminological axioms. These axioms can be used as a powerful modeling tool but are currently not supported in $\mathcal{ALCRP(D)}$ w.r.t. decidability.

Semantics An *interpretation* $\mathcal{I} = (\Delta_{\mathcal{I}}, \cdot^{\mathcal{I}})$ consists of a set $\Delta_{\mathcal{I}}$ (the abstract domain) and an interpretation function $\cdot^{\mathcal{I}}$. The sets $\Delta_{\mathcal{D}}$ (see above) and $\Delta_{\mathcal{I}}$ must be disjoint. The interpretation function maps each concept name C to a subset $\mathsf{C}^{\mathcal{I}}$ of $\Delta_{\mathcal{I}}$, each role name R to a subset $\mathsf{R}^{\mathcal{I}}$ of $\Delta_{\mathcal{I}} \times \Delta_{\mathcal{I}}$, and each feature name f to a partial function $\mathsf{f}^{\mathcal{I}}$ from $\Delta_{\mathcal{I}}$ to $\Delta_{\mathcal{D}} \cup \Delta_{\mathcal{I}}$, where $\mathsf{f}^{\mathcal{I}}(a) = x$ will be written as $(a, x) \in \mathsf{f}^{\mathcal{I}}$. If $u = \mathsf{f}_1 \cdots \mathsf{f}_n$ is a feature chain, then $u^{\mathcal{I}}$ denotes the composition $\mathsf{f}_1^{\mathcal{I}} \circ \cdots \circ \mathsf{f}_n^{\mathcal{I}}$ of the partial functions $\mathsf{f}_1^{\mathcal{I}}, \ldots, \mathsf{f}_n^{\mathcal{I}}$. Let the symbols C, D, R, P, $\mathsf{u}_1, \ldots, \mathsf{u}_m$, and $\mathsf{v}_1, \ldots, \mathsf{v}_m$ be defined as above. Then the interpretation function can be extended to arbitrary concept and role terms as follows:

$$(\mathsf{C} \sqcap \mathsf{D})^{\mathcal{I}} := \mathsf{C}^{\mathcal{I}} \cap \mathsf{D}^{\mathcal{I}}, \quad (\mathsf{C} \sqcup \mathsf{D})^{\mathcal{I}} := \mathsf{C}^{\mathcal{I}} \cup \mathsf{D}^{\mathcal{I}}, \quad (\neg \mathsf{C})^{\mathcal{I}} := \Delta_{\mathcal{I}} \setminus \mathsf{C}^{\mathcal{I}}$$

$$(\exists \mathsf{R} . \mathsf{C})^{\mathcal{I}} := \{a \in \Delta_{\mathcal{I}} \mid \exists b \in \Delta_{\mathcal{I}} : (a, b) \in \mathsf{R}^{\mathcal{I}}, b \in \mathsf{C}^{\mathcal{I}}\}$$

$$(\forall \mathsf{R} . \mathsf{C})^{\mathcal{I}} := \{a \in \Delta_{\mathcal{I}} \mid \forall b : (a, b) \in \mathsf{R}^{\mathcal{I}} \Rightarrow b \in \mathsf{C}^{\mathcal{I}}\}$$

$$(\exists \mathsf{u}_1, \ldots, \mathsf{u}_n . \mathsf{P})^{\mathcal{I}} :=$$

$$\{a \in \Delta_{\mathcal{I}} \mid \exists x_1, \ldots, x_n \in \Delta_{\mathcal{D}} : (a, x_1) \in \mathsf{u}_1^{\mathcal{I}}, \ldots, (a, x_n) \in \mathsf{u}_n^{\mathcal{I}}, (x_1, \ldots, x_n) \in \mathsf{P}^{\mathcal{D}}\}$$

$$(\exists (\mathsf{u}_1, \ldots, \mathsf{u}_n)(\mathsf{v}_1, \ldots, \mathsf{v}_m) . \mathsf{P})^{\mathcal{I}} :=$$

$$\{(a, b) \in \Delta_{\mathcal{I}} \times \Delta_{\mathcal{I}} \mid \exists x_1, \ldots, x_n, y_1, \ldots, y_m \in \Delta_{\mathcal{D}} : (a, x_1) \in \mathsf{u}_1^{\mathcal{I}}, \ldots, (a, x_n) \in \mathsf{u}_n^{\mathcal{I}},$$
$$(b, y_1) \in \mathsf{v}_1^{\mathcal{I}}, \ldots, (b, y_m) \in \mathsf{v}_m^{\mathcal{I}}, (x_1, \ldots, x_n, y_1, \ldots, y_m) \in \mathsf{P}^{\mathcal{D}}\}$$

An interpretation $\mathcal{I}$ is a *model* of a TBox $\mathcal{T}$ iff it satisfies $\mathsf{A}^{\mathcal{I}} = \mathsf{C}^{\mathcal{I}}$ for all terminological axioms $\mathsf{A} \doteq \mathsf{C}$ in $\mathcal{T}$, and $\mathsf{A}^{\mathcal{I}} \subseteq \mathsf{C}^{\mathcal{I}}$ for $\mathsf{A} \sqsubseteq \mathsf{C}$ respectively.

Figure 3 illustrates the idea behind the semantics of the role-forming predicate operator for the domain $\mathcal{S}_2$. The spatial predicates (e.g. g_inside) operate on concrete domain values (e.g. polygons) that are attached to corresponding abstract individuals via features. If a role (e.g. is_s_inside) is *defined* by a predicate (e.g. s_inside), then every pair $(\mathsf{p}_1, \mathsf{p}_2)$ of polygons that are fillers of has_area for two abstract individuals i_1 and i_2 is tested whether the binary predicate s_inside$(\mathsf{p}_2, \mathsf{p}_1)$ is fulfilled. In case of a successful test the role membership (e.g. is_s_inside) is also established for the abstract individuals i_1 and i_2, i.e. it holds is_s_inside$(\mathsf{i}_2, \mathsf{i}_1)$. This also applies for the opposite direction: if a role membership is asserted for a pair of abstract individuals, their associated concrete feature fillers are either established with the corresponding predicate or verified if concrete feature fillers already exist.

2.2   The Assertional Language

The *assertional language* of a DL is designed for stating concept or role memberships of named individuals that are used to describe the factual world.
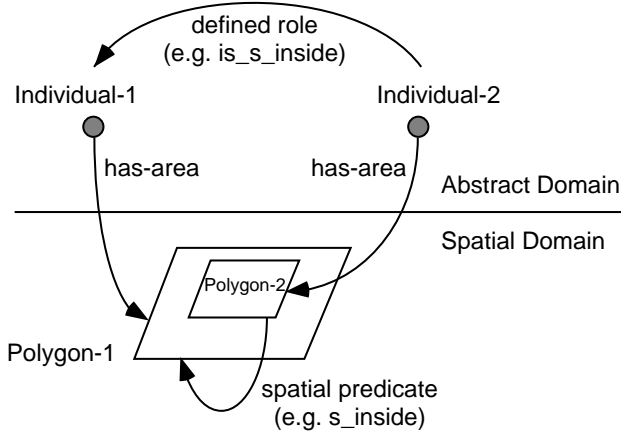
Figure 3: Relationship between abstract and spatial domain.

With respect to concrete domains we distinguish *abstract* and *concrete* individuals. Abstract individuals are elements of the abstract domain. Concrete individuals are elements of the concrete domain and are used as parameters for predicates. Both types of individuals can be feature fillers while only abstract individuals can be role fillers. For instance, in the VL domain abstract individuals can be used to represent geometric figures such as circles, rectangles, etc. that, in turn, represent *syntactic* elements of a query language. Additionally, abstract individuals can also be used to represent *semantic* elements (e.g. lake, estate, etc.) of a GIS. The descriptions might be associated via features with reals defining geometric properties such as the diameter of a circle (lake), the width and height of a rectangle (estate), etc.

The set of assertions (ABox) has to comply to the definitions declared in the TBox. An *ABox* of $\mathcal{ALCRP}(\mathcal{D})$ is a finite set of assertions defined as follows.

Syntax Let I_A and I_D be two disjoint sets of individual names for the abstract and concrete domain. If C is a concept term, R an atomic or complex role term, f a feature name, P a predicate name with arity n, a and b are elements of I_A and x, x_1, ... ,x_n are elements of I_D, then the following expressions are *assertional axioms*: $\mathsf{a} : \mathsf{C}$ (*concept membership*), $(\mathsf{a}, \mathsf{b}) : \mathsf{R}$ (*role filler*), $(\mathsf{a}, \mathsf{x}) : \mathsf{f}$ (*feature filler*), $(\mathsf{x}_1, \ldots, \mathsf{x}_n) : \mathsf{P}$ (*predicate membership*).

Semantics For specifying the semantics of ABox assertions we have to extend the interpretation function $\mathcal{I}$. An *interpretation* for the assertional language is an interpretation for the concept language which additionally maps every individual name from I_A to a single element of $\Delta_{\mathcal{I}}$ and every

individual name from I_D to a single element from $\Delta_{\mathcal{D}}$. We assume that the unique name assumption does not hold, that is $a^{\mathcal{I}} = b^{\mathcal{I}}$ may hold even if $a \neq b$.

$$a : C \text{ iff } a^{\mathcal{I}} \in C^{\mathcal{I}}, \ (a, b) : R \text{ iff } (a^{\mathcal{I}}, b^{\mathcal{I}}) \in R^{\mathcal{I}}$$
$$(a, x) : f \text{ iff } f^{\mathcal{I}}(a^{\mathcal{I}}) = x^{\mathcal{I}}, \ (x_1, \ldots, x_n) : P \text{ iff } x_1^{\mathcal{I}}, \ldots, x_n^{\mathcal{I}} \in P^{\mathcal{D}}$$

ABox Example Using the cottage scenario the following ABox $\mathcal{A}$ illustrates the four different types of ABox assertions. Based on the semantics given above a $\mathcal{ALCRP}(\mathcal{D})$ reasoner will infer that c is a member of normal_cottage_in_forest. S_c and S_f denote the associated area polygon of c and f.

$\mathcal{A} = \{c : \text{cottage}, (c, 60) : \text{has\_space}, (c, S_c) : \text{has\_area}, f : \text{forest},$
$\qquad (f, S_f) : \text{has\_area}, (S_c, S_f) : \text{g\_inside}\}.$

2.3   Reasoning Services and Complexity

The notion of a *model* (see above) is used to define the reasoning services that a DL inference engine has to provide, i.e. it proves for every concept specification that the following conditions hold:

- a term A *subsumes* a term B if and only if for every model $\mathcal{I}$, $B^{\mathcal{I}} \subseteq A^{\mathcal{I}}$;

- a term A is *coherent/satisfiable* if and only if there exists at least one model $\mathcal{I}$ such that $A^{\mathcal{I}} \neq \emptyset$;

- terms A and B are *disjoint* if and only if for every model $\mathcal{I}$, $A^{\mathcal{I}} \cap B^{\mathcal{I}} = \emptyset$;

- terms A and B are *equivalent* if and only if for every model $\mathcal{I}$, $A^{\mathcal{I}} = B^{\mathcal{I}}$.

- An *ABox* $\mathcal{A}$ is *consistent* if and only if there exists a model $\mathcal{I}$ of $\mathcal{A}$.

- An individual a in an ABox $\mathcal{A}$ is a *member* of a concept C if and only if $\mathcal{A} \cup \{a : \neg C\}$ is not consistent.

The process of computing the direct subsumption relationships between all concept names in a TBox $\mathcal{T}$ is called *classification* and creates a concept taxonomy. The process of computing the direct concept membership for all individuals in an ABox $\mathcal{A}$ is called *realization*.

The expressiveness and computational complexity of a particular DL depends on the variety of employed description constructors. Various complexity results for subsumption algorithms for specific description logics are summarized in [14]. The complexity for deciding satisfiability of $\mathcal{ALCRP}(\mathcal{D})$ concepts is NExpTime-complete [10].

The incoherence of a concept is illustrated by the following situation. We define a paradise cottage as a fishing cottage located in a mosquito-free forest, i.e. the forest is not spatially connected with a river.

fishing_cottage $\doteq$ cottage $\sqcap$ $\exists$ is_touching . river

mosquito_free_forest $\doteq$ forest $\sqcap$ $\forall$ is_connected . $\neg$river

paradise_cottage $\doteq$ fishing_cottage $\sqcap$ $\exists$ is_g_inside . forest $\sqcap$
$\forall$ is_g_inside . mosquito_free_forest

However, a fishing cottage is defined as a cottage that touches a river. It follows that the forest containing a fishing cottage must also be spatially connected with this river. Obviously, the paradise cottage is only a dream that can not exist in the real world. This is due to the intended semantics of the underlying spatial relations: A situation where a region r_1 (cottage) is *g_inside* another region r_2 (forest) and this region r_1 is also *touching* a third region r_3 (river) implies that r_2 is *connected* to r_3, i.e. g_inside$(r_1, r_2) \wedge$ touching$(r_1, r_3) \Rightarrow$ connected(r_2, r_3). This implies that a paradise cottage cannot be located inside a mosquito-free cottage.

3   Semantics of Spatial Queries

The previous sections defined the description logic $\mathcal{ALCRP}(\mathcal{D})$ and demonstrated its usefulness for spatial reasoning. We introduced semantic entities such as buildings, cottages, forests, etc. These entities are suitable candidates for elements of visual spatial queries. In VISCO we assume that these and other basic map objects are predefined in a GIS. Furthermore, spatial areas are defined by polygons. Map elements (e.g. polylines, polygons) are annotated with labels such as "forest", "building", "river" etc. that directly correspond to the semantic entities characterized above.

We imagine a VISCO application scenario for querying a GIS as follows. Instead of textually writing a complicated SQL query, a user simply draws a constellation of spatial entities that resemble the intended constellation of interest. Using the basic vocabulary provided by the GIS the user has to annotate drawing elements by concept names (e.g. this polygon represents

(a) Incomplete query   (b) Completion 1   (c) Completion 2   (d) Inconsistency

Figure 4: Automatic completion of visual queries by application of default rules.

a forest). The parser of VISCO would analyze the drawing and create a corresponding ABox as semantic representation. Thus, the *semantics* of a query is defined by an ABox derived from a spatial constellation.

Sometimes it might be hard for users to fully specify a query. Therefore, a completion facility is needed to resolve semantic ambiguities or to complete underspecified information by using default rules for further specialization. The next subsections describe the usefulness of spatial default reasoning and the query processing and reasoning process.

3.1 Completion of Queries

Default knowledge is used to make queries more precise if it can be applied in a consistent way. Due to space limitations we omit any discussions about the formal representation of default knowledge and its rules of inference. This is discussed elsewhere [11]. First of all, the process of formulating (visual) queries can be facilitated by *automatically completing* queries in a meaningful way, therefore reducing the number of mouse interactions. The process of selecting semantic concept descriptors for objects involved in a query (e.g. cottage, river, forest) can partly be automated by interpreting a partially specified query. For instance, in its current development stage, VISCO users can select concept descriptors from a list of over 300 predefined concepts. Thus, even a situation-adapted reduction of the complete list of possibilities to a suitable subset or an order relation for sorting groups of possible concept candidates would be very appropriate.

In order to analyze the modeling problems in this context, we begin with a more detailed discussion of a visual query example. Let us assume a person is interested in buying a cottage located in a forest. In Figure 4(a) the user just started to formulate the query. After (s)he has specified that the type of the surrounding polygon A should be a forest, the type of the small polygon B must be specified. A smart interface should use formal derivation processes for computing plausible candidates for object "type" specifications.

10

For narrowing the set of possibilities we assume that two default rules are applicable: one is saying that the interior small polygon B could be a cottage (Figure 4(b)) and another is stating that B could be a lake (Figure 4(c)) if this does not lead to inconsistencies. Since an object can be either a lake or a cottage, there is no way to believe in both possibilities at a time. This kind of default rule interaction is a simple example demonstrating the necessity of considering different *possible worlds* which must be maintained by the reasoning system. Depending on the default rule being used to conclude new knowledge, different subsequent conclusions might be possible.

Other potentially active default rules might produce inconsistencies with the set of current assertions without providing a possibility of using multiple worlds to avoid inconsistencies. For example, if a default rule is applied that the small polygon B might as well be a forest (Figure 4(d)), we would get a contradiction if an axiom (as part of our conceptual background knowledge) states that a forest can never contain another forest. Thus, in our query context, the latter default cannot be applied and, as a consequence of computing and appropriately interpreting the set of possible worlds, we can compose a situation-adapted menu for the graphical user interface and the user can select between meaningful concepts for object B. In our specific example, the menu will contain items for cottage and lake but not for forest.

If more than one possible world is computed, an intuitive criterion would be to select the world originating from a default with the more specific pre-condition or conclusion. E.g., in the query shown in Figure 5(a) we would prefer a default concluding that the thin graphical object might be a 'river flowing into a lake' (which might be a useful concept in our scenario) instead of a more general default concluding only that the object is an ordinary river.

The automatic augmentation of visual queries by conclusions of applied default rules can be seen as a specialization process. Therefore, this process might not only be useful during the construction of a visual query, but also useful as a tool for query refinement after a query has been executed that yields too many results. In addition, not only conceptual information is important. In our GIS context we also have to consider the spatial relations between domain objects

In the context of sketch-based visual querying, on the one hand it is sometimes useful to leave some spatial relations between graphical objects unspecified because they are unknown or simply because the user is not willing to specify them. On the other hand, in order to actually draw a picture, the user *must* specify each spatial relation, even if it is just one of several possible (base) relations. The problem of how to specify "don't care relations" or "example relations" is well known and inherent in diagrammatic representations. It is similar to the problem of visualizing visual disjunctions.

Figure 5: Scenarios for situation-adapted completion of queries (see text).

For example, in the query shown in Figure 5b, we have a visible disjoint relation between the river and the lake. If we intended the river to be disjoint from the lake, the query answering system would not find any rivers flowing into this lake. The problem is how can we specify that the river should be strictly inside the forest but leave the relation to the lake unspecified. As a possible solution to this problem, we could simply *ignore* each visible disjoint relation. But, with this interpretation, we can now no longer state a query searching for rivers *not* flowing into this specific lake, which might be a very useful concept. We propose the following solution. For objects like the river that are drawn with a specific drawing attribute such as dashing, the universal spatial relation to other objects (disjunction of all base relations) is asserted. Dashed objects introduce no spatial query constraints. However, in some cases this would usually not match the users intention as there will be too many matches, i.e. the answer set will be too large. With the help of default knowledge we can automatically refine the query in a way that is appropriate according to the semantics of the objects involved in a query. So, we can guide the interpretation of spatial aspects by the help of conceptual background knowledge and application of defaults, yielding different hypotheses as possible worlds. A river flows into a lake or not, i.e. graphically both objects are either touching, see also Figure 5c) or they are disjoint (see Figure 5d). With respect to a lake, there are no other possibilities. In our world model a river never overlaps with a lake (see also Figure 5e). This is assumed to be stated as an axiom as part of our general conceptual background knowledge. Besides defaults involving concept constraints we also have to take care of default rules with conclusions yielding new *relation* constraints. For instance, one rule could conclude the relationship touching, the other disjoint (see [11] for a discussion).

3.2   Reasoning about Visual Spatial Queries

We flesh out the scenario for the GIS query introduced above. The cottage should be located in a forest with a river in the immediate vicinity. The buyer

Figure 6: Spatial sketches representing spatial queries.

also want a cottage that provides about 60 m^2 floor space. The estate itself should have about 400 m^2. Having these requirements in mind we sketch a query (see Figure 6a) reflecting the intended spatial and geometric constraints. The parser translates the sketch to an equivalent ABox on the basis of a taxonomy containing concept descriptions for the spatial vocabulary of this GIS domain. We get the following Abox:

$\mathcal{A}_0 = \{$c : cottage $\sqcap \exists$ has_space . $\lambda_{\mathcal{R}} x$. $(x > 40 \wedge x < 70)$,
e : estate_area $\sqcap \exists$ has_space . $\lambda_{\mathcal{R}} x$. $(x > 350 \wedge x < 450)$,
r : river, (r, e) : is_touching, (c, e) : is_g_inside, f : forest, (e, f) : is_g_inside$\}$

We use concept and role expressions as defined in the previous examples. The cottage is described by the individual c with a predicate-exists restriction asserting a floor space between 40 and 70 m^2. The cottage c has to be inside of an estate with a size between 350 and 450 m^2. As a simplification we assume that the river r has to touch the estate e that is inside of a forest f. Additionally, we assume the following new or revised concept definitions.

estate $\sqsubseteq$ spatial_area $\sqcap \exists$ has_space . $\lambda_{\mathcal{R}} x$. $(\top_{\mathcal{R}}(x))$

estate_in_forest $\doteq$ estate $\sqcap \exists$ is_g_inside . forest

fishing_cottage $\doteq$ cottage $\sqcap \exists$ is_g_inside . (estate $\sqcap \exists$ is_touching . river)

The realizing component of the $\mathcal{ALCRP}(\mathcal{D})$ reasoner will compute the following direct concept memberships of the individual c: normal_cottage_in_forest and fishing_cottage. The individual e will be the direct concept member of estate_in_forest. The other individuals r and f keep their asserted concept memberships.

In the following we employ a so-called *abstraction process* that can reduce particular ABoxes to corresponding ABoxes consisting of a single concept assertion representing the original query. The concept used in the assertion represents an abstraction of the ABox and, in turn, of the original query. It defines the *semantics* of this query and is used as a *query concept*. We assume

13

that the GIS provides a query facility (implemented by the $\mathcal{ALCRP(D)}$) reasoner) that can retrieve all stored individuals that are members of a given query concept. Therefore, the concept mentioned in a reduced ABox is used to retrieve all "matching" individuals and to answer the query.

For instance, using the abstraction process we can replace ABox $\mathcal{A}_0$ by ABox $\mathcal{A}_1 = \{c : cottage_{c_1}\}$. $\mathcal{A}_1$ uses the synthesized concept $cottage_{c_1}$. The concept $estate_{e_1}$ is only introduced to improve the readability of $cottage_{c_1}$.

$$estate_{e_1} \doteq estate \sqcap \exists\, is\_g\_inside\,.\, forest \sqcap \exists\, is\_touching\,.\, river$$

$$cottage_{c_1} \doteq cottage \sqcap \exists\, is\_g\_inside\,.\, estate_{e_1} \sqcap$$
$$\exists\, has\_space\,.\, \lambda_{\mathcal{R}}\, x\,.\, (x > 40 \wedge x < 70)$$

The query concept $cottage_{c_1}$ is classified by the reasoner. The semantic validity of this query is automatically verified during classification, i.e. to check whether the query concept is coherent (see Section 2.3). For instance, if the forest f were required to be 'mosquito-free' (see above), the $\mathcal{ALCRP(D)}$ reasoner would immediately recognize the incoherence of $cottage_{c_1}$. This information could be used by the spatial parser for generating an explanation to the user and for identifying the source of the contradiction.

Let us assume that the query with the concept $cottage_{c_1}$ returns more than 100 matches (i.e. individuals). The next step for the user might be to refine the query by adding more constraints.[1] One could add more requirements to the estate, e.g. we ask for a garage connected to the cottage. The extended sketch (see Figure 6b) corresponds to the ABox $\mathcal{A}_2$ (ignoring the lake) by adding to ABox $\mathcal{A}_0$ new assertions: $\mathcal{A}_2 = \mathcal{A}_0 \cup \{g : garage, (c, g) : is\_touching\}$. The abstraction process reduces ABox $\mathcal{A}_2$ to ABox $\mathcal{A}_3 = \{c : cottage_{c_2}\}$ using the following synthesized concept description.

$$cottage_{c_2} \doteq cottage \sqcap \exists\, has\_space\,.\, \lambda_{\mathcal{R}}\, x\,.\, (x > 40 \wedge x < 70) \sqcap$$
$$\exists\, is\_g\_inside\,.\, estate_{e_1} \sqcap \exists\, is\_touching\,.\, garage$$

$$cottage_{c_3} \doteq cottage_{c_1} \sqcap \exists\, is\_touching\,.\, garage$$

The $\mathcal{ALCRP(D)}$ reasoner recognizes the taxonomic relationship that $cottage_{c_1}$ subsumes $cottage_{c_2}$. It can be rewritten as $cottage_{c_3}$ that even textually demonstrates the subsumption relationship.

---

[1]Of course, one of the most important criteria is the price of the estate. This is currently neglected due to the non-spatial nature of this part of the query (but see our proposal in Section 4).

Subsumption between query concepts immediately leads to subsumption between GIS queries and to its utilization for query optimization. In order to answer the refined query a query optimizer can benefit from the detected query subsumption and reduce the search space to the set of query matches already computed for the query concept from ABox $\mathcal{A}_1$. Note that these query matches are members of the concept $\mathsf{cottage}_{c_1}$. This type of query optimization is an important aspect in applying description logics to database theory (see [1] for an introduction to these topics).

The benefits of computing a concept subsumption taxonomy can be even more subtle. Imagine a query from another user looking for a cottage located in a forest that is connected to a river. The ABox $\mathcal{A}_4$ is derived from the sketch. The abstraction process creates the following two concept definitions from ABox $\mathcal{A}_4$.

$\mathcal{A}_4 = \{ \mathsf{c} : \mathsf{cottage}, \mathsf{e} : \mathsf{estate\_area}, (\mathsf{c}, \mathsf{e}) : \mathsf{is\_g\_inside}, \mathsf{r} : \mathsf{river}, \mathsf{f} : \mathsf{forest},$
$\qquad (\mathsf{f}, \mathsf{r}) : \mathsf{is\_connected}, (\mathsf{e}, \mathsf{f}) : \mathsf{is\_g\_inside} \}$

$\mathsf{estate}_{e_2} \doteq \mathsf{estate} \sqcap \exists \, \mathsf{is\_g\_inside} \, . \, (\mathsf{forest} \sqcap \exists \, \mathsf{is\_connected} \, . \, \mathsf{river})$

$\mathsf{cottage}_{c_4} \doteq \mathsf{cottage} \sqcap \exists \, \mathsf{is\_g\_inside} \, . \, \mathsf{estate}_{e_2}$

The resulting ABox $\mathcal{A}_5 = \{ \mathsf{c} : \mathsf{cottage}_{c_4} \}$ uses the derived concepts. It turns out that the concept $\mathsf{cottage}_{c_4}$ subsumes the other concepts $\mathsf{cottage}_{c_i}$ although the concept descriptions are textually different. This is a rather complex proof based on the interaction between the spatial relations: $\mathsf{g\_inside}(\mathsf{e}, \mathsf{f}) \wedge \mathsf{touching}(\mathsf{e}, \mathsf{r}) \Rightarrow \mathsf{connected}(\mathsf{f}, \mathsf{r})$.

The abstraction process works rather well for ABoxes containing no joins or cycles, i.e. the same individual is a filler of several roles or even related to itself through a cycle of role assertions. If joins or cycles are present in an ABox, it depends on the expressiveness of the description logic whether an ABox can be reduced to a single concept membership assertion. For instance, joins can be expressed by restricting the number of possible role fillers or by equality restrictions for feature fillers. As mentioned above, other DLs also support the definition of cyclic concepts that might be required to fully reduce some ABoxes. Due to unknown decidability results $\mathcal{ALCRP}(\mathcal{D})$ currently does not allow cyclic concepts or number restrictions. Therefore, in case of ABoxes with joins or cycles, we can only partially reduce these ABoxes. This is illustrated in Figure 6b by adding a lake. The river has to flow into the lake and the same lake is touching the forest. This is an example for a join in a corresponding ABox. However, the reasoning with $\mathcal{ALCRP}(\mathcal{D})$ as described above is still valid and usable for query processing.

Only the subsumption between ABox queries requires a more sophisticated approach.

4 Using ABox Patterns for n-ary Queries

We have demonstrated that the abstraction process of rewriting a query ABox to a concept term provides a means to specify the semantics of a visual query. Unfortunately, there are also some drawbacks with this approach. First of all, since the semantics of a concept description is only the set of individuals that satisfy the concept, no n-ary query results can be returned. This is not surprising, since concept descriptions correspond to first-order formulae with only one free variable. Thus, the abstraction process is only successfully applicable to ABoxes where it can select exactly one primary "target individual" from the query. The target individual remains as the single individual and the other individuals are represented by the query concept derived by the abstraction process. In the case of the cottage example, the target object of the query is the cottage the user is looking for. However, considering VISCO, the semantics of a VISCO query is a set of n-tuples, so one would need more than one free variable to fully specify the semantics of VISCO's query language which can handle aggregates.

As a solution to this problem we have developed so-called *ABox patterns*, which are ordinary $\mathcal{ALCRP}(\mathcal{S}_2)$ ABoxes that may –in addition to ordinary individuals– contain *variables*, e.g. $x?$, $y?$. Intuitively speaking, given a "GIS database ABox" $\mathcal{A}$ and a query ABox $\mathcal{Q}$ that contains variables, the *ABox pattern retrieval service* returns a set σ of *substitutions* which are mappings from variables in $\mathcal{Q}$ to ABox individuals in the database ABox $\mathcal{A}$ such that $\mathcal{A} \models_{\mathcal{T}} \sigma(\mathcal{Q})$ (w.r.t. the TBox $\mathcal{T}$).

$abox\_pattern\_retrieval(\mathcal{A}, \mathcal{Q}) := \{\sigma \mid \mathcal{A} \models_{\mathcal{T}} \sigma(\mathcal{Q})\}.$

$\sigma(Q)$ applies the given substitution σ to the query ABox $\mathcal{Q}$, replacing its variables with individuals from $\mathcal{A}$. As a subproblem, we need to decide the *ABox entailment problem*, e.g. the question whether $\mathcal{A} \models_{\mathcal{T}} \sigma(\mathcal{Q})$. This problem is decidable for $\mathcal{ALCRP}(\mathcal{S}_2)$ (see [11]). The substitutions can simply be enumerated as mappings from $Vars(\mathcal{Q}) \rightarrow Individuals(\mathcal{Q})^{\|Vars(\mathcal{Q})\|}$ and applied to $\mathcal{Q}$ yielding $\sigma(\mathcal{Q})$. If $\sigma(\mathcal{Q})$ is entailed by $\mathcal{A}$, then the range (image) of σ is the n-ary query result. Note that computing the set of individuals that are members of a concept C is a special case of an ABox pattern retrieval.

$concept\_members(\mathcal{A}, \mathsf{C}) := abox\_pattern\_retrieval(\mathcal{A}, \{x?\!:\!\mathsf{C}\}).$

As an example, we reconsider our query where a user is looking for a cottage in a forest. Obviously, the price of the estate as well as the cottage and the

forest itself are of interest, so instead of returning just the cottage the query should return triples such as $\langle\mathsf{cottage}, \mathsf{estate}, \mathsf{forest}\rangle$ which can be further inspected. The corresponding query ABox $\mathcal{Q}$ might be defined as

$\mathcal{Q} = \{x? : \mathsf{cottage}, y? : \mathsf{estate}, z? : \mathsf{forest}, (x?, y?) : \mathsf{g\_inside}, (y?, z?) : \mathsf{g\_inside}\}$.

The user can also refer to specific individuals, e.g.

$\mathcal{Q} = \{x? : \mathsf{cottage}, y? : \mathsf{estate}, (x?, y?) : \mathsf{g\_inside}, (y?, \mathsf{black\_forest}) : \mathsf{g\_inside}\}$, where $\mathsf{black\_forest}$ is a specific database ABox individual. Obviously, joins can easily be specified. Suppose we are looking for two cottages within the same estate that are located at the same river:

$\mathcal{Q} = \{x? : \mathsf{cottage}, y? : \mathsf{cottage}, x? \neq y?, z? : \mathsf{estate}, (x?, z?) : \mathsf{g\_inside},$
$\quad (y?, z?) : \mathsf{g\_inside}, r? : \mathsf{river}, (x?, r?) : \mathsf{touching}, (y?, r?) : \mathsf{touching}\}$.

Since the unique name assumption also does not hold for variables, we introduce an additional assertion $x? \neq y?$ in order to ensure that $\sigma(x?) \neq \sigma(y?)$. This simply constrains the substitution σ and has no impact on $\mathcal{ALCRP}(\mathcal{S}_2)$.

5 Conclusion

The formalism presented in this paper can be used to define the semantics of visual spatial queries, to reason about query subsumption, and to deal with multiple worlds or query completion with the help of default reasoning. The proposed ABox pattern retrieval solves the problem with joins and/or cycles in ABoxes and supports n-ary query results. However, the price is an increased query execution complexity due to the more complex inference problem.

References

[1] A. Borgida. Description logics in data management. *IEEE Transactions on Knowledge and Data Engineering*, 7(5):671–682, 1995.

[2] R.J. Brachman and J.G. Schmolze. An overview of the KL-ONE knowledge representation system. *Cognitive Science*, pages 171–216, August 1985.

[3] V. Haarslev. A fully formalized theory for describing visual notations. In K. Marriott and B. Meyer, editors, *Visual Language Theory*, pages 261–292. Springer Verlag, Berlin, 1998.

[4] V. Haarslev. A logic-based formalism for reasoning about visual representations (extended abstract). In *Proceedings, AAAI Workshop*

on Formalizing Reasoning with Visual and Diagrammatic Representations, AAAI Fall Symposium Series 1998, October 23-25, Orlando, Florida/USA. Technical Report FS-98-04, pages 57–66. AAAI Press, 1998.

[5] V. Haarslev. A logic-based formalism for reasoning about visual representations. *Journal of Visual Languages and Computing*, 10(4):421–445, 1999.

[6] V. Haarslev, C. Lutz, and R. Möller. Foundations of spatioterminological reasoning with description logics. In T. Cohn, L. Schubert, and S. Shapiro, editors, *Proceedings of Sixth International Conference on Principles of Knowledge Representation and Reasoning (KR'98), Trento, Italy, June 2-5, 1998*, pages 112–123, June 1998.

[7] V. Haarslev, C. Lutz, and R. Möller. A description logic with concrete domains and a role-forming predicate operator. *Journal of Logic and Computation*, 9(3):351–384, June 1999.

[8] V. Haarslev, R. Möller, and M. Wessel. On specifying semantics of visual spatial query languages. In *1999 IEEE Symposium on Visual Languages, Tokyo, Japan, Sep. 13-16*, pages 4–11. IEEE Computer Society Press, Los Alamitos, September 1999.

[9] V. Haarslev and M. Wessel. Querying GIS with animated spatial sketches. In *1997 IEEE Symposium on Visual Languages, Capri, Italy, Sep. 23-26*, pages 197–204. IEEE Computer Society Press, Los Alamitos, September 1997.

[10] C. Lutz. On the complexity of terminological reasoning. Technical Report LTCS-99-04, RWTH Aachen, 1999.

[11] R. Möller and M. Wessel. Terminological default reasoning about spatial information: A first step. In *Proc. of COSIT'99, International Conference on Spatial Information Theory, Stade*, pages 172–189. Springer Verlag, Berlin, 1999.

[12] A. Tarski. *A Decision Method for Elementary Algebra and Geometry.* University of California Press, Berkeley, CA, 1951.

[13] M. Wessel and V. Haarslev. VISCO: Bringing visual spatial querying to reality. In *1998 IEEE Symposium on Visual Languages, Halifax, Canada, Sep. 1-4*, pages 170–177. IEEE Computer Society Press, Los Alamitos, September 1998.

[14] W.A. Woods and J.G. Schmolze. The KL-ONE family. In F. Lehmann, editor, *Semantic Networks in Artificial Intelligence*, pages 133–177. Pergamon Press, Oxford, 1992.