Towards Computer Vision with Description Logics: Some Recent Progress

Ralf Möller, Bernd Neumann, and Michael Wessel
University of Hamburg, Computer Science Department
Vogt-Kölln-Str. 30, 22527 Hamburg, Germany

Abstract

A description logic (DL) is a knowledge representation formalism which may provide interesting inference services for diverse application areas. This paper first gives an overview of the benefits which a DL may provide for Computer Vision. The main body of the paper presents recent work at Hamburg University on extending DLs to handle spatial reasoning and default reasoning.

1: Why is Description Logic Interesting for Computer Vision?

This contribution discusses the merits of description logics (DLs) for Computer Vision (CV) and reports about some recent work on DL extensions at Hamburg University. Our goal is to make DLs more useful for diverse applications, in particular those involving concrete real-life phenomena which play a part in diagnosis, configuration and – last not least – in CV. This work extends results previously published in [9].

DL is the family name of object-based knowledge-representation formalisms in the spirit of KL-ONE which have been introduced 20 years ago [4] with the main purpose of providing formal semantics for semantic nets, thus providing the logical foundations for knowledge-based inferences. A particular DL typically realizes a particular subset of First Order Predicate Logic (FOPL). Much of the research on DLs has dealt with decidability and complexity properties depending on the expressiveness of the language. Different from a full FOPL prover, the goal is generally that a DL system should provide decidable inference services.

Typical inference services offered by a DL system are

- subsumption check,
- consistency check,
- classification,
- abstraction.

These services can be applied to the conceptual descriptions in a terminological knowledge base (TBox) and may provide obvious advantages regarding the construction and maintenance of large knowledge bases. In particular, automatic classification of concept terms allows for the semantic-based construction of concept hierarchies (taxonomies).

For the representation of factual knowledge, a DL system provides facilities for the declaration of knowledge about individual objects in the assertional knowledge base (ABox), which refers to a TBox. With an ABox it is possible to express conceptual properties of instances, for example, of the contents of a particular scene. Furthermore, relations between individuals are described. The

TBox background knowledge determines what can be inferred from the explicit declarations in an ABox. For example, an ABox object can be shown to be an instance of certain TBox concepts (instance checking inference service). In addition, the set of most specific concept names of which an individual is an instance can be computed (this process is sometimes called object classification).

For CV researchers who are primarily interested in geometric and photometric aspects of a CV problem, formal knowledge representation may seem a distant topic, relevant mainly for symbolic processing in high-level vision tasks. We will argue in the following that formal knowledge representation and, in particular, DLs may play a more significant part for vision systems than commonly recognized.

Standardized services: From a very general point of view, a DL is attractive not so much because of object-based declarative knowledge representation (this is also possible with other tools) but rather because of the standardized services which are available to a system developer. It will be a significant economical advantage if provably correct and reusable software components can be used instead of complex application-dependent software. For example, one could use the object classifier of a DL instead of programming an object recognition procedure. At this stage, we are still at the beginning of exploring the potential of DL services for specific tasks, including image interpretation. But the goal is to express knowledge-based operations of a vision system by standardized inference procedures.

Pattern classification: One way of using a DL for image interpretation is to employ the DL classifier for pattern classification tasks. This has been investigated in [10] for change detection in aerial image sequences. The idea is to conceptually define classes (in this case types of changes) in terms of sufficient conditions which must be fulfilled by image features. Given image features, a DL classifier can then automatically deduce which classes apply. As discussed in [10], this approach does not permit a hypothesize-and-test control which is indispensable for complex vision tasks. On the other hand, DL classifiers guarantee soundness, completeness and termination and can be obtained off-the-shelves, providing considerable software engineering benefits.

The logics of understanding images: Several researchers have tried to clarify the underlying logics of an image interpretation task. This is an important issue in view of the difficulties of our field to establish a consensus about the expected results and performance of vision systems. In a recent dissertation [11] a DL has been used to provide the knowledge-representation framework for model-construction (in the logical sense) which has been identified as the formal task underlying image interpretation. While this work does not provide the basis for an immediate implementation, it identifies several functional building blocks of a vision system at the logical level. The work also contains interesting evaluations of existing image interpretation formalisms and points out several deficiencies.

Interfacing to knowledge bases: Vision is often part of more complex tasks, where symbolic knowledge-representation is indispensable. For example, a robot may involve vision in planning and plan execution, based on beliefs, desires and intentions in his knowledge base. There is obviously the need to interface vision with other AI components, and it is interesting to look at the requirements for this interface. First, one must notice that factual knowledge (encoded in the ABox of a knowledge base) often provides the situational context for vision. Hence ABox reasoning is an important facility, for example, checking vision results for consistency with contextual knowledge.

Second, in order to support hypothesize-and-test processes of a vision system, the representation system should provide more services than a simple ask-and-tell interface. In particular, it must be possible to generate expectations which restrict and prioritize possible hypotheses. One approach we pursue is to generate expectations using default reasoning. Important foundations for defaults in DLs are due to [2]. Another approach is to extend DLs with probabilities so that "soft" classifications and ordered hypotheses can be supported [7]. Probabilistic hypothesis generation would certainly meet important requirements from the CV side. On the other hand, central notions of formal knowledge representation, e.g. consistency, lose their traditional meaning.

Dealing with space and time: The generality and application-independence of symbolic logic formalisms is an advantage with respect to validity and reusability, but may be a severe impediment when domain-specific properties and laws must be exploited for a task. One of the most interesting extensions of DLs has been the incorporation of "concrete domains" [1]. Under certain conditions, objects and relations of a concrete domain (e.g. real numbers, strings, polygons) can be built into a DL so that knowledge representation and reasoning can be performed with other than purely symbolic objects. For several application areas including vision it is important to reason about space and time. In particular in high-level vision, many interesting concepts can be described as spatial and temporal aggregates of objects. For example, an overtake event can be described conceptually as an aggregate of individual object motions which are temporally and spatially related. The construction of a recognition system for overtake occurrences and other spatiotemporal aggregates could be facilitated considerably if consistency checks and other services could be extended to incorporate spatial and temporal theories.

In the following sections we present recent work at Hamburg University on extending DLs in the spirit described above. In particular, we investigate reasoning about spatial information with the DL $\mathcal{ALCRP}(\mathcal{S}_2)$ where $\mathcal{S}_2$ denotes a particular concrete domain which is used to model topological relations. We show that the inbuilt topological reasoning power of this DL can be used to control default reasoning, for example for hypothesis generation. While this may be only a modest advance on the way to a full-fledged DL-based vision system, we present the work in some detail to demonstrate some of the problems and subtleties with which one has to deal when extending DLs.

2: Description Logics

In order to demonstrate modeling with description logics, we briefly discuss some examples. Let *man* be a concept. Then the concept *man* $\sqcap$ $\exists$*offspring*.*human* describes all men which are related to at least one human via the offspring role (existential quantification). Thus, the concept term given above could have been named *father* with the terminological axiom *father* $\doteq$ *man* $\sqcap$ $\exists$*offspring*.*human*.

It can be seen as a limitation that standard description logics can only handle abstract knowledge. Imagine that we want to represent the ages of humans as natural numbers. This cannot be done in most description logics. There are, however, some DL formalisms which overcome this limitation and are able to additionally represent knowledge about so-called concrete objects such as numbers and polygons. One important formalism of this type is the language $\mathcal{ALC}(\mathcal{D})$ defined by [1]. With this language, one could define an old person as *human* $\sqcap$ $\exists age. >_{60}$. Here, *age* is a single-valued role (those roles are called features). The feature *age* attaches concrete objects that represent natural numbers to abstract objects (in this case of type *human*). The extension of the above-mentioned concept term is "All humans who are older than 60 years." The example demonstrates that defining concept terms based on predicates over concrete objects (e.g. "$>_{60}$") greatly extends the expressiveness of the knowledge representation formalism.

The language $\mathcal{ALCRP}(\mathcal{D})$ defined in [8] goes one step further. It also allows one to define roles based on predicates over concrete objects. Like in the $\mathcal{ALC}(\mathcal{D})$ example above, predicates over concrete objects that are attached to abstract objects via features can be seen as properties of these abstract objects. Take again humans and their ages as an example. The age is a property of each object which is of type *human* (it is a concrete object attached via the *age* feature). Assume that we would like to define the concept *oldest-person*. The extension of this concept does not have a cardinality greater than one unless there are some people which have the same age. In $\mathcal{ALCRP}(\mathcal{D})$, one could use the term $human \sqcap \neg\exists older.human$, where *older* is a defined role whose extension is the set of those pairs of objects (a, b) such that the natural number attached to object b via the feature *age* is greater then the natural number attached to object a via the same feature. Thus, only those objects of type *human* are inside the extension of the concept for which no other object exists that is older and also of type *human*. An equivalent formalization that takes all domain objects into account cannot be expressed using $\mathcal{ALC}(\mathcal{D})$.

$\mathcal{ALCRP}(\mathcal{D})$ is a very powerful language for reasoning about abstract and concrete knowledge. Like $\mathcal{ALC}(\mathcal{D})$ it can be parameterized with a concrete domain, which is a set of concrete objects plus a set of predicates over these concrete objects. Unfortunately, reasoning in $\mathcal{ALCRP}(\mathcal{D})$ is undecidable in general as proven in [8]. In [5] syntactic restrictions to be posed on $\mathcal{ALCRP}(\mathcal{D})$-terminologies are introduced. It is shown that w.r.t. these so-called restricted terminologies sound and complete algorithms for deciding the common reasoning problems exist. Decidability is achieved by restricting the free combinability of operators in restricted terminologies. Some combinations of value and exists restrictions are not allowed if they quantify over defined roles. Furthermore, the use of the concept forming predicate operator known from $\mathcal{ALC}(\mathcal{D})$ has to be restricted, too. These restrictions are solely motivated by decidability issues. From the knowledge engineer's point of view they are relatively strong constraints on the possible structure of concept terms. Another approach for defining a decidable version of $\mathcal{ALCRP}(\mathcal{D})$ would have been to pose limitations on the allowed set of predicates that can be used with concept- and role-forming operators. But this seems to be less promising because the intended areas of application, representing time and space, already require fairly complex predicates which presumably cause undecidability of the resulting language.

In the following we define the syntax of role and concept terms in $\mathcal{ALCRP}(\mathcal{D})$. The formal semantics is given in [6].

Definition 1 Let R and F be disjoint sets of role and feature names, respectively. For brevity we also use the terms roles and features. Any element of $\mathsf{R} \cup \mathsf{F}$ is an *atomic* role term. A composition of features (written $f_1 f_2 \cdots$) is called a feature chain. A simple feature can be viewed as a feature chain of length 1. If P is a predicate name from $\mathcal{S}_2$ with arity $n+m$ and $u_1, \ldots, u_n$ as well as $v_1, \ldots, v_m$ are feature chains, then the expression $\exists(u_1, \ldots, u_n)(v_1, \ldots, v_m).P$ (*role-forming predicate restriction*) is a *complex* role term. Let S be a role name and let T be a role term. Then $S \doteq T$ is a *terminological axiom*.

Definition 2 Let C be a set of concept names which is disjoint to R and F. Any element of C is a *concept term* (*atomic* concept term). If C and D are concept terms, R is a role term, P is a predicate name from $\mathcal{S}_2$ with arity n, and $u_1, \ldots, u_n$ are feature chains, then the following expressions are also concept terms: $C \sqcap D$ (*conjunction*), $C \sqcup D$ (*disjunction*), $\neg C$ (*negation*), $\exists R.C$ (*exists restriction*), $\forall R.C$ (*value restriction*), and $\exists u_1, \ldots, u_n.P$ (*predicate exists restriction*).

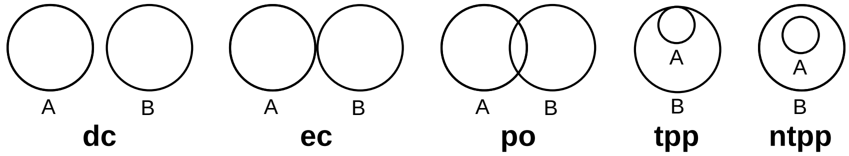

Figure 1. Elementary relations between two regions A and B.

For all kinds of exists and value restrictions, the role term or the list of feature chains may be written in parentheses. Let A be a concept name and let D be a concept term. Then $A \doteq D$ (equivalence) and $A \sqsubseteq D$ (implication) are terminological axioms as well. A finite set of terminological axioms $\mathcal{T}$ is 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.

$\mathcal{ALCRP}(\mathcal{S}_2)$ is the description logic resulting from the instantiation of the description logic $\mathcal{ALCRP}(\mathcal{D})$ with the concrete domain $\mathcal{D} = \mathcal{S}_2$ (see [5, 6]).

Definition 3 The concrete domain $\mathcal{S}_2$ is defined w.r.t. the topological space $\langle \mathbb{R}^2, 2^{\ \ 2} \rangle$. The domain $\Delta_{\mathcal{S}_2}$ contains all non-empty, regular closed subsets of $\mathbb{R}^2$ which are called *regions* for short. The set of predicate names is defined as follows:

- A unary $concrete\_domain\_top$ predicate is-region with $\text{is-region}^{\mathcal{S}_2} = \Delta_{\mathcal{S}_2}$ and its negation is-no-region with $\text{is-no-region}^{\mathcal{S}_2} = \emptyset$.
- The 8 basic predicates dc, ec, po, tpp, ntpp, tppi, ntppi and eq correspond to the RCC-8 relations (Figure 1). Due to space restrictions we would like to refer to [6] for a formal definition of the semantics.
- In order to name disjunctions of base relations, we need additional predicates. Unique names for these "disjunction predicates" are enforced by imposing the following canonical order on the basic predicate names: dc, ec, po, tpp, ntpp, tppi, ntppi, eq. The tppi relation (ntppi) is the inverse of the tpp (resp. ntppi) relation; all others are symmetrical. For each sequence $p_1, \ldots, p_n$ of basic predicates in canonical order ($n \geq 2$), an additional predicate of arity 2 is defined. The predicate has the name $\mathsf{p_1}\text{-}\cdots\text{-}\mathsf{p_n}$ and we have $(r_1, r_2) \in \mathsf{p_1}\text{-}\cdots\text{-}\mathsf{p_n}^{\mathcal{S}_2}$ iff $(r_1, r_2) \in \mathsf{p_1}^{\mathcal{S}_2}$ or $\ldots$ or $(r_1, r_2) \in \mathsf{p_n}^{\mathcal{S}_2}$. The predicate dc-ec-po-tpp-ntpp-tppi-ntppi-eq is also called spatially-related.
- A binary predicate inconsistent-relation with $\text{inconsistent-relation}^{\mathcal{S}_2} = \emptyset$ is the negation of spatially-related.

3: Terminological Default Reasoning

In the following we investigate a Reiter-based approach to terminological default reasoning about spatial information. Originally, a default rule has the form

$$\frac{\alpha : \beta_1, \beta_2, \ldots, \beta_n}{\gamma}$$

(also written $\alpha : \beta_1, \beta_2, \ldots, \beta_n \ / \ \gamma$), where α, β_i and γ are FOPL formulae. α is called the *precondition* of the rule, the β_i terms are called *justifications,* and γ is the *consequent*. Intuitively the idea behind default reasoning is the following: starting with a world description A of what is known to be true, default rules can be applied such that they augment A by default rule conclusions γ to yield a *set of beliefs*. A default can be applied, i.e. its conclusion γ can be added to the set of current beliefs iff α is entailed by this set, each formula β_i is consistent with the current set of beliefs and γ is not already entailed.

Defaults may interact and depending on the set of default rules being applied, different "possible worlds" or hypotheses can be computed. These possible worlds are referred to as *extensions* (see below for a formal definition). Depending on the reasoning mode the *consequence problem* for terminological default theories is to decide whether a given assertional axiom is member of all extensions (skeptical mode) or of at least one extension (credulous mode).

Using description logic *concept terms* in default rules instead of first-order or propositional logic formulae has been extensively considered in [2]. A *terminological default theory* is a pair (A, D) where A is an ABox, and D is a finite set of *terminological* default rules whose preconditions, justifications and consequents are concept terms. Because concept terms correspond to unary predicates ranging over a free variable, these defaults are called *open* defaults. In contrast, *closed* defaults do not contain any free variables. Unlike Reiter's original proposal, the approach of [2] applies defaults only to those individuals that are explicitly mentioned in the world description (ABox). Default rules are never applied to implicit individuals introduced by $\exists$-restrictions. With this kind of semantics the consequence problem for (A, D) is decidable (see [2] for details). Closed default rules can be obtained by instantiating the free variable in the concept expressions with all explicitly mentioned ABox individuals (see [2] for a formal definition). Thus, for closed defaults, α, β_i and γ are *concept membership assertions* (ABox concept axioms).

Once we have a closed default theory, a set of consequences of such a theory is referred to as an *extension* which is a set of deductively closed formulae defined by a fixed point construction. In the case of terminological default reasoning about spatial information it is also interesting to conclude spatial relations by default. Therefore, we extended the approach presented in [2] to be able to deal with role assertions in default rules. This can be achieved by allowing $\mathcal{ALCRP}(\mathcal{S}_2)$ *ABoxes* inside the default rules as α, β_i and γ. Before discussing the computation of extensions of such closed default theories, we first consider some examples of using defaults in the context of terminological reasoning about spatial information.

4: Examples for Spatioterminological Default Reasoning

We will now illustrate the use of a DL with integrated topological reasoning for an example which could be part of an aerial image interpretation task. The idea is to use defaults for hypothesis generation regarding the classification of areas in an image. The default reasoning component of the DL will generate extensions of the ABox representing hypothesized classifications which are consistent with the rest of the knowledge base. The consistency check involves spatial reasoning. Additionally, also spatial relationships between areas could be hypothesized, for example, in case of partial object occlusions (see below).

4.1: Example 1

Suppose we have incomplete knowledge regarding the classification of the object b in Figure 2(a). We already know that a is a country, but area b is only known to be an area. The image interpretation system may want to generate possible hypotheses for b – b could be a city (Figure 2(b)), but could

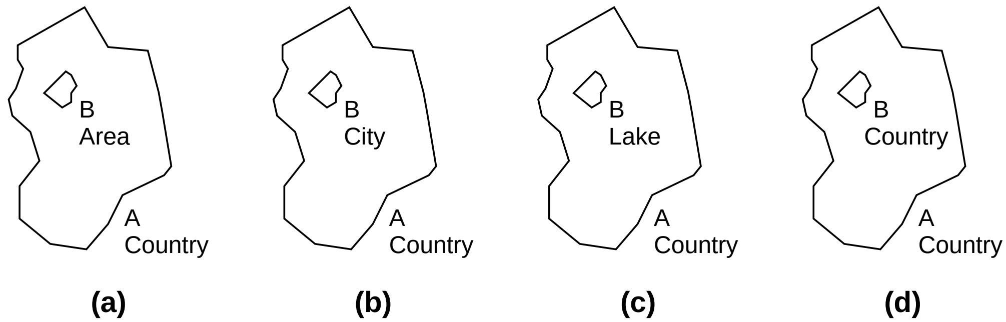

Figure 2. Generation of hypotheses for object B.

also be a lake (Figure 2(c)), both plausible hypotheses w.r.t. the size of area b. Obviously, these different hypotheses are disjoint, since b cannot be both a city and a lake. Other hypotheses are not generated although these might be plausible at first sight. In particular, since we require that countries are always disjoint (relation dc) or touching (relation ec), the system deduces that the hypothesis shown in Figure 2(d) should not be generated.

4.1.1: Formalizing the Example

Using $\mathcal{ALCRP}(\mathcal{S}_2)$'s role-forming predicate-based operator, we define a set of complex roles according to the mentioned RCC-8 $\mathcal{S}_2$ predicates:

$$\begin{array}{rcl} inside & \doteq & \exists(has\_area)(has\_area).\mathsf{tpp\text{-}ntpp} \\ contains & \doteq & \exists(has\_area)(has\_area).\mathsf{tppi\text{-}ntppi} \\ overlaps & \doteq & \exists(has\_area)(has\_area).\mathsf{po} \\ touches & \doteq & \exists(has\_area)(has\_area).\mathsf{ec} \\ disjoint & \doteq & \exists(has\_area)(has\_area).\mathsf{dc} \end{array}$$

The following definitions of concepts are required to model domain objects representing different kinds of regions in a TBox which satisfies the $\mathcal{ALCRP}(\mathcal{D})$ restrictedness criteria. This conceptual background knowledge also applies to the subsequent examples.

$$\begin{array}{rcl} area & \doteq & \exists(has\_area).\mathsf{is\text{-}region} \\ natural\_region & \doteq & \neg administrative\_region \\ country\_region & \dot{\sqsubseteq} & administrative\_region \sqcap \\ & & large\_scale \sqcap area \\ city\_region & \dot{\sqsubseteq} & administrative\_region \sqcap \\ & & \neg large\_scale \sqcap area \\ lake\_region & \dot{\sqsubseteq} & natural\_region \sqcap area \\ river\_region & \dot{\sqsubseteq} & natural\_region \sqcap area \end{array}$$

An $area$ is a two-dimensional region with some extent. Furthermore, we distinguish between $administrative\_regions$ and $natural\_regions$ which are disjoint concepts. The difference between a $country\_region$ and a $city\_region$ is that the former is $large\_scale$, but the latter is not. Thus, these two concepts are disjoint as well. The intention behind the other concepts should be obvious.

$$\begin{array}{rcl} country & \doteq & country\_region \sqcap \\ & & \forall contains.\neg country\_region \sqcap \\ & & \forall overlaps.\neg country\_region \sqcap \\ & & \forall inside.\neg country\_region \\ city & \doteq & city\_region \sqcap \\ & & \exists inside.country\_region \\ lake & \dot{\sqsubseteq} & lake\_region \\ river & \doteq & river\_region \sqcap \\ & & \forall overlaps.\neg lake\_region \sqcap \\ & & \forall contains.\bot \sqcap \\ & & \forall inside.\neg lake\_region \end{array}$$

A $country$ is a $country\_region$ that can never contain or be contained within other $country\_regions$. Also, $countries$ never overlap other $country\_regions$. Each $city$ must belong to a specific $country$, i.e. must lie within a $country$. Unfortunately, we cannot write this directly as $\exists inside.country$ because the unfolded resulting term is no longer restricted. So, we have to use the somewhat weaker version with the base concept $country\_region$. In our world model a $city$ must be inside a $country$. For a $river$ we require that it never $overlaps$ or is $inside$ with a $lake\_region$.

$$river\_flowing\_into\_a\_lake \doteq river \sqcap \exists touches.lake\_region$$

A $river\_flowing\_into\_a\_lake$ is a specific $river$ that $touches$ a $lake\_region$ (recall that the RCC-8 relations ec and po and also ec and ntpp-tpp are disjoint). It would be reasonable to also state that cities do not overlap other cities etc., but this is ignored here for the sake of brevity.

We have seen that $\mathcal{ALCRP}(\mathcal{S}_2)$ provides the necessary expressiveness to model domain objects in our geographic information system scenario. In [6] more examples for the use of $\mathcal{ALCRP}(\mathcal{D})$ are given, which also demonstrate the influence of spatial reasoning on TBox reasoning (subsumption of concepts).

Formalizing hypothesis generation in the way we already discussed informally, we now consider the following spatioterminological *default rules* d_1, d_2 and d_3:

$$d_1 = \frac{area\ :\ city}{city} \qquad d_2 = \frac{area\ :\ lake}{lake} \qquad d_3 = \frac{area\ :\ country}{country}$$

Suppose we have an ABox according to our world description as shown in Figure 2(a):

$$\{a : country, b : area, (a,b) : contains, (b,a) : inside\}$$

Closing defaults, i.e. instantiating the defaults d_1, d_2, d_3 over the ABox individuals a and b yields 6 different closed defaults. Now, let us assume α, β and γ have been replaced by the corresponding *assertional axioms* (e.g. instantiating the default $area\ :\ city\ /\ city$ with the individual a yields the closed default rule $\{a\ :\ area\}\ :\ \{a\ :\ city\}\ /\ \{a\ :\ city\}$ – expressions like $a\ :\ city$ are called assertional axioms or ABox axioms). We use the notation $d_i(ind)$ to refer to a default that is instantiated with the individual ind. Given our 6 closed default rules let us examine the status of each:

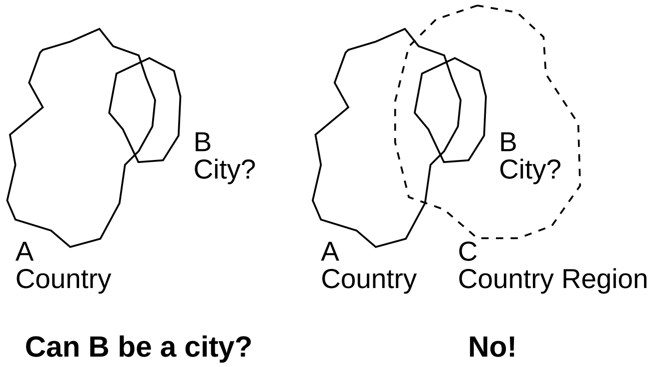

Figure 3. Subtle inferences due to topological constraints.

- Default $d_1(a)$ cannot be applied because adding $a : city$ to the ABox yields a contradiction with $a : country$. The concepts $country\_region$ and $city\_region$ are disjoint (due to $large\_scale$ and $\neg large\_scale$).
- Default $d_1(b)$ can be applied. We get an augmented ABox or *Extension 1*, see Figure 2(b):

$$\{a : country, b : area, b : city, (a, b) : contains, (b, a) : inside\}$$

- Default $d_2(a)$ cannot be applied because adding $a : lake$ to the ABox yields a contradiction with $a : country$. A $country$ is an $administrative\_region$ and a $lake$ is defined as a $natural\_region$, and both are disjoint concepts.
- Default $d_2(b)$ can be applied. Thus, we can get an augmented ABox or *Extension 2*, see Figure 2(c):

$$\{a : country, b : area, b : lake, (a, b) : contains, (b, a) : inside\}$$

  However, if we have an ABox already augmented by the conclusion of default $d_1(b)$, $b : city$, we cannot apply $d_2(b)$. So, only one of $d_1(b)$ or $d_2(b)$ can be applied, resulting in two different *extensions.*
- Default $d_3(a)$ cannot be applied, because its conclusion is already entailed by the ABox.
- Default $d_3(b)$ cannot be applied even if no other default has been applied before. Adding the default's consequent $b : country$ would yield an inconsistent ABox because a is already known to be a $country$ and so, among others, $a : \forall contains.\neg country\_region$ holds. Because $(a, b) : contains$ holds and $b : country$ would imply $b : country\_region$, the default cannot be applied. Thus, we cannot get an extension corresponding to the wrong interpretation in Figure 2(d).

4.2: Example 2

Another subtle inference can be demonstrated by showing that the default $d_1(b)$ (as defined above) cannot be applied to conclude that object b in Figure 3 is a $city$. Figure 3 corresponds to the ABox or world description

$$\{a : country, b : area, (a, b) : overlaps, (b, a) : overlaps\}$$

Trying to assert $b : city$ would result in a constraint $b : city\_region \sqcap \exists inside.country\_region$. Therefore, polygon a cannot be the appropriate $country\_region$ because $(b, a) : overlaps$ holds.

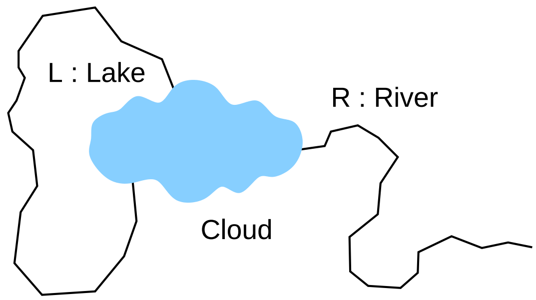

Figure 4. Incomplete spatial information.

Due to the exists restriction there exists an implicit individual c which is a $country\_region$ such that $(b, c) : inside$ holds. As can be seen in Figure 3, there is no way to find a spatial arrangement such that b is inside c and c does not overlap with a or does not contain a. Because a is a $country$ and, therefore, may not overlap or may not be contained in another $country\_region$, there is no way to conclude that b could possibly be a $city$.

4.3: Example 3

As already mentioned, terminological default rules like the ones used in the previous examples have already been exploited by Baader & Hollunder (but not in an image interpretation context).

Let us consider Figure 4. In this case, we only have *incomplete spatial information* w.r.t. to the topological relationship between r and l, because a cloud occludes relevant parts of the two objects. The corresponding ABox is

$$\{l : lake, r : river\}$$

Since we already know that l is a lake and r is a river (perhaps this is also a hypothesis generated by previous default rule applications), we can conclude from our conceptual background knowledge that the spatial relationship between the river and the lake must be either ec (touches) or dc (disconnected or disjoint). There are no other possibilities, e.g., a river never overlaps a lake and is never contained within a lake. We can therefore hypothesize these two possible spatial relationships by default rule applications. This shows that not only concept or class memberships can be deduced by defaults. The important insight is the following duality: We can either use spatial relations between object pairs to conclude their concept memberships, or we can use already known concept memberships to conclude particular spatial relations between objects.

Unfortunately, relationship conclusions cannot be expressed with the terminological default rules introduced so far, because α, β_i and γ are limited to *concept expressions*. This shows why we extended the terminological default rules introduced in [2] by permitting so-called *ABox patterns* instead of concept expressions for α, β_i and γ ([9]). ABox patterns are basically ABoxes with placeholders for individuals (written with capital letters). Closing the default rules instantiates the patterns with all possible combinations of individuals yielding closed defaults whose α, β_i and γ are $\mathcal{ALCRP}(\mathcal{S}_2)$ ABoxes:

$$d_4 = \frac{\{X : lake, Y : river\} \ : \ \{(X, Y) : disjoint\}}{\{(X, Y) : disjoint\}}$$

$$d_5 = \frac{\{X : lake, Y : river\} \ : \ \{(X, Y) : touches\}}{\{(X, Y) : touches\}}$$

Closing the patterns, i.e. instantiating X, Y over the ABox $A = \{l : lake, r : river\}$ would yield eight different closed defaults whose α, β_i and γ are $\mathcal{ALCRP}(\mathcal{S}_2)$ ABoxes, e.g. instantiating d_4 with $X \leftarrow l, Y \leftarrow r$ yields the closed default rule

$$\frac{\{l : lake, r : river\} \ : \ \{(l, r) : disjoint\}}{\{(l, r) : disjoint\}}$$

Additionally, as well as allowing variables such as X and Y, one might also be able to refer to specific ABox individuals in the ABox patterns (for instance, the individual "Bodensee").

4.3.1: Default Reasoning with Specificity

Let us consider the world description

$$A = \{r : river\_flowing\_into\_a\_lake, l : lake\}$$

Since it is already known that r is really a $river\_flowing\_into\_a\_lake$ and not only a $river$, we would like to conclude that the lake l in A should be *the* lake. That is, the complex role assertion $(l, r) : touches$ should be added:

$$d_6 = \frac{\{X : lake, Y : river\_flowing\_into\_a\_lake\} \ : \ \{(X, Y) : touches\}}{\{(X, Y) : touches\}}$$

In the case of d_6, we would like to render the application of d_4 and d_5 *invalid*, because they are "less specific" than d_6 (even if d_5 yields the same conclusion, $touches$).

A default d_a is said to be more specific than d_b, $d_a \prec d_b$ iff $(\alpha(d_a) \models \alpha(d_b)) \wedge (\alpha(d_b) \not\models \alpha(d_a))$ where $\alpha(d)$ denotes the precondition of the default d. Algorithms for computing the so-called *S-extensions* (*S* for specificity) have already been developed by Baader and Hollunder [3]. There is a strong conjecture that these algorithms can be applied in our $\mathcal{ALCRP}(\mathcal{S}_2)$ context as well. In contrast, the ordinary extensions are called *R-extensions* (*R* for Reiter). In our example, we would get two different R-extensions, but only one S-extension containing the ABox axiom $(r, l) : touches$. The other *R-extension* containing $(r, l) : disjoint$ could not be derived, since only the most specific active defaults are applied when computing S-extensions. This would render the application of d_4 and d_5 impossible because d_6 is also active and more specific than both d_4 and d_5.

This concludes the illustrating examples. We have shown that standardized reasoning services of a DL can be used to generate hypotheses consistent with the available knowledge. This is, of course, only one of several building blocks required for image interpretation. Questions regarding the ordering of extensions, the verification of possible extensions with additional evidence, the incorporation of metric information a.o. have not been treated and in many cases cannot be answered. Below the line, however, we hope that the value of inference services has been demonstrated.

In the next section we will show that the *consequence problem* is decidable for terminological default theories with default rules containing $\mathcal{ALCRP}(\mathcal{S}_2)$ ABoxes. Since we can always obtain ordinary ABoxes from our ABox patterns by closing them, the consequence problem is decidable for defaults with ABox patterns as well.

5: Computing Extensions

Intuitively, given a closed terminological default theory (A, D) a deductively closed set of consequences of such a theory is referred to as an *extension*. As usual, the exact definition is given by a fixpoint construction. We cite a formal definition taken from [2]. $Th(\Gamma)$ stands for the deductive closure of a set of formulae Γ. In a description logic context Γ is an ABox.

Definition 4 Let E be a set of closed formulae and (A, D) be a closed default theory. We define $E_0 := A$ and for all $i \geq 0$

$$E_{i+1} := E_i \cup \{\gamma \mid \alpha : \beta_1, \ldots, \beta_n/\gamma \in D, \alpha \in Th(E_i), \neg\beta_1, \ldots, \neg\beta_n \notin Th(E)\}$$

Then, $Th(E)$ is an extension of (A, D) iff

$$Th(E) = \bigcup_{i=0}^{\infty} Th(E_i)$$

Note that, in principle, this definition for an extension $Th(E)$ has a non-constructive nature because in the definition the deductive closure $Th(E)$ is already used in each iteration step. Nevertheless, as we will see below, the definition induces an algorithm for actually computing extensions if the implicit entailment subproblems in the definition are decidable (see also [2]).

In order to be able to infer spatial relations between domain objects, the basic terminological default reasoning approach described in [2] is adapted. The basic idea is that the precondition, the justifications and the consequent of a default can be ABoxes.

Definition 5 A *spatioterminological default rule* d (or spatioterminological default for short) has the form $d = \alpha : \beta_1 \ldots \beta_n \; / \; \gamma$ where α, β_i and γ are consistent and restricted $\mathcal{ALCRP}(\mathcal{S}_2)$ ABoxes which may, among others, contain predicate-based role axioms of the form (a, b) : $\exists$(*has_area*)(*has_area*).*P* with P being an $\mathcal{S}_2$ predicate of arity two. A *spatioterminological default theory* is a tuple (A, D) where D is a set of spatioterminological default rules and A is a consistent and restricted $\mathcal{ALCRP}(\mathcal{S}_2)$ ABox.

Lemma 1 A restricted $\mathcal{ALCRP}(\mathcal{S}_2)$ ABox axiom δ is logically entailed by a restricted $\mathcal{ALCRP}(\mathcal{S}_2)$ ABox A,

$$A \models \delta, \qquad \text{iff} \qquad \begin{cases} \delta = a : C \longrightarrow \\ \qquad \neg SAT(A \cup \{a : \neg C\}) \\ \delta = (a, b) : R \longrightarrow \\ \qquad \neg SAT(A \cup \{a : \forall R.X_{new}, b : \neg X_{new}\}), \\ \delta = (a, b) : f \longrightarrow \\ \qquad \neg SAT(A \cup \; \{a : \forall f.X_{new} \sqcup \\ \qquad\qquad \exists(f).\textsf{is-region}, b : \neg X_{new}\}), \\ \delta = (a, x) : f \longrightarrow \\ \qquad \neg SAT(A \cup \; \{a : \exists(f).\Psi \sqcup \exists f.\top \sqcup \forall f.\bot, x : \overline{\Psi}\}), \\ \delta = (x_1, x_2) : P \longrightarrow \\ \qquad \neg SAT(A \cup \{(x_1, x_2) : \overline{P}\}) \\ \delta = (a, b) : \exists(u)(v).P \longrightarrow \\ \qquad \neg SAT(A \cup \{(a, b) : \exists(u)(v).\overline{P}\}) \wedge \\ \qquad \neg SAT(A \cup \{a : \forall u.\top\}) \wedge \neg SAT(A \cup \{b : \forall v.\top\}), \\ \qquad \text{where } u = v = has\_area, \end{cases}$$

where X_{new} is a new atomic concept that does not appear elsewhere in the ABox A. X_{new} is used as a "marker" concept. Analogously, Ψ (resp. $\overline{\Psi}$) is a new (otherwise unused) concrete domain "marker" predicate. These two predicates have the property that they do not interact with the other concrete domain predicates P_i. Therefore, the two arbitrary conjunctions of concrete domain predicates $\bigwedge_{i=1}^{k} P_i \wedge \Psi$ and $\bigwedge_{i=1}^{k} P_i \wedge \overline{\Psi}$ are satisfiable iff $\bigwedge_{i=1}^{k} P_i$ is satisfiable. However, $\bigwedge_{i=1}^{k} P_i \wedge \Psi \wedge \overline{\Psi}$ is always unsatisfiable, regardless of the satisfiability of $\bigwedge_{i=1}^{k} P_i$. Additionally, R is a primitive role and f is a feature. $SAT(A)$ decides the ABox consistency problem for an ABox A. Please note that a, b are interpreted as abstract domain objects, unlike x, x_1, x_2 which are interpreted as concrete domain objects. The concrete domain S_2 and the abstract domain are disjoint.

Proof 1 (Sketch) The first case is the instance checking problem, which is decidable because C is a restricted concept term. The second case deals with primitive role assertions. In case that b is an R successor of a, the assertion $a : \forall R.X_{new}$ would entail $b : X_{new}$, where X_{new} is a new (otherwise unused) atomic "marker" concept. This would obviously contradict the assertion $b : \neg X_{new}$. The same trick can be applied to check whether $(a, b) : f$ holds. Unlike primitive roles, the f successor of a might be a concrete domain object, which would also contradict the assertion $a : \forall R.X_{new}$. However, we can check for the presence of a concrete domain filler f of a by asserting $\exists(f).\mathsf{is\text{-}region}$. To check if $(a, x) : f$ holds, we cannot propagate a X_{new} marker, since $x : X_{new}$ yields an immediate contradiction (recall that the concrete domain and the abstract domain are disjoint). We therefore have to propagate a new, otherwise unused concrete domain "marker" predicate Ψ. As stated above, Ψ (resp. $\overline{\Psi}$) does not affect the satisfiability of the other concrete domain predicates P_i, and therefore the only possibility to get a contradiction with respect to Ψ ($\overline{\Psi}$) is to have asserted both $\Psi(x)$ and $\overline{\Psi(x)}$ for a concrete domain object x. However, we do not want to infer $(a, x) : f$ if a has an f successor in the abstract domain or can not have an f successor in the concrete or abstract domain. We therefore check for the presence of an abstract domain filler f of a by asserting $\exists f.\top$ and additionally check whether it is known that a can't have an f successor by asserting $a : \forall f.\bot$. In the fifth case we must decide whether the binary concrete domain predicate P holds for the concrete domain objects x_1, x_2. There exists a concrete domain predicate $\overline{P}$, the negation of P. The last case is more problematic, because the $\mathcal{ALCRP}(\mathcal{S}_2)$ language does not provide a negation operator for predicate-based role axioms. However, we can check whether $(a, b) : \exists(has\_area)(has\_area).\overline{P} \;\vee\; a : \neg\exists(has\_area).\mathsf{is\text{-}region} \;\vee\; b : \neg\exists(has\_area).\mathsf{is\text{-}region}$ holds. The NNF of $\neg\exists(has\_area).\mathsf{is\text{-}region}$ is $\exists(has\_area).\mathsf{is\text{-}no\text{-}region} \sqcup \forall(has\_area).\top$. Since $\exists(has\_area).\mathsf{is\text{-}no\text{-}region}$ is inconsistent, the resulting term is $(a, b) : \exists(has\_area)(has\_area).\overline{P} \vee a : \forall has\_area.\top \vee b : \forall has\_area.\top$. Obviously, this is not an $\mathcal{ALCRP}(\mathcal{S}_2)$ ABox. However, $A \cup \{a_1 \vee a_2 \vee \cdots \vee a_n\}$ is inconsistent iff $\forall a_i : A \cup \{a_i\}$ is inconsistent. Note that the predicate name $\overline{P}$ exists because the concrete domain is required to be admissible. $\Box$

Theorem 1 The consequence problem for a spatioterminological default theory (A, D) is decidable.

Proof 2 Considering the sound and complete tableaux calculus for deciding the consistency of restricted $\mathcal{ALCRP}(\mathcal{S}_2)$ ABoxes, $\delta \in Th(\Gamma)$ iff $\Gamma \models \delta$. Thus, instead of taking $Th(E)$ we can view the ABox E as a representative for an extension. The fixpoint construction in Definition 4 can be used as a tester for determining whether a given ABox E really *is* an extension of a default theory (A, D). Since each extension E is an ABox having the form $A \cup \{\gamma \,|\, \alpha : \beta_1 \dots \beta_n/\gamma \in D'\}$ for a set of so-called *generating defaults* $D' \subseteq D$, we can simply check for each element E of $\{A \cup X \mid X \in 2^{\{\gamma \,|\, \alpha:\beta_1 \dots \beta_n/\gamma \in D\}}\}$ whether it is an extension or not.

The following inference problems need to be decided:

1. $\alpha \in Th(E_i)$: This can be easily tested by checking whether $E_i \models \alpha$ where $\alpha = \{a_1, a_2, \ldots, a_n\}$. We can decide this *ABox entailment problem* iff we can decide whether each assertional axiom a_i is logically entailed by E_i, i.e. $\forall a_i \in \alpha : E_i \models a_i$. This can be decided according to Lemma 1.
2. $\neg\beta_i \notin Th(E)$: This can be checked by testing whether $E \not\models \neg\beta_i$. However, $E \not\models \neg\beta_i$, where $\beta_i = \{b_1, b_2, \ldots, b_n\}$ iff $A \cup \beta_i$ is consistent. The ABox consistency problem for restricted $\mathcal{ALCRP}(\mathcal{S}_2)$ ABoxes is decidable.
3. $Th(E) = \bigcup_{i=0}^{\infty} Th(E_i)$: The fixpoint can be constructed in a finite number of steps because we consider only a finite number of defaults. In principle, we have to decide the *ABox equivalence problem*. An ABox A_1 is equivalent to an ABox A_2, $A_1 \equiv A_2$ iff $A_1 \models A_2$ and $A_2 \models A_1$, i.e. the ABox equivalence problem can be reduced to two ABox entailment problems. □

In [2] another algorithm is discussed for computing extensions. This algorithm seems to be more efficient in the average case. There is a strong conjecture that the algorithm is also applicable in the $\mathcal{ALCRP}(\mathcal{S}_2)$ context. Furthermore, it can easily be seen that the results for spatioterminological default theories wrt. $\mathcal{ALCRP}(\mathcal{S}_2)$ can be extended to $\mathcal{ALCRP}(\mathcal{D})$ as well.

6: Conclusion

To the best of our knowledge we have proposed a first theory for spatioterminological default reasoning. Our spatioterminological default approach extends previous work done in [5, 6]. The new contributions to [2] are: As a base language, the expressive spatioterminological description logic $\mathcal{ALCRP}(\mathcal{S}_2)$ is used. Allowing not only concept terms as formulae in default rules but also restricted $\mathcal{ALCRP}(\mathcal{S}_2)$ ABoxes with complex role assertions is necessary from an application-oriented point of view but imposes a number of theoretical problems. We have shown that the possible extensions of a closed $\mathcal{ALCRP}(\mathcal{S}_2)$ spatioterminological default theory can be effectively computed.

An implementation of $\mathcal{ALCRP}(\mathcal{D})$ is described in [12]. With the implementation of the $\mathcal{ALCRP}(\mathcal{D})$ default reasoning substrate, an implementation of an $\mathcal{ALCRP}(\mathcal{D})$ TBox and ABox management system as well as an RCC-8 relation network consistency checker is also available for research purposes. Qualitatively speaking, tests with the current implementation indicate that for small problems with few ABox assertions, results can be expected in a reasonable time but runtimes dramatically increase when more than only a few individuals are involved.

As pointed out before, spatioterminological default reasoning is an important service for constrained hypothesis generation in vision systems. To develop the underlying foundations is a necessary step towards knowledge-based vision system architectures, where powerful inference services can be employed instead of costly and error-prone application-specific programming.

7: Acknowledgments

We would like to thank Volker Haarslev, Carsten Lutz and Anni-Yasmin Turhan for valuable discussions on the topics covered by this paper. We would like to thank Carsten Lutz for pointing out an unnecessary restriction concerning the default reasoning substrate. Furthermore, thanks go to Omar Nuri for his work on the implementation of the default reasoning substrate. All deficiencies are due to our own faults, of course.

References

[1] F. Baader and P. Hanschke. A scheme for integrating concrete domains into concept languages. In *Twelfth International Conference on Artificial Intelligence, Darling Harbour, Sydney, Australia, Aug. 24-30, 1991*, pages 452–457, August 1991.

[2] F. Baader and B. Hollunder. Embedding defaults into terminological representation systems. *J. Automated Reasoning*, 14:149–180, 1995.

[3] F. Baader and B. Hollunder. Priorities on defaults with prerequisites, and their application in treating specificity in terminological default logic. *J. Automated Reasoning*, 15:41–68, 1995.

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

[5] 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.

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

[7] D. Koller, A. Levy, and A. Pfeffer. P-Classic: A tractable probabilistic description logic. In *Proc. of AAAI 97*, pages 390–397, Providence, Rhode Island, 1997.

[8] C. Lutz and R. Möller. Defined topological relations in description logics. In M.-C. Rousset et al., editor, *Proceedings of the International Workshop on Description Logics, DL'97, Sep. 27-29, 1997, Gif sur Yvette, France*, pages 15–19. Universite Paris-Sud, Paris, September 1997.

[9] R. Möller and M. Wessel. Terminological default reasoning about spatial information: A first step. In D.M. Mark C. Freksa, editor, *Proc. International Conference on Spatial Information Theory, COSIT '99, Stade, Germany*, pages 189–204. Springer-Verlag, 1999.

[10] B. Neumann and C. Schröder. How useful is formal knowledge representation for image interpretation. In E. Guelch K. Torleg, editor, *Photogrammetric Reports Nr. 63, Proceedings of the Joint Workshop of ISPRS WG III/2 and IC WG II/III, August 30 - September 1, Stockholm, Sweden*, Royal Institute of Technology, Stockholm, 1995.

[11] C. Schröder. *Bildinterpretation durch Modellkonstruktion: Eine Theorie zur rechnergestützten Analyse von Bildern*. Dissertation (in German), Universität Hamburg, FB Informatik, 1998.

[12] A.-Y. Turhan. Entwurf und Implementation eines Tableaux-Beweisers für die Beschreibungslogik $\mathcal{ALCRP}(\mathcal{D})$. Studienarbeit (in German), 1998.