Non-Standard Inferences for Knowledge-Based Image Retrieval

Tommaso Di Noia[†], Eugenio Di Sciascio[†], Francesco M. Donini[‡], Francesco di Cugno[†], Eufemia Tinelli[†]

Politecnico di Bari[†]
Via Re David, 200, I-70125, Bari, Italy
{t.dinoia,disciascio,f.dicugno,e.tinelli}@poliba.it

Università della Tuscia[‡]
via San Carlo, 32, I-01100, Viterbo, Italy
donini@unitus.it

Abstract

As more multimedia content is semantically annotated there is the need to fully benefit from the efforts devoted to the annotation. Such benefit is obvious when thinking of standard inference services such as classification and satisfiability, yet we believe there are other, smart, services that can complement basic ones and provide smarter queries on annotated image repositories. In this paper we show and motivate how Concept Abduction and Concept Contraction, recently introduced inference services in Description Logics, can be used for semantic based query and query refinement of annotated images.

1 Introduction

Content-Based Image Retrieval (CBIR) has been in recent years a hot research topic, it has been widely studied and a number of methodologies, techniques and tools, related to image content processing, have been studied for identification and comparison of image features, in order to develop classification and retrieval systems based on – almost– automatic interpretation of image content [1, 10, 13]. The purpose of such investigations was the development of Multimedia Information Systems (MMIS) able to retrieve images coherent –to some extent– with a graphically sketched query or with a reference image, in a query by example fashion. To achieve this goal typical features considered were color, texture, shape, and position of objects within images, extracted from database images and used as query terms.

As a matter of fact classification, indexing and retrieval of images, based on their actual content, calls for a semantic interpretation of images.

So, the computation of visual features is a surrogate of semantic interpretation, to be used for identifying images similar to a given query image, using features similarity functions. In fact the problem of finding images composing a conceptual query describing users' needs becomes a similarity search in some feature space, which maintains a semantic gap between users' needs and the possibility to model it.

To overcome limitations of such an approach, also motivated by the growing interest in the Semantic Web initiative, efforts have been devoted to describe images with techniques exploiting automated knowledge representation [23, 9, 21]. The need to annotate multimedia content to ease machine interpretation is obviously not new, and also standards as MPEG-7 [17] and MPEG-21 [16] tried to tackle the need to describe in some unambiguous way multimedia content. Nevertheless such standards basically aimed at providing structured description of low-level features, and descriptions are not endowed of well-defined associated semantics, which prevents automated reasoning on descriptions.

Ontologies founded on logical languages are hence emerging to describe images, and images high level clues, within various domains, see *e.g.*, [22, 19, 12].

Approaches to fully automate the annotation stage, *i.e.*, from actual image content to its semantic description are far from having reached satisfactory results, but encouraging results have been obtained when images belong to a restricted domain of interest, see *e.g.*, [3]. It is also noteworthy, that proposals have been made also based on manual and collaborative annotation as in Photostuff [11].

The advantage of having images annotated using languages with well defined syntax and semantics should be the possibility to carry out automated reasoning tasks on them, and the possibility to compose conceptual queries. Nevertheless we believe the effort of semantic annotation has to be compensated with services smarter than plain satisfiability and classification tasks.

In this contribution, borrowing from ongoing activity in other Semantic Web research fields that share as central the notion of semantically annotated resource, we present an innovative Description Logics framework for semantic-based image retrieval that introduces non-standard inferences for conceptual query and query refinement.

Description Logics have been already proposed in the framework of content-based image retrieval, see, *e.g.*, [7, 15, 24, 8, 14]. The novelty of our proposal lies in the fact that our approach allows both execution of reasoning tasks using well-known reasoning services and using new, non standard inferences.

Such services, named Concept Abduction and Concept Contraction have been recently defined in the framework of Description Logics and allow for semantic-based ranking and intelligent conceptual query refinement [4, 5].

Our setting adopts an ontology modeled using a subset of OWL-DL [20] and allows queries by structured textual description and/or by sketch, *i.e.*, by composition of iconic descriptions corresponding to elements of the ontology.

The remaining of this paper is structured as follows: next section summarizes basics of Description Logics; then we recall above mentioned non-standard inference services and motivate their utility in semantic-based image retrieval; Section 4 presents our framework. Conclusions and ongoing work close the paper.

2 Description Logics basics

Our approach adopts Description Logics (DLs) [2] as a formal framework. DLs are a family of logic formalisms devised to model complex hierarchical structures and provide reasoning services on these structures. DLs basic syntax elements are *concept* names, such as `Vehicle`, `Car`, `Object`, and *role* names, such as `is_main_subject`. Intuitively, concepts stand for sets of objects, and roles link objects in different concepts. Individuals are used for special named elements belonging to concepts.

Formally, a semantic *interpretation* is a pair $\mathcal{I} = (\Delta, \cdot^{\mathcal{I}})$, which consists of the *domain* Δ and the *interpretation function* $\cdot^{\mathcal{I}}$, which maps every concept to a subset of Δ, every role to a subset of $\Delta \times \Delta$, and every individual to an element of Δ. The Unique Name Assumption is usually made, *i.e.*, different individuals are mapped to different elements of Δ: $a^{\mathcal{I}} \neq b^{\mathcal{I}}$ for individuals $a \neq b$. Basic elements can be combined using *constructors* to form concept and role *expressions*; each DL has its own set of constructors. Every DL is endowed of *conjunction* of concepts, usually denoted as $\sqcap$; some DL include also disjunction $\sqcup$ and complement $\neg$ to close concept expressions under boolean operations. Roles can be combined with concepts using *existential role quantification*, and *universal role quantification*. Other constructs may involve counting, as number restrictions. Many other constructs can be defined, increasing the expressive power of the DL, up to n-ary relations. Expressions are given a semantics by defining the interpretation function over each construct. For example, concept conjunction is interpreted as set intersection: $(C \sqcap D)^{\mathcal{I}} = C^{\mathcal{I}} \cap D^{\mathcal{I}}$, and also the other boolean connectives $\sqcup$ and $\neg$, when present, are given the usual set-theoretic interpretation of union and complement. The interpretation of constructs involving quantification on roles needs to make domain elements explicit.

Concept expressions can be used as axioms in *inclusion assertions*, and *definitions*, which impose restrictions on possible interpretations according to the knowledge elicited for a given domain. Definitions are useful to give a meaningful name to particular combinations, Historically, sets of such inclusions are called TBox (Terminological Box), which basically amounts to an ontology. In simple DLs, only a concept name can appear on the left-hand side of an inclusion. The semantics of inclusions and definitions is based on set containment: an interpretation $\mathcal{I}$ satisfies an inclusion $C \sqsubseteq D$ if $C^{\mathcal{I}} \subseteq D^{\mathcal{I}}$, and it satisfies a definition $C \equiv D$ when $C^{\mathcal{I}} = D^{\mathcal{I}}$. A *model* of a TBox $\mathcal{T}$ is an interpretation satisfying all inclusions and definitions of $\mathcal{T}$.

DL-based engines usually provide at least two basic reasoning services:

1. *Concept Satisfiability.* $C \equiv_{\mathcal{T}} \bot$: given a TBox $\mathcal{T}$ and a concept C, does there exist at least one model of $\mathcal{T}$ assigning a non-empty extension to C?

2. *Subsumption.* $D \sqsubseteq_{\mathcal{T}} C$: given a TBox $\mathcal{T}$ and two concepts C and D, is C more general than D in any model of $\mathcal{T}$?

Informally, satisfiability accounts for the internal coherency of the description of a concept (no contradictory properties are present), and subsumption accounts for the more general/more specific relation among concepts, that forms the basis of a taxonomy. Such services provide basically a yes/no result. As it is typical of generic resource retrieval this might be not enough when such services have to provide query answering. For example, in content-based image retrieval we believe that at least a ranking function should exist, and this should be –if possible– based on logical criteria.

In the next section we introduce some novel non-standard inference services, hihglighting thier behavior and motivating their usefulness in semantic-based image retrieval in computing a semantic-based score for a query.

3 Why Non-standard Inference Services?

We now informally introduce previously cited novel inference services [4, 5] and sketch their rationale in semantic-based image retrieval.

Abduction is a form of non-monotonic reasoning, modeling commonsense reasoning, usually aimed at finding an explanation for some given symptoms or manifestations. *Concept Abduction* captures the reasoning mechanism – namely, making hypotheses – involved when some of the features required in a description are not specified within available image descriptions.

Contraction is the first step in belief revision. *Concept Contraction* captures the possibility to relax some of the features requested in a query Q when they are in conflict with those of an available image description Im, that is, using DLs syntax, when $Q \sqcap Im$ is an unsatisfiable concept. Formally, Concept Contraction can be defined as follows:

Definition 1 *Let $\mathcal{L}$ be a DL, Im, Q, be two concepts in $\mathcal{L}$, and $\mathcal{T}$ be a set of axioms in $\mathcal{L}$, where both Im and Q are satisfiable in $\mathcal{T}$. A* Concept Contraction Problem *(CCP), identified by $\langle \mathcal{L}, Q, Im, \mathcal{T} \rangle$, is finding a pair of concepts $\langle G, K \rangle \in \mathcal{L} \times \mathcal{L}$ such that $\mathcal{T} \models Q \equiv G \sqcap K$, and $K \sqcap Im$ is satisfiable in $\mathcal{T}$. We call K a* contraction *of Q according to Im and $\mathcal{T}$.*

Obviously, there is always a trivial solution $\langle G, K \rangle = \langle Q, \top \rangle$ to a CCP. This solution corresponds to the most drastic contraction, that gives up everything of Q. Concept contraction can be hence said to extend satisfiability.

In an analogous way we can formally define Concept Abduction:

Definition 2 *Let $\mathcal{L}$ be a DL, Im, Q, be two concepts in $\mathcal{L}$, and $\mathcal{T}$ be a set of axioms in $\mathcal{L}$, where both Im and Q are satisfiable in $\mathcal{T}$. A* Concept Abduction Problem *(CAP), identified by $\langle \mathcal{L}, Q, Im, \mathcal{T} \rangle$, is finding a concept $H \in \mathcal{L}$ such that $\mathcal{T} \models Im \sqcap H \sqsubseteq Q$, and moreover $Im \sqcap H$ is satisfiable in $\mathcal{T}$. We call H a* hypothesis *about Im according to Q and $\mathcal{T}$.*

As Concept Contraction extends satisfiability, Concept Abduction extends subsumption. We use such inferences in our approach both to provide interaction and refinement in the query / retrieval process and to rank images based on their semantic-based distance from the query. We motivate and clarify the rationale of our approach with the aid of a simple example.

Let us consider a simple conceptual query to the system: **Q:** *I'm looking for images showing a skyscraper, the sun and the sea as background.*
Let us suppose that, among various images, there are three in the knowledge base whose –simplified– descriptions are as follows:

Im1: *Landscape with a Person and the sun.*

Im2: *Person near a building and a garden as background.*

Im3: *Building in background with a garden, and the sun.*

Expressing both the query and image descriptions with respect to a Description Logic we have

Q: $Skyscraper \sqcap Sun \sqcap \exists background \sqcap \forall background.Sea$

Im1: $Person \sqcap Sun$

Im2: $Building \sqcap Person \sqcap \exists background \sqcap \forall background.Garden$

Im3: $Building \sqcap Sun \sqcap \exists background \sqcap \forall background.Sea$

.

Now suppose to have a simple ontology containing only two axioms, *i.e.*, $\mathcal{T} = \{Skyscraper \sqsubseteq Building; Sea \sqsubseteq \neg Garden\}$.

With respect to the previous ontology we can evaluate the match degree between **Q** and each image description as well as some semantic-based explanations, exploiting the relations among concepts as formalized within the ontology axioms, for the match score.

Notice that the approach is based on an Open World Assumption (OWA). With respect to an OWA, the absence, in an image description, of a feature requested in the query is not interpreted as a constraint of absence, or, in other words, as if the feature was negated. The OWA allows to handle incomplete information. It is hence possible to deal with underspecified descriptions.

SIM(Im1,Q): **Im1** can potentially satisfy **Q**. No feature in **Im1** is conflicting with those represented in **Q**, but characteristics depicted in **Im1** do not completely fulfill **Q**. Using DLs syntax we say $Q \sqcap Im1 \not\equiv \bot$ and $Im1 \not\sqsubseteq Q$. Explanation hypothesis are needed for the non-subsumption relation. Solving the related *Concept Abduction Problem*, a solution is $H_{Im1} = Skyscraper \sqcap \exists background \sqcap \forall background.Sea$. Considering **Im1**, **Q** and the ontology $\mathcal{T}$, a semantic-based score for H_{Im1} has to be computed.

SIM(Im2,Q): Features in **Im2** are conflicting with what is requested in **Q** *i.e.*, $Im2 \sqcap Q \equiv \bot$. A revision of the query may be needed if we are interested in images similar to **Im2**. Solving the related *Concept Contraction Problem* a solution is $Q = G_Q \sqcap K_Q$, where $G_Q = \forall background.Sea$ and $K_Q = Skyscraper \sqcap Sun \sqcap \exists background$. Now the relation $K_Q \sqcap Im2 \not\equiv \bot$ holds, but we have $Im2 \not\sqsubseteq K_Q$. Again, solving the related Concept Abduction Problem a solution $H_{Im2} = Skyscraper \sqcap Sun$ results.

The ranking function, computing the overall score, in this case must consider G_Q, K_Q, H_{Im2} and/or the relative scores for each of them.

SIM(Im3,Q): $Q \sqcap Im3 \not\equiv \bot$ and $Im3 \not\sqsubseteq Q$. $H_{Im3} = Skyscraper$.

Notice that because of the axiom $Skyscraper \sqsubseteq Building$ in $\mathcal{T}$, the concept $Skyscraper$ in H_{Im2} must be evaluated differently from $Skyscraper$ in H_{Im2}. To cope with this property, both monotonicity and anti-monotonicity behavior [6, 18] has to be respected by the penalty function computing the overall score.

With respect to the above considerations, ranking $Im1$, $Im2$ and $Im3$ w.r.t. Q, the first image in the list is $Im3$. Notice also that if there was one or more images ImX, such that $ImX \sqsubseteq Q$, then ImX would be obviously the best result for Q.

4 Proposed Framework

In the former example we used the whole annotation associated to an image to define elements to be used in the evaluation of both a semantic based match degree and semantic based explanation in case of non-exact matches. Obviously, the explanation services can be used in query reformulation and refinement. This is the general approach one can adopt when nothing is known about the nature, the context and/or the structure of both the annotation and the annotated resource.

We would like to point out we are not expressing here an explicit ranking function. Various ranking functions can be in fact defined, having as arguments concepts provided by G, K and H, which can be used to determine the match degree. We currently have ongoing experiments with various penalty functions aimed at finding those that best model human common sense.

Yet, in image retrieval the structure of the annotation, at an abstract level, can be exploited in order to perform a matching process which takes into account the nature of the annotated resources – the images – and the annotation itself.

To this aim we adopt a simple ontology model for images semantic-annotation. We started modeling a core ontology, aimed at capturing the concepts that describe –and are used to describe– images.

In fact, before building this core image ontology we asked some volunteers to annotate, with natural language sentences, several pictures related to different context and subsequently to formulate queries on images in the collection. Looking at both the annotation and the queries, we observed that they followed a typical pattern. Typical annotations and queries were in the form "`A red car on the` **`left side`** `of the picture and a box in the` **`background`**" or "`Birthday picture with a cake` **`in the center`** `and people` **`all around`**", and so on. Content information are related to their position within the picture and position-based patterns. While the first are strictly related to the picture context, the second are context-independent and can be used for any image.

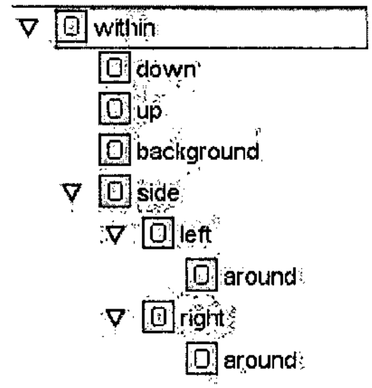

Figure 1. Simplified role hierarchy for the core image ontology

Based on the above observations, we modeled a core ontology for images in which only spatial relations are represented. Such ontology is strongly property oriented, that is the notions of `background`, `up`, `down`, `on the left side`, `all around` and so on, are modeled as properties – from now on R^i_{core} – whose restrictions describe the content of a particular region in the picture. Both the image annotations and the queries are expressed as a conjunction of quantified properties. For example, the annotation for `A red car on the` **`left side`** `of the picture and a box in the` **`background`** can be roughly expressed as

$\exists \texttt{left} \sqcap \forall \texttt{left}.(Car \sqcap \exists hasColor \sqcap \forall hasColor.Red) \sqcap$
$\exists \texttt{background} \sqcap \forall \texttt{background}.Box$

The above introduced properties, R^i_{core}, are related with each other. For instance the semantics of `around` is related to the ones of `left` and `right`. If something is all around the picture, then it is on the `left` side and on the `right` one. Hence, the image ontology contains the set $\mathcal{R} = \{R^i_{core}\}$ and axioms on such properties stating their mutual relation. In Figure 1 the role hierarchy is depicted representing the core ontology for our approach.

The obvious question is: if the image ontology contains only spatial relations, where are the content information? How to model and describe the content of the picture?

We started from the observation that usually, the content is strongly related to the picture knowledge domain and we developed different domain ontologies for different knowledge domains. Currently we modeled ontologies for the landscapes, groups and touristic places domains. Then, using the `<owl:imports/>` OWL TAG, not only we make possible to use different sets of knowledge domains (*i.e.*, groups and touristic places, or landscapes and touristic places) but we allow to introduce the content within an im-

$\forall$background .(*b_annot*)$\sqcap$
$\forall$around.(*a_annot*).

Figure 2. A semantic-based annotation Im for an image

$\forall$within.(*w_query*)$\sqcap$
$\forall$left.(*lx_query*)$\sqcap$
$\forall$right.(*rx_query*).

Figure 3. A semantic-based query Q for an image

age annotation. Using <owl:imports/> we import the domain ontologies within the image core ontology, reusing such ontologies for image content description.

Since the matching process (see Section 3) involves basically the image content, domain ontologies were developed using the $\mathcal{ALN}$ subset of OWL-DL, for which algorithms to solve Concept Abduction and Concept Contraction problems exist [4, 5].

Importing in the core image ontology the domain ontologies, it is possible to annotate images with reference both to spatial relations and actual image content. We recall that both annotations and queries are in conjunctive form with respect to the core image ontology, as sketched in Figure 2 and Figure 3.

Now we describe how the semantic-based image retrieval process is performed, with the aid of the above sketched annotation and query.

In Figure 4 is depicted a snapshot of the prototype user interface for both annotation and querying.

1. The user selects the content domain (or a set of content domains). This selection corresponds to the identification of the domain ontology (or a set of domain ontologies) to be imported in the core image ontology.

2. The user selects the properties she is interested in, with respect to the core image ontology, and composes restriction of such properties (*i.e.*, the content of selected image regions). With respect to Figure 3, in Q the user first selects the property `within` and then composes its restriction – *w_query* – using domain ontologies, then `left` with its *left_side_request* and finally `right` and the related *rx_query*.

3. For each R^i_{core} in the query, the corresponding restriction is selected both in the query and in the image annotation and an extended matchmaking process is performed. In the image annotation, also the restriction related to properties R^j_{core} such that $R^j_{core} \sqsubseteq R^i_{core}$ are selected and put in conjunction with the one of R^i_{core}. Considering the query Q in Figure 3 and the annotation Im in Figure 2 the following steps are executed:

   a. *w_query* is selected as the restriction of `within` in Q.

   b. in Im, due to the sub-property relations in the image core ontology, the restriction of both `background` – *b_annot* – and `around`– *a_annot* – are selected.

   c. an extended matchmaking process is performed between *a_annot*$\sqcap$*b_annot* and *w_query* returning $\langle G_{\texttt{within}}, K_{\texttt{within}}, H_{\texttt{within}} \rangle$.

   d. *lx_query* is selected as the restriction of `left` in Q.

   e. in $Im$the restriction of `around`– *a_annot* – is selected.

   f. an extended matchmaking process is performed between *a_annot* and *lx_query* returning $\langle G_{\texttt{left}}, K_{\texttt{left}}, H_{\texttt{left}} \rangle$.

   g. a similar process is performed for *rx_query* as the restriction of `right`in Q returning $\langle G_{\texttt{right}}, K_{\texttt{right}}, H_{\texttt{right}} \rangle$.

   Notice that for R^i_{core} in Q, if neither R^i_{core} nor any $R^j_{core} \sqsubseteq R^i_{core}$ are in Im, then $\top$ is considered as the restriction of R^i_{core} in Im.

4. a score is computed by means of a penalty function using $\langle G_{\texttt{within}}, K_{\texttt{within}}, H_{\texttt{within}} \rangle$, $\langle G_{\texttt{left}}, K_{\texttt{left}}, H_{\texttt{left}} \rangle$ and $\langle G_{\texttt{right}}, K_{\texttt{right}}, H_{\texttt{right}} \rangle$.

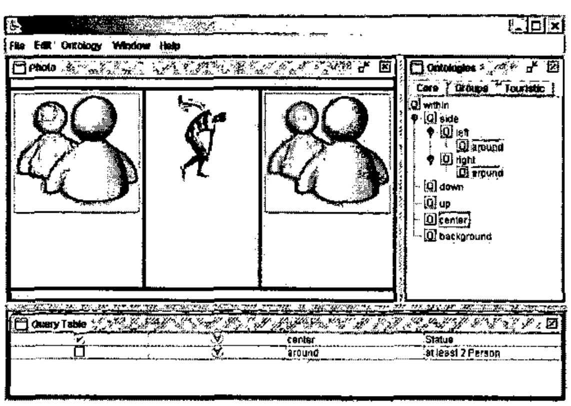

Figure 4. The prototype user interface

In the query the user is also able to express a strict constraint, that is the restriction of a property R^i_{core} in the query must not be in conflict with the corresponding restriction in the retrieved image. The above user specification is modeled in our framework as a condition on $G_{R^i_{core}}$. If the user states R^i_{core} as a strict property, then $G_{R^i_{core}} \equiv \top$ must occur. No Concept Contraction has to be performed on the strict property restriction. For instance, with respect to the above example, if the user states that $R^i_{core} = \texttt{left}$ is a strict constraint in the query, then the result of the step 3.f. **must** be $\langle \top, lx\_query, H_{\texttt{left}} \rangle$ otherwise the image is not selected and is discarded. That means that if the `left`specifications are not in conflict with the related ones in the image annotation then no contraction is needed – $G_{\texttt{left}} = \top$, nothing as to be contracted, and $K_{\texttt{left}} = lx\_query$, all the `left`specification is kept – and then the strict constraint is respected, making the image selected and presented within the ranked result list to the user.

Acknowledgments

The authors wish to acknowledge partial support of project PON CNOSSO. They also wish to thank Vito Nicola Oro for useful discussions and software implementations.

5 Conclusion

In this paper we have presented a framework for semantic-based image retrieval, which exploits a core ontology that comprises only spatial relationships in the image, while specific domain ontologies are imported for different domains. We take advantage of recently devised non standard inference services in Description Logics. Such services, namely Concept abduction and Concpet Contraction allow to determine a semantic-based ranking of annotated images with respect to a conceptual query, and provide explanation of retrieval results, which can be used for query refinement and reformulation. Based on our framework we have built a simple prototype, which includes a semi-automatic image annotation tool, a query composition interface, and exploits the above briefly described inference mechanisms to provide logic-based results ranking and query refinement. Currently, experiments are ongoing using various penalty functions, aimed at finding those that best model human common sense.

References

[1] S. Antani, R. Kasturi, and R. Jain. A survey on the use of pattern recognition methods for abstraction, indexing and retrieval of images and video. *Pattern Recognition*, 35(4):945–965, 2002.

[2] F. Baader, D. Calvanese, D. Mc Guinness, D. Nardi, and P. Patel-Schneider, editors. *The Description Logic Handbook*. Cambridge University Press, 2002.

[3] S. Bloehdorn, K. Petridis, C. Saathoff, N. Simou, V. Tzouvaras, Y. Avrithis, S. Handschuh, Y. Kompatsiaris, S. Staab, and M. Strintzis. Semantic annotation of images and videos for multimedia analysis. In *The Semantic Web: Research and Applications (proc. of ESWC 2005)*, volume 3532 of *LNCS*, pages 592–607. 2005.

[4] S. Colucci, T. Di Noia, E. Di Sciascio, F.M. Donini, and M. Mongiello. Concept Abduction and Contraction in Description Logics. In *Proceedings of the 16th International Workshop on Description Logics (DL'03)*, volume 81 of *CEUR Workshop Proceedings*, 2003.

[5] S. Colucci, T. Di Noia, E. Di Sciascio, F.M. Donini, and M. Mongiello. A Uniform Tableaux-Based Method for Concept Abduction and Contraction in Description Logics. In *Proceedings of the 16th European Conference on Artificial Intelligence (ECAI '04))*, pages 975–976. IOS Press, 2004.

[6] T. Di Noia, E. Di Sciascio, F.M. Donini, and M. Mongiello. A system for Principled Matchmaking in an Electronic Marketplace. *International Journal of Electronic Commerce*, 8(4):9–37, 2004.

[7] E. Di Sciascio, F.M. Donini, and M. Mongiello. Structured knowledge representation for image retrieval. *Journal of Artificial Intelligence Research*, 16:209–257, 2002.

[8] E. Di Sciascio, F.M. Donini, and M. Mongiello. A logic for SVG documents query and retrieval. *Multimedia Tools and Applications*, 24(2):125–153, 2004.

[9] Joost Geurts, Jacco van Ossenbrugen, and Lynda Hardman. Requirements for practical multimedia annotation. In *Multimedia and the Semantic Web Workshop at ESWC 2005*, 2005.

[10] T. Gevers and A.W.M. Smeulders. Content-based image retrieval: An overview. In G. Medioni and S. B. Kang, editors, *Emerging Topics in Computer Vision*. Prentice Hall, 2004.

[11] Christian Halaschek-Wiener, Andrew Schain, Jennifer Golbeck, Michael Grove, Bijan Parsia, and Jim Hendler. A flexible approach for managing digital images on the semantic web. 5th International Workshop on Knowledge Markup and Semantic Annotation, 2005.

[12] S. Hubner, R. Spittel, U. Visser, and T. J. Vogele. Ontology-based search for interactive digital maps. *IEEE Intelligent Systems*, 19(3):80–86, 2004.

[13] M. L. Kherfi, D. Ziou, and A. Bernardi. Image Retrieval from the World Wide Web: Issues, Techniques, and Systems. *ACM Computing Surveys*, 36(1):35–67, 2004.

[14] Nicolas Maillot, Monique Thonnat, and Alain Boucher. Towards ontology based cognitive vision. *Machine Vision and Applications*, 16(1):33–40, 2004.

[15] Carlo Meghini, Fabrizio Sebastiani, and Umberto Straccia. A model of multimedia information retrieval. *Journal of ACM*, 48(5):909–970, 2001.

[16] MPEG-21. MPEG-21 Overview. http://www.chiariglione.org/mpeg/standards/mpeg-21/mpeg-21.htm.

[17] MPEG-7. MPEG-7 Overview. http://www.chiariglione.org/mpeg/standards/mpeg-7/mpeg-7.htm.

[18] T. Di Noia, E. Di Sciascio, and F.M. Donini. Extending Semantic-Based Matchmaking via Concept Abduction and Contraction. In *EKAW 2004, Engineering Knowledge in the Age of the Semantic Web*, volume 3257 of *LNAI*, pages 307–320. 2004.

[19] N. F. Noy, D. L. Rubin, and M. A. Musen. Making biomedical ontologies and ontology repositories work. *IEEE Intelligent Systems*, 19(6):78–81, 2004.

[20] OWL. www.w3.org/TR/owl-features/.

[21] K. Petridis, I. Kompatsiaris, M. G. Strintzis, S. Bloehdorn, S. Handschuh, S. Staab, and N. Simou. Knowledge representation for semantic multimedia content analysis and reasoning. In *1 st European Workshop on the Integration of Knowledge, Semantics and Digital Media Technology*, 2004.

[22] A. Schreiber, B. Dubbeldam, J. Wielemaker, and B. Wielinga. Ontology-based photo annotation. *IEEE Intelligent Systems*, 16(3):66–74, 2001.

[23] Giorgos Stoilos, Giorgos Stamou, Vassilis Tzouvaras, Jeff Pan, and Ian Horrocks. A fuzzy description logic for multimedia knowledge representation. In *Multimedia and the Semantic Web Workshop at ESWC 2005*, 2005.

[24] Umberto Straccia. Reasoning within fuzzy description logics. *Journal of Artificial Intelligence Research*, 14:137–166, 2001.