

# Multimedia Information Extraction in Ontology-based Semantic Annotation of Product Catalogues

Roberto Bartolini, Emiliano Giovannetti, Simone Marchi, and Simonetta Montemagni, *ILC-CNR*

Claudio Andreatta and Roberto Brunelli, *ITC-irst*

Rodolfo Stecher, *Fraunhofer IPSI*

Paolo Bouquet, *DIT-University of Trento*

**Abstract**—The demand for efficient methods for extracting knowledge from multimedia content has led to a growing research community investigating the convergence of multimedia and knowledge technologies. In this paper we describe a methodology for extracting multimedia information from product catalogues empowered by the synergetic use and extension of a domain ontology. The methodology was implemented in the Trade Fair Advanced Semantic Annotation Pipeline of the VIKE-framework.

**Index Terms**—Semantic Web Technologies, Ontology Creation, Ontology Extraction, Ontology Evolution, Semantic Annotation of Multimedia Content

## I. INTRODUCTION

**E**FFECTIVE acquisition, organization, processing, use and sharing of the knowledge embedded in textual and multimedia content play a major role for competitiveness in the modern information society and for the emerging knowledge economy. However, this wealth of knowledge implicitly conveyed in the vast amount of available digital content is nowadays only accessible, if considerable manual effort has been invested into its interpretation and semantic annotation, which is possible only for a small fraction of the available content.

Manuscript received October 27, 2006.

Roberto Bartolini, Emiliano Giovannetti, Simone Marchi, and Simonetta Montemagni work at the Istituto di Linguistica Computazionale (ILC-CNR) in Pisa, via Moruzzi 1, 56124, Pisa, Italy (emails: {roberto.bartolini, emiliano.giovannetti, simone.marchi, simonetta.montemagni}@ilc.cnr.it)

Claudio Andreatta and Roberto Brunelli work at the Istituto per la Ricerca Scientifica e Tecnologica ITC-irst in Trento, via Sommarive 18, 38050, Trento, Italy (emails: {andreatta, brunelli}@itc.it)

Rodolfo Stecher works at the L3S Research Center, Appelstrasse 9a, Hannover, Germany (email:stecher@l3s.de)

Paolo Bouquet works at the Department of Information and Communication Technologies of the University of Trento, via Sommarive 14, 38050 Trento, Italy (email: paolo.bouquet@unitn.it)

The field of semi-automatic information extraction from multimedia corpora is central for overcoming the so-called “knowledge acquisition bottleneck”. Multimedia sources of information, such as product catalogues, contain text (captions) and images (pictures of the products) thus requiring information extraction approaches combining several different techniques, ranging from Natural Language Processing to Image Analysis and Understanding. In our approach we have three main aspects to consider: 1) the information extraction per se, 2) the ontology, its use and creation, and 3) the usage of the ontology in the information extraction process and the synergy between different kinds of extraction processes.

The development of adequate ontologies itself is one of the knowledge acquisition bottlenecks: the use of (semi-) automatic tools for semantic information extraction from multimedia corpora is very promising but, to be efficiently exploited, must have access to a formal representation of a given domain, i.e., an ontology. We support the ontology creation process in two different and complementary ways, ontology learning and reuse of existing ontologies. The ontology learning approach takes advantage of the results of the extraction to enrich the ontology, and the reuse support provides methods and tools to re-use already existing ontologies which capture the target domain under a similar modelling perspective as the one of interest for the extraction task. This (apparent) vicious circle (between the need of having the domain represented in the ontology for an extraction process and the enrichment of the ontology based on the results obtained from the extraction) can be turned to a virtuous circle if the necessary conditions are set to let the evolving ontology and the information extraction tool interact in a synergetic way.

After a brief introduction to the Vike-Framework the general methodology is described in section III, including specific details about the four different components of the system pipeline. Some conclusions will be presented in section

IV.

## II. THE VIKEF-FRAMEWORK

The methodology we present is developed inside the Vikef project (Virtual Information and Knowledge Environment Framework, IST-2002-507173 - <http://www.vikef.net/>), which creates an advanced software framework for enabling the integrated development of semantic-based Information, Content, and Knowledge (ICK) management systems. Apart from the scientific and academic interest related to these fields of research we have also registered a growing need from industrial parties for automated knowledge elicitation tools to be applied to their commercial resources, such as product catalogues.

VIKEF bridges the gap between the partly implicit knowledge and information conveyed in scientific and business content resources (e.g. text, speech, images) and the explicit representation of knowledge required for a targeted and effective access, dissemination, sharing, use, and annotation of ICK resources by scientific and business communities and their information- and knowledge-based work processes.

R&D within VIKEF builds on and significantly extends the current Semantic Web efforts by addressing crucial operationalisation and application challenges in building up real-world semantically enriched virtual information and knowledge environments.

## III. THE METHODOLOGY

The task of (semi) automatically annotating content objects with semantic information requires a multi-phased process, where multimedia entities discovered within a content object are coupled with domain knowledge represented by an ontology. For effective semantic annotation support, linguistic, image-related and knowledge representation aspects, approaches, and formats, have to be combined in a synergetic way. The proposed methodology can be presented as a pipeline (together with the employed representation formats within the pipeline), which supports semantic annotation in a flexible and pragmatic way.

The pipeline has been implemented as a prototype developed as part of the VIKEF project and evaluated for content from the Trade Fair domain.

The main components of the pipeline are four, and can be functionally summarized in this way:

- A) Annotation of text – the ontology based semantic annotation of the textual part of the catalogue;
- B) Annotation of images – the ontology based semantic annotation of the images appearing in the catalogue;
- C) Elicitation and refinement – to make information extracted by the annotation component machine-understandable and to enrich the ontology for further annotations;
- D) Reuse of existing ontologies – to support the ontology

creation and refinement by exploiting existing ontologies.

The approach has been thought and implemented to provide the possibility of triggering a “virtuous circle”: once information extracted in the annotation steps are integrated inside the ontology, the whole process can be restarted, thus allowing the textual and image annotators to exploit the novel information added to the ontology during the previous run.

### A. Annotation of Text

Semantic annotation of content is a crucial task (probably the most important) for processing documents to be accessed inside the Semantic Web. To semantically annotate a text it is necessary to develop (semi) automatic Information Extraction techniques capable of overcoming the so-called “knowledge acquisition bottleneck” typical of Semantic Web related applications.

Semantic annotation of product catalogues poses different challenges at different levels. Concerning the textual part, relative to product descriptions, catalogues do not contain linguistically sound text: very often, sentences are constituted by strictly nominal descriptions, thus discouraging the recourse to traditional NLP techniques. On the other hand, product descriptions appear as semi-structured texts where product features appear in a fixed (or at least regular) order. Semantic annotation of product catalogues appears therefore as a complex task requiring the combination of different types of techniques. Previous works about product information semantic annotations are quite scarce: the two main works

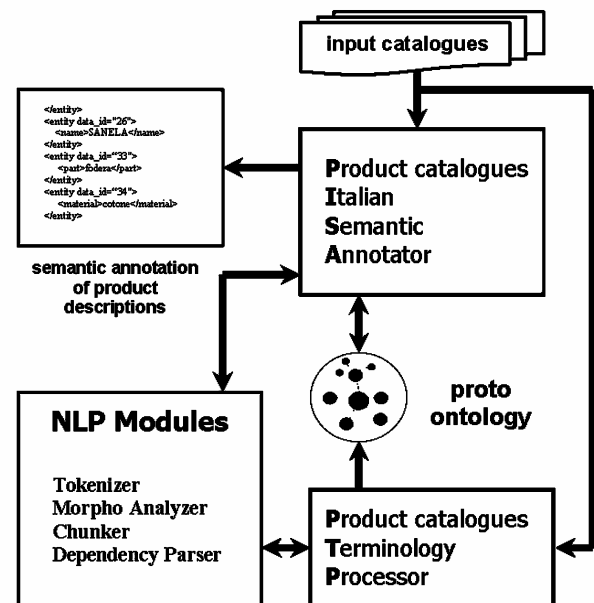


Fig. 1: The general architecture of the annotation system, producing a semantic annotation of product descriptions extracted from the input catalogue.

being the European project CROSSMARC [1] and the Czech national project Rainbow [2]. The project CROSSMARC aims to the electronic-retail product comparison, using a

combination of language engineering, machine learning and user modelling, where a domain ontology is used as “semantic glue” to link together the various analysis modules.

Within the Rainbow project, a multi-layered ontology has been defined, to integrate the more abstract aspects of the domain (domain-neutral), relative to web-sites in general, with the more specific ones (domain-dependent) relative to concepts found in sites of small organization offering products or services. Concerning the information extraction task, Rainbow makes use of lexical indicators and, depending on the document to analyze, applies HTML-centred or free-text-centred extractors, in the latter case using shallow parsing techniques.

The hybrid methodology we propose (which has been applied to Italian product catalogues belonging to the furniture domain) makes use of two different approaches: first, pattern matching techniques are resorted to for isolating individual product descriptions within the textual flow and for identifying their basic building blocks (e.g. the product *name*, its *price* as well as its natural language description). Then, for each identified product, the natural language description is processed by a battery of NLP tools ([3] [4]) in charge of identifying relevant entities (e.g. *colour*, *material*, *parts* of a given product) and the relations holding between them (e.g. *part\_of*, *colour\_of* which can be referred either to the product itself or to individual parts).

The architecture in Fig. 1 includes two main components, the Product catalogue Italian Semantic Annotator (PISA), and the Product catalogues Terminology Processor (PTP), both exploiting the battery of NLP modules, the former to

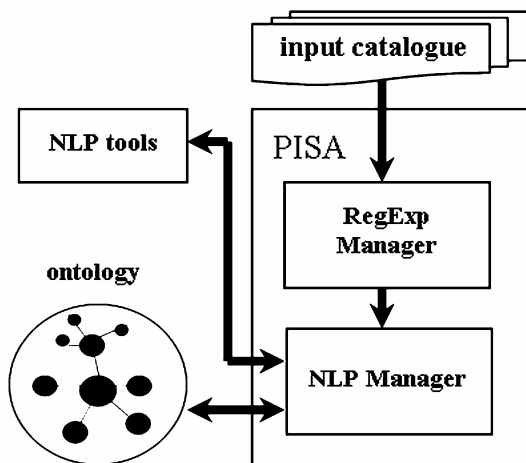


Fig. 2: The general architecture of PISA component, including the RegExp Manager committed for the pattern matching step, and the NLP Manager in charge for the analysis of the “free text” part of the product description.

linguistically analyze the free text part of the product descriptions, the latter to obtain the *TermBank* [5] upon which a simple application ontology can be constructed, to be exploited for disambiguation. This “proto”-ontology forms the

terminological basis for the development of the final (project) ontology, to be populated using the information derived from the multimedia semantic annotation tasks.

Once this proto-ontology has been constructed using the PTP module, it is possible to run the PISA component (Fig. 2): each product description is firstly extracted by pattern matching starting from a set of regular expressions, each one matching a particular product structure. Once a description has been isolated, some of its components, interpreted on the basis of particular “groups” of the matching regular expressions (“name”, “type”, “product id”, etc.) can be detected.

The remaining “free text” part of the description is then processed by the NLP Manager module, which is able to access the NLP tools and the ontology, the former to linguistically analyze the text, the latter to resolve possible syntactical ambiguities found during the analysis.

Consider the example in Fig. 3 which is relative to the annotation of the description in the box about a given product. Through pattern matching it is possible to extract its name (“Sanela”), the type (cushion), the price (€12,95), its dimensions (40 cm of width and 60 cm of length) and the product unique identifier (900.582.56), as well as the relations

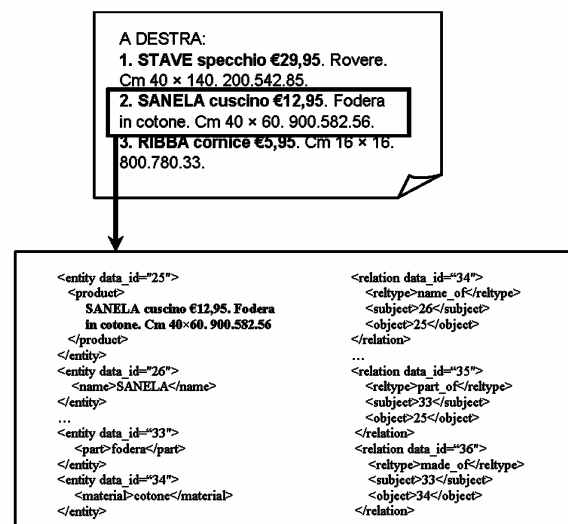


Fig. 3: Example of product semantic annotation, including entities and relations among entities.

between this information and the product itself (i.e. *name\_of*, *price\_of*, etc). The natural language description identified at this stage is then passed to the NLP Manager which is in charge of acquiring, with the support of the application ontology, further information about the product: in this example, the system detects a part (“cover”) and a material (“cotton”), as well as a relation holding between them (“made\_of”). The box below contains a snippet of the (XML-style) final annotated product description, where some of the extracted features of the product are listed, including the fact, for instance, that *the cover of the cushion is made of cotton*.

A preliminary evaluation of the system has been done analyzing the semantic annotation of the “IKEA 2006” italian

furniture catalogue. There are two main levels of evaluation that have been taken into consideration, relative to the two main components of the system: PTP and PISA. Due to the lack of a “golden standard” furniture ontology to compare to the one obtained with the help of the PTP component, a “task based” evaluation technique has been adopted, where the coverage of the ontology has been indirectly evaluated on the basis of the quality of the obtained annotation.

To sum up the results of the preliminary evaluation we can say the ontology is able to detect approximately the 70% of the terms appearing in the free text description of the extracted product and put them in the correct relation, thus scoring a 0,7 of recall. On the other hand, considering that only terms included in the ontology can be extracted, the system scores a precision of 1. Concerning the disambiguation functionality, our analysis has proved that every time two terms are correctly detected and recognized, the disambiguation works. After all, we can assert that the quality of the linguistic analysis is strongly related to the ontology coverage.

From the pattern matching point of view, the system has scored a 0,9 of precision and 0,8 of recall, extracting 800 products out of 1000 applying 9 different regular expressions.

Concerning the task of semantic text annotation future directions of research include, on the one hand the application of the presented technique to different product catalogues (not necessarily in Italian) and, on the other hand, the evolution of the methodology by the integration of a more sophisticated ontology learning from text system currently under development within the DylanLab of the ILC-CNR in Pisa.

### B. Annotation of Images

The information conveyed by a multimedia document is analyzed and extracted at two different levels: the document level, in which the document (geometrical) layout is investigated considering both text and images, and the image level, in which pictures are examined in order to describe their visual content and to recognize the depicted objects. Two interacting methodologies were adopted to analyze the document at the two different levels, both of them exploiting a domain ontology.




Fig. 4. The graph structure representing the page layout. Each graph vertex is associated to a text caption or to a product image. Images are extracted from the page background by means of a segmentation algorithm.



Fig. 5. The final result of the geometrical layout analysis algorithm. The figure depicts the subgraphs extracted and the thickness of the edges is related to the score obtained by the caption-figure association.

A catalogue document usually provides a rich source of structure. As far as the first algorithm is based on a modified version of the data mining system SUBDUE ([6]). SUBDUE is a system that discovers interesting substructures in structural data based on inexact graph matching techniques and computationally constrained beam search algorithm guided by heuristics. In particular, a substructure is evaluated based on how well it can compress the graph, according to the minimum description length principle. Highly compressing structures can be considered as building block of the entire data set.

To apply the algorithm to a document, we construct a graph representing the page layout (Fig.4). The graph representation consists of two types of vertices, image (classified in three categories: highlight, scene and miscellaneous) and text paragraph, and one type of arc representing the spatial relationships between the vertices (e.g. top-left, overlapping).

The algorithm can provide abstract structured components resulting in a hierarchical view of the document that can be analyzed at many levels of detail and focus. Usually in catalogues a common recurring structure is formed by image and caption (Fig.5). It is remarkable fact that, while text and images may be separately ambiguous, jointly they tend not to be. Establishing from the catalogue structure meaningful links between images and text paragraphs allows exploiting the semantic annotation of the textual part to semantically annotate the images or to guide image processing algorithms in order to recognize the object depicted or to infer correspondence between words and particular image structures (as in [7][8]). Domain knowledge can be also added to guide the discovery process and to separate the important substructures from the irrelevant ones.

The methodology adopted to describe the image content relies on MEMORI ([9][10]), a system for the detection and recognition of objects in digital images using pre stored visual information obtained from shots or 3D models. Object recognition plays a crucial role in Computer Vision, especially in the semantic description of visual content. Although object recognition has been intensely studied, it still remains a hard and computational expensive problem. The main difficulty in




Fig. 6. The MEMORI system analyzes color digital images in order to detect and recognize objects. The figure shows on the left an image extracted from a catalog, and on the right the recognized object along with the resulting annotation compliant to the MPEG-7 standard.

the description of image content is the lack of information about the kind and the number of objects possibly present. Moreover, objects can appear at different locations in the image and they can be deformed, rotated, rescaled, differently illuminated or even occluded with respect to a reference view. In order to simplify object detection and to reduce computational cost, many systems (e.g. [11]) limit the recognition to specific classes of objects. In these cases, a priori knowledge permits to select the most descriptive features for the objects at hand and to circumscribe the search space. However, even under this restriction, high classification performance is seldom reached. Moreover, many object recognition systems rely on user interaction to label as wrong or correct the returned items or to improve system response [12].

The MEMORI system tackles the object recognition problem by segmenting the input image into regions, applying a region grouping algorithm, which interacts with an object classifier, producing a set of object candidates and filtering the candidate list (Fig.6). The object segmentation and recognition modules need domain knowledge in the form of object snapshots from multiple view points.

Supporting the extraction and recognition process with a domain ontology allows the development of context-aware strategies in order to guide and focus the multimedia semantic analysis. The knowledge of the image context permits the object recognition module to restrict the search to a limited subset of objects and to refine the heuristics weighting object hypotheses according to their own context. The other way round content based image analysis allows the acquisition and exploitation of similarity relations among multimedia entities thus allowing to refine and enrich the knowledge representation modeled in the domain ontology.

The experiments and tests conducted so far show that the proposed approach is a promising method for the detection and recognition of objects and for image annotation. An extensive assessment of the performance was conducted on a synthetic test-bed. A set of synthetic images has been created by drawing on a non uniform background rotated and rescaled

version of objects taken from COIL-100 [13], a database widely adopted in the object recognition community.

The achieved recognition rates were compared with other state of the art recognition methods ([14][15]), MEMORI performs best in all experiments, regardless of the number of training views. The recognition rate is over 96% with as few as only 4 training views, demonstrating the robustness of the method. New techniques to make the system robust with respect to illumination changes and partial occlusions are currently under development. Future work will extend the evaluation to more complex and non synthetic images. The potential synergy between visual similarity and semantic similarity measures based on ontologies will also be investigated and exploited.

### C. Elicitation and Refinement

Once information has been extracted from text and images, and stored in the XML-based format partially described in Figure 3, the next task is to make its meaning explicit (and thus machine-understandable) by transforming it into a suitable RDF/OWL representation.

This transformation has two parts<sup>1</sup>:

- The first step, called *syntactic elicitation*, aims at producing a collection of RDF statements, namely an explicit representation of the knowledge content of the XML elements in the schema. For example, we want to construct the fact (expressed by an RDF statement) that the string “900.582.56” is the ID of “SANELA” (actually, of some entity which happens to be a product named “SANELA”).
- The second step, called *semantic elicitation and refinement*, aims at leveraging the RDF representation to a full semantic level, where entities are assigned a proper identifier (a URI), properties are associated to some data type or object property in some domain ontology (more precisely, are replaced by the URI of their ontological counterpart), and finally entities are assigned to the most appropriate class (e.g. “SANELA” should be assigned to the class of cushions, which belongs to a hierarchy of classes which is very likely to include among its ancestors the class of products and of physical entity).

The first step, in our implementation, is rather simple, as it is performed via a simple XSL transformation from the XML schema depicted in Figure 3 to a collection of RDF statements expressed in the RDF XML syntax. The only tricky part is the decision of what statements are to be produced from the XML file (in fact, we observe that the simple snippet in Figure 3 contains a large number of implicit statements, and therefore one needs to select the most useful ones). The outcome of this step is an RDF file containing a potentially large number of

<sup>1</sup> More details are provided in [11], where the entire VIKEF knowledge pipeline is described thoroughly.

facts about the product catalog.

The second step is by far more interesting. Indeed, the goal is to qualify the RDF statements generated in the previous step by linking their constituents to some pre-existing ontology<sup>2</sup> (this is what we call semantic elicitation). In our approach, this task can be decomposed in two different sub-tasks: (1) entity-level elicitation, and (2) class/property-level elicitation.

The first sub-task is implemented as a problem of matching entity descriptions onto the entities stored in an repository called OKKAM. A full description of OKKAM is beyond the scope of this paper; the main idea is that it creates and stores URIs (which can then be reused through multiple applications) together with additional information, including any known description of the entity itself. Once a new entity is recognized in any digital document (plain text documents, relational databases, HTML pages, and so on), OKKAM can be queried to check whether that entity already has a known URI, which is returned for reuse; if no match is found, a new entity is created, and stored with its URI and all available descriptions.

The second task uses a tool, called CtxMatch2.0 [16] for schema and ontology matching, which is a VIKEF motivated extension of a pre-existing tool. The details of semantic elicitation with CtxMatch2.0 are provided in other papers (see e.g. [17]). In the scenario described in this paper, the tool is used to match categories and relations extracted from catalogs to classes and properties found in any available domain ontology. The matching method uses two sources of information: the hierarchical structure of elements (which is particularly important in product catalogs) and lexical information associated both to catalog labels and ontology elements.

The outcome of this second step is a refined RDF file where linguistic descriptions of entities are replaced by unique URIs (which can be used later to merge RDF graphs produced from different catalogs or in general from different collections of documents), category names are replaced by the URIs of ontological classes (if available), and relation names are replaced by the URIs of data type or object properties (if available). We notice that entities may be classified using complex concepts, compositionally constructed from their linguistic descriptions (for example, "SANELA" will correspond to a cushion made of cotton).

A final remark is that the outcome of this elicitation and refinement phase includes a mixture of knowledge coming from multimedia sources, namely text and pictures, in a single representation.

#### D. Reuse of existing ontologies

Ontology engineering is a very time consuming and subjective task. It is time consuming because it is hard to model not only new, but also already known domains, and subjective since in most of the cases the same domains are

modeled in different ways by different ontology engineers. Even if the domains are the same, the perspectives under which they are modeled are usually different. One way to detect the differences (or similarities) of the modeling perspectives is by analyzing the relations existing between the concepts. Due to the already mentioned intrinsic difficulties of the ontology engineering task, advanced research is being performed in order to find a solution to (partially) overcome these problems. The method and tools aim to support the ontology engineer in finding existing Semantic Web ontologies written in OWL (or parts of them) that model the targeted domain in a similar way; this is, with the same modeling perspective. The approach relies on the existence of searchable pool(s) of ontologies where candidate ontologies can be searched for and pre-selected based on a user-specified set of desired ontological classes (this set is called fragment). The pre-selection process searches for a specific percentile match of labels and synonyms of them appearing in the user specified fragment and ontologies in the pool of ontologies. Currently we are using the SWOOGLE (<http://swoogle.umbc.edu/>) ontology repository, but any other similar repository can be easily incorporated. The pre-selected ontologies are then further analyzed in detail by analyzing the labels of classes and relations in combination with a lexical resource (currently WordNet). Labels are analyzed and word meanings are related

by various means in order to compute the likelihood of one or the other possible sense of a word to hold in the given context. This is then represented in two ways, in a logical formulae (following the [18] approach) used for detecting similarity of concepts using a reasoner and as weights of relevance for each possible sense in order to compare concepts and decide if they are similar and to which degree this similarity holds. The relevance

measures are combined with the logical results in order to give a measure of the closeness of the considered concepts. The architecture of this approach is presented in Fig. 7. The ontology engineer can then decide to use any of the proposed ontologies as a basis for performing the required extension of the initially specified fragment. There exist several different techniques for matching ontologies, using very different approaches. Ontology mapping, ontology alignment and ontology matching fall all under a broader research area where correspondences between two schemata need to be discovered.




Fig. 7. The architecture of the described approach for the re-use of existing ontologies.

<sup>2</sup> See next section for what concerns the support we provide to build an ontology by reusing existing ontologies.

In some cases rules for mapping schemata need to be created, in others it is enough to know the relations existing between two or more elements from different schemata. The area of alignment or matching of schemata has been long studied and many approaches exist in integrating database schemata, XML schemata in general, query mediation, etc. The closest approaches to ontology alignment are the approaches for schemata alignment, especially taxonomic hierarchies or classification hierarchies in general. A comprehensive study is presented in [19]. In the following we will present some approaches that are more closely related to the one presented in this section.

- The iPrompt [20] tool suite contains tools for ontology merging and alignment among others. Its input are a set of pairs of related/similar concept from different ontologies, based on them it proceeds to find other pairs of similar concepts relying heavily on the structure.
- The GLUE approach [21] uses a machine learning approach to analyze instances. The similarity of concept meaning is defined based on the joint probability distribution of the involved concepts.
- The QOM [22] approach is based on a combination of several approaches that analyze features of different ontologies. The similarity of concepts is later computed by aggregation of different similarity. The approach is iterative; every iteration uses results obtained in the previous ones in order to enhance quality of the results.
- Cupid [23] presents an approach that considers schemata and not instances information for matching XML Schemata mainly. It uses linguistic information and the structure of the schema. This approach uses also keys, referential constraints and other schema information to derive more precise matching results. Cupid also handles context dependent matches of shared type definitions that are used in several larger structures in the schema.
- The MoA [24] is an approach to merge and align OWL ontologies. This approach uses linguistic methods to disambiguate the meaning of concepts and proposes an algorithm to detect and represent in a *Semantic Bridge* semantic equivalence between concepts. The *Semantic Bridge* can represent equivalence of concepts and properties, subconcept and subproperties and identity of instances. A merging algorithm is also presented.
- The OMEN [25] approach uses Bayesian Networks (BN) for deciding the match of concepts based on an initial *hand made* match. Based on known matches it analyzes the structure (e.g. domain and range of properties) to derive further matches. The algorithm operates iteratively and produces in every iteration a new match which is used in the following interactions.

#### IV. CONCLUSIONS

In this paper we presented a methodology for supporting the use of background ontologies in the task of information

extraction from multimedia sources, in particular from product catalogues. Our methodology tries to enable a virtuous circle by which domain ontologies are used in the extraction process, and at the same time the extraction process becomes a way for creating or extending the available ontologies. The result of the extraction process is a semantically rich representation of the content of catalogs, where knowledge extracted from texts (e.g. product descriptions) is integrated with knowledge extracted from pictures, and made available for any service one may want to build on top of it.

#### REFERENCES

- [1] M. T. Pazienza, A. Stellato, M. Vindigni, Cross-lingual multi-agent retail comparison, AWSS 2003 Workshop on Applications, Products and Services of Web-based Support Systems, 13 Oct. 2003, Halifax Canada in conjunction with Web-Intelligence international Conference WI2003.
- [2] V. Svatek, J. Kosek, M. Labsky, J. Braza, M. Kavalec, M. Vacura, V. Vavra, V. Snasel: Rainbow – multiway semantic analysis of websites, in *Proceedings of the 14th International Workshop on Database and Expert Systems Applications (DEXA'03)*, 2003, p. 635.
- [3] A. Lenci, S. Montemagni, V. Pirrelli, Chunk-it. An Italian shallow parser for robust syntactic annotation. In A. Zampolli, N. Calzolari, L. Cignoni, (eds.), *Computational Linguistics in Pisa - Linguistica Computazionale a Pisa. Linguistica Computazionale, Special Issue, XVI-XVII, (2003)*, Pisa-Roma, IEPI. Tomo I, 2003, pp. 353–386.
- [4] R. Bartolini, A. Lenci, S. Montemagni, V. Pirrelli, Hybrid Cosntraints for Robust Parsing: First Experiments and Evaluation, in *Proceedings of LREC 2004: Fourth International Conference on Language Resources and Evaluation*. Lisbon, Portugal, 26th, 27th & 28 May 2004, Volume III, Paris, pp. 859-862.
- [5] R. Bartolini, D. Giorgetti, A. Lenci, S. Montemagni, V. Pirrelli, Automatic Incremental Term Acquisition from Domain Corpora, in *Proceedings of the 7th International conference on Terminology and Knowledge Engineering (TKE2005)*, Copenhagen Business School, Copenhagen, Denmark, 2005, pp. 17–18.
- [6] J. Coble, R. Rathi, D. J. Cook and L. B. Holder, Iterative Structure Discovery in Graph-Based Data, *International Journal on Artificial Intelligence Tools*, 14(1-2):101-124, 2005
- [7] P. Duygulu, K. Barnard, N. de Freitas, and D. Forsyth. Object Recognition as Machine Translation: Learning a lexicon for a fixed image vocabulary. In *European Conference on Computer Vision (ECCV) Copenhagen, 2002*
- [8] C. Carson, S. Belongie, H. Greenspan, and J. Malik. Blobworld: Image segmentation using expectation-maximization and its application to image querying. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 24(8):1026-1038, 2002.
- [9] M. Lecca, Object Recognition in Color Images by the Self Configuring System MEMORI, *International Journal of Signal Processing*, Vol. 3, No. 3, pp. 176-185, World Enformatika Society, 2006
- [10] C. Andreatta, M. Lecca and S. Messelodi, Memory-based Object Recognition in digital Images, 10th International Fall Workshop - Vision, Modeling, and Visualization - VMV 2005, Erlangen, Germany, November 16-18, 2005
- [11] A. Hoogs, R. Collins, R. Kaucic, and J. Mundy. A common set of perceptual observables for grouping, figure - ground discrimination, and texture classification. *IEEE Transaction on Pattern Analysis and Machine Intelligence*, (4):458–474, 2003.
- [12] B. Ko and H. Byun. Extracting Salient Regions And Learning Importance Scores In Region-Based Image Retrieval. *International Journal of Patter Recognition and Artificial Intelligence*, (17(8)):1349–1367, 2003.
- [13] S. A. Nene, S. K. Nayar, and H. Murase. Columbia object image library (COIL-100). In *Technical Report CUCS-006-96*, Columbia University, 1996.
- [14] S. Obdrzalek and J. Matas. Object recognition using local affine frames on distinguished regions. In Paul L. Rosin and David Marshall, editors,

- Proceedings of the British Machine Vision Conference, volume 1, pages 113-122, London, UK, September 2002. BMVA.
- [15] Ming-Hsuan Yang, D. Roth and N. Ahuja. Learning to Recognize 3D Objects with SNoW. In ECCV-2000, The Proceedings of the Sixth European Conference on Computer Vision, volume 1, pages 439-454, 2000.
  - [16] Bouquet, P., Serafini, L., Zanobini, S.: Peer-to-peer semantic coordination. *Journal of Web Semantics* 2 (2005)
  - [17] Bouquet, P., Serafini, L., Zanobini, S., Sceffer S.: Bootstrapping semantics on the web. meaning elicitation from schemas. WWW2006, Edinburgh (Scotland, UK), May 2006
  - [18] R. Stecher, C. Niederee, P. Bouquet, T. Jaquin, S. Ait-Mokhtar, S. Montemagni, R. Brunelli, G. Demetriou, "Enabling a knowledge supply chain: from content resources to ontologies", in *Proceedings of the ESWC2006 workshop on "Mastering the Gap: from information extraction to semantic representation"*, CEUR workshop proceedings, vol. 187, Budva (Montenegro), 11 June 2006.
  - [19] Pavel Shvaiko, Jérôme Euzenat: "A Survey of Schema-Based Matching Approaches", in *Journal of Data Semantics IV*, 2005.
  - [20] N. Noy and M. Musen: "The PROMPT suite: Interactive tools for ontology merging and mapping", in Technical report, SMI, Stanford University, CA, USA (2002)
  - [21] A. Doan, P. Domingos and A. Halevy: "Learning to match the schemas of data sources: A multistrategy approach", 2003.
  - [22] Marc Ehrig and Steffen Staab: „QOM - Quick Ontology Mapping“, in *Proceedings of the International Semantic Web Conference*, 2004.
  - [23] Jayant Madhavan, Philip A. Bernstein and Erhard Rahm: "Generic Schema Matching with Cupid", in *The VLDB Journal*, 2001.
  - [24] Jaehong Kim, Minsu Jang, Young-Guk Ha, Joo-Chan Sohn and Sang Jo Lee: "MoA: OWL Ontology Merging and Alignment Tool for the Semantic Web" in *Innovations in Applied Artificial Intelligence: 18th International Conference on Industrial and Engineering Applications of Artificial Intelligence and Expert Systems*, 2005.
  - [25] Prasenjit Mitra, Natasha F. Noy and Anuj Jaiswal: "OMEN: A Probabilistic Ontology Mapping Tool", in *Proceedings of the International Semantic Web Conference*, 2005.