
A hybrid ontology and visual-based retrieval model for cultural heritage multimedia collections

Stefanos Vrochidis*, Charalampos Doulaverakis,
Anastasios Gounaris, Evangelia Nidelkou,
Lambros Makris and Ioannis Kompatsiaris

Informatics and Telematics Institute,
Centre for Research and Technology Hellas,
Thessaloniki, Greece

E-mail: stefanos@iti.gr

E-mail: doulaver@iti.gr

E-mail: gounaris@iti.gr

E-mail: nidelkou@iti.gr

E-mail: lmak@iti.gr

E-mail: ikom@iti.gr

*Corresponding author

Abstract: This paper introduces a hybrid multimedia retrieval model that is capable of retrieving cultural heritage multimedia content, based on their semantic annotation with the help of an ontology and on low level visual features with a view to finding similar content. The main novelty is the way in which these techniques cooperate transparently during the evaluation of a single query in a hybrid fashion, making recommendations to the user. A search engine has been developed implementing this model, which is capable of searching through cultural heritage multimedia collections, and indicative examples are discussed, along with insights into its performance.

Keywords: hybrid; ontology; content-based; retrieval; cultural heritage libraries.

Reference to this paper should be made as follows: Vrochidis, S., Doulaverakis, C., Gounaris, A., Nidelkou, E., Makris, L. and Kompatsiaris, I. (2008) 'A hybrid ontology and visual-based retrieval model for cultural heritage multimedia collections', *Int. J. Metadata, Semantics and Ontologies*, Vol. 3, No. 3, pp.167–182.

Biographical notes: Stefanos Vrochidis received the Diploma Degree in Electrical and Computer Engineering from Aristotle University of Thessaloniki, Greece and the MSc Degree in Radio Frequency Communication Systems from the University of Southampton, UK, in 2000 and 2001, respectively. He is a research associate with the Informatics and Telematics Institute. His research interests include semantic multimedia analysis, retrieval, audiovisual search engines, and telecommunications. He is the co-author of three papers in refereed journals, three book chapters and more than ten papers in international conferences. He is a member of the Technical Chamber of Greece.

Charalampos Doulaverakis received his Diploma Degree in Electronic and Computer Engineering from the Technical University of Crete, Greece in 2003 and his MSc Degree in Advanced Computing from the Aristotle University of Thessaloniki, Greece in 2008. Since 2004, he has been a research associate with the Informatics and Telematics Institute/Centre for Research and Technology Hellas, Thessaloniki, Greece. His research interests include image processing, ontology engineering and semantic web. He is a member of the Technical Chamber of Greece.

Anastasios Gounaris is a Lecturer at the Department of Informatics of the Aristotle University of Thessaloniki, Greece. Prior to that, he was a Visiting Lecturer with the University of Cyprus, a research associate with the School of Computer Science of the University of Manchester and he has also collaborated with the Centre of Research and technology – Hellas CERTH. He received his PhD from the University of Manchester, UK in 2005. His research interests are in the area of grid and distributed data management, resource scheduling in grids, autonomic computing, adaptive query processing and semantic technologies. He is a member of the Technical Chamber of Greece, the Association for Computing Machinery, USA, and the Computer Society of IEEE.

Evangelia Nidelkou received the Diploma Degree in Electrical and Computer Engineering and the MSc Degree in Medical Informatics from Aristotle University of Thessaloniki, Greece in 2005 and 2007, respectively. She is a research associate with the Informatics and Telematics Institute. Her research interests include semantic web, biomedical informatics and ontology engineering. She is a member of the Technical Chamber of Greece.

Lambros Makris holds a Diploma in Electrical Engineering and a PhD in Multimedia Communications from the Aristotle University of Thessaloniki. He is currently a research associate at the Informatics and Telematics Institute. His main research interests include hypermedia and multimedia technologies, databases, data security, and mobile applications. He is a member of the IEEE and the Technical Chamber of Greece.

Ioannis Kompatsiaris received the Diploma Degree in Electrical Engineering and the PhD Degree in 3-D model based image sequence coding from Aristotle University of Thessaloniki, Greece, in 1996 and 2001, respectively. He is a senior researcher (Researcher C') with the Informatics and Telematics Institute. His research interests include semantic multimedia analysis, indexing and retrieval, multimedia and the Semantic Web. He is the co-author of six book chapters, 21 papers in refereed journals and more than 60 papers in international conferences. He is an IEEE member and a member of the Technical Chamber of Greece.

1 Introduction

Multimedia content management plays a key role in modern information systems. Usually in the past, medium size, mainly textual archives used to be the only resources for knowledge management, whereas nowadays, organisations and companies handle very large quantities of multimedia information, basically due to the dramatic decrease of digital storage cost. As a result, ranging from personal photo collections to media archives, cultural heritage collections and bio-medical applications, an extremely valuable information asset is in the form of images and video. To enable the same functionalities for the manipulation and knowledge retrieval from such visual content as those provided for text processing, a key aspect is the development of more efficient search engines for image and video files.

To date, two main approaches to image search engine techniques have been proposed, namely annotation-based and content-based. The former is based on image metadata or keywords that annotate the visual content or they refer to the properties of the image file. Examples of image file properties include the name of the image file, its creation date, copyright information, image format, resolution and so on. On the other hand, content metadata correspond to the properties of the entities depicted. For instance, a photograph from someone's holidays may be tagged with concepts such as 'beach' and 'mountain'. Several variants of annotation-based multimedia search engines have been proposed. Some of them assume manual annotation (Flickner et al., 1995), while others provide support for automatic annotation (e.g., Smeulders et al., 2000), by exploiting, for example, relevant text associated to the image and by applying natural language processing techniques. A well known example that falls into this category is Google Image Search.¹

This search approach has benefitted significantly from the advances in the Semantic Web and ontologies

(Sinclair et al., 2005b), so that annotations can have well-defined semantics. Ontologies are "an explicit specification of a conceptualisation" (Gruber, 1993), and they guarantee firstly a shared understanding of a particular domain, and secondly, a formal model that is amenable to unsupervised, machine processing. The use of ontologies has also made possible the integration of different content under a unified description base where various collections can be accessed using a common querying framework. For example, some museums use ontologies for storing and describing their collections, so that users can browse and explore the museum collections, and understand the way in which the items are described and organised, based on their semantics. Indicative examples of such systems are Artefacts Canada² and MuseoSuomi.³

However, annotation and semantic-based search are often insufficient when dealing with visual content. When searching through cultural heritage material for example, it is very common to look for images that are visually similar but may be annotated in a different way (thus a strict annotation search cannot retrieve all relevant images). In addition, in many cases, manual annotation is not available or is incomplete. To tackle this problem, a second complementary approach has been devised: content-based search. The core idea is to apply image processing and feature extraction algorithms to the visual content and extract low-level visual features, such as colour layout and edge histogram (Smeulders et al., 2000). The retrieval is performed based on similarity metrics between such features, attempting to imitate the way humans perceive visual similarity (Mich et al., 1999). The most common way for conducting a search is using the query-by-example methodology (Izquierdo et al., 1993; Naphade et al., 2000) where a user is expected to select a visual multimedia item depicting the desired object or scene of interest and retrieve visually similar multimedia content. As this scenario is not always

feasible, different strategies have also been followed as the query-by-sketch (Chang et al., 1998; Flickner et al., 1995) paradigm, where the user can draw a shape or select a colour layout similar to that of the desired images thus negating the need of an existing starting query image. While content-based search techniques provide fairly acceptable results, they fail to capture the underlying conceptual associations of the visual content and they often generate visually similar but semantically unrelated retrieval results. This is due to the difficulty, which is also known as the *Semantic Gap*, of automatically associating visual features to semantic concepts. While a substantial research effort, especially using Semantic Web technologies, to address this problem, which has produced significant results and solutions (Dasiopoulou et al., 2005; Little and Hunter, 2004; Athanasiadis et al., 2005), has relatively recently started, until today none of the proposed techniques has been adopted for analysis and automatic visual content-based annotation purposes as they are usually applicable to a specific domain only (e.g., football or holiday) or are strongly dependent on content quality.

This paper focuses on a hybrid retrieval model by combining in a novel way the content and annotation based approaches to images from the culture domain. This new model does not just provide an integration or mixing of results but uses the results from each stand-alone approach to generate new searches and enhance the final list of images presented to the user taking into account both semantic and visual similarity. A search engine has been developed implementing this model using images from the culture domain. The ontology-based search builds upon the knowledge base storing the pre-existing semantic annotations of the collection items. Regarding content-based search, the engine employs state-of-the-art techniques, which involve automatic segmentation of 2D visual content and MPEG-7 features extraction, while the novel, hybrid search functionality is capable of extending either the content-based search by making user suggestions of additional, potentially interesting results that are semantically related to the visually similar ones, or the semantic search by recommending images, which are visually similar with the results. Extending the procedure for video files, it is assumed that temporal aspects are not important in the retrieval task, thus the search engine can analyse shot keyframes independently, which leads to a unified way of handling image and video.

The remainder of the paper is structured as follows. Section 2 introduces the retrieval model on which the engine is built and describes the visual, the semantic and the hybrid retrieval. The implementation and evaluation procedures including the presentation of the search engine, the ontology structure, the visual content and the ground truth definition are presented in Section 3, while results of the hybrid engine and insight into the performance of the different approaches appear in Section 4. Section 5 deals with the related work and eventually, Section 6 concludes the paper.

2 Hybrid retrieval of visual content

As mentioned previously, the search engine described hereby supports three modes of queries and retrieval of images and video, namely:

- content-based retrieval
- ontology-based retrieval
- hybrid retrieval, which builds upon the combination of the two aforementioned methods.

In the remainder of this section, the technical details of the three different flavours of the search engine will be discussed.

2.1 Combining visual and semantic information

The main objective behind this retrieval model is to allow a user to complement a query primarily addressed using one of the visual or semantic mode with the other. With a view to supporting such functionality, the hybrid search engine provides a novel retrieval method, in which both visual and ontology search are employed for the same query. The novel retrieval method, employed by the model, automatically combines different types of search results, and complements content-based search with ontology-based search and vice versa. Starting with one mode, information arising from the complementary mode is used to enhance the results. The additional results presented, are considered to be a set of recommendations, which are generated in a transparent to the user way and they are visually or semantically related to the initial results, broadening in that way the desirable query. It is important to note that the hybrid engine generates the new queries involved to retrieve more results in a way transparent to the user. The final result sets are integrated and are presented to the user in a unified manner.

The mathematical model of the hybrid retrieval system is described below for both cases. We assume that the function $Sem(data_{rdf}, q_{sem})$ is producing the desired output (i.e., the artefacts that satisfy the query) given the RDF⁴ data stored in the knowledge base and the semantic query:

$$Res_{sem} = Sem(data_{rdf}, q_{sem}) \quad (1)$$

where $data_{rdf}$ are the metadata stored in RDF in the Knowledge Base and q_{sem} is the query string in RDQL or SeRQL. Here it has to be mentioned that the Res_{sem} , which represent the resulted artefacts, are ranked randomly by the function $Sem(data_{rdf}, q_{sem})$ as all the output results satisfy 100% the submitted query.

In a similar way, the function: $Vis(data_{desc}, q_{vis})$ outputs the results from content-based search using as data the extracted descriptors of the multimedia content and the visual query:

$$Res_{vis} = Vis(data_{desc}, q_{vis}) \quad (2)$$

where $data_{desc}$ represent the extracted descriptors of the multimedia content, q_{vis} represents the desirable input

(i.e., one or a set of images), for which visually matching content expected to be retrieved and displayed, and Res_{vis} represent the resulted artefacts that satisfy the query. The function $Vis(data_{desc}, q_{vis})$ outputs Res_{vis} in a specific ranking based on the similarity coefficient which derives from the calculation of the distances of the extracted descriptors for the objects included in the query.

Subsequently, two cases of hybrid search are defined:

- the visual search, where the system, given the desirable query, produces visually similar results to the initial object accompanied by a set of recommendations deriving from the transparent semantic query that the visual results produce
- the semantic search, where a user can submit a query by browsing the ontology fields and acquire the results that illustrate the content, which satisfies the constraints of the query, complemented by recommendations based on visual similarity of the initial results.

In the case of visual search the initial visual results are processed in order to provide complementary semantic recommendations based on the semantic feature that is most commonly shared between the initial results. On the other hand, when a semantic search occurs, the results that are produced consist of the initial semantic set accompanied by the visual one, which is based on the visual similarity of a hypothetical image averaging the MPEG-7 features of the initial results. The whole procedure is supported by a detailed mathematical model.

In visual search the output consists of two sets of results: the initial results produced by Res_{vis} and the set of recommendations Rec_{sem} based on semantic search and given by:

$$Rec_{sem} = Sem(data_{rdf}, ResToQ_{sem}(Res_{vis})) \Leftrightarrow (3)$$

$$Rec_{sem} = Sem(data_{rdf}, ResToQ_{sem}(Vis(data_{desc}, q_{vis}))) \quad (4)$$

where the function $ResToQ_{sem}$ creates a new query based on the first set of the results in order to retrieve the semantically related content. This function could be based on a combination of several algorithms in order to produce a query which is based on the information of the visual results. The algorithm adopted for the tests and the evaluation with the search engine for cultural content was based on the semantic concept which appeared more frequently between the results. More specifically the aforementioned function exploits the initial set of results by processing the ontology fields of every output object in order to define the semantic annotation which is mostly shared by these results (e.g., such common annotation could be the the creator, the period, etc.). The query produced leads to a semantic search for content that shares the specific common value in the chosen ontology fields with the results.

The final set of results Res is the set of results from visual similarity Res_{vis} enhanced by the recommendation results Rec_{sem} :

$$Res = Res_{vis} \cup Rec_{sem}. \quad (5)$$

Here the union operator is selected to show that the final set of results is a union of two subsets (i.e., initial results and recommendations). However these two subsets are presented separately as they serve different purposes.

In a similar way, when a semantic search occurs, the results that are produced consist of: the first set provided by Res_{sem} and the second set Rec_{vis} illustrating the recommendations:

$$Rec_{vis} = Vis(data_{desc}, ResToQ_{vis}(Res_{sem})) \Leftrightarrow (6)$$

$$Rec_{vis} = Vis(data_{desc}, ResToQ_{vis}(Sem(data_{rdf}, q_{sem}))) \quad (7)$$

where the function $ResToQ_{vis}$ constructs a query taking into account the visual features of the initial results. The algorithm employed for this function produces descriptors of a hypothetical object by averaging the descriptors of the results. Assuming that the resulted objects would share common visual features due to their semantic relation, the query constructed by the function would search for objects visually similar with the hypothetical one.

The final set of results is:

$$Res = Res_{sem} \cup Rec_{vis}. \quad (8)$$

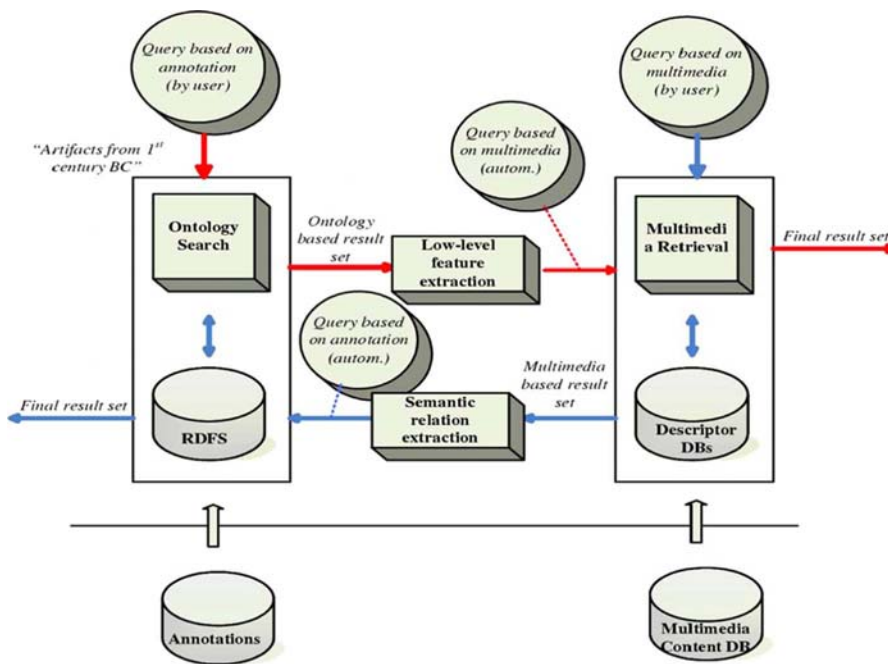
The architecture of the hybrid retrieval model is illustrated in Figure 1. While methods for fusing together plain metadata and low-level features mainly in web-based search systems have been employed (Petraakis et al., 2006; Cai et al., 2004), to the best of the authors' knowledge this is the first attempt to combine semantic and visual-based feature sources in order to provide a more efficient retrieval strategy.

2.2 Content-based retrieval

In this retrieval mode, described by equation (2), users are able to perform a visual-based search by taking advantage of low-level multimedia content features. The retrieval system can handle 2D still images and potentially video. In this mode, the user provides, as the input query, an example of the multimedia content she or he is interested in, and, based on the extracted descriptors of the input and the indexed offline-generated descriptors of the content repository, the system performs a visual similarity-based search and the relevant results are retrieved. The analysis of 2D images is performed using the approach described in Mezaris et al. (2004) and involves the extraction of MPEG-7 features.

For proper handling of the various content types, different strategies are employed for each type in the offline analysis process. This process involves the following steps:

Figure 1 Hybrid retrieval model architecture (see online version for colours)




- Step 1: Segmentation algorithms are applied.
- Step 2: Low-level visual features are firstly extracted, and secondly normalised.
- Step 3: The results generated by this offline process are indexed and stored in a database.

Figure 2 illustrates the 2D image analysis with potential application to video files.

At runtime, the first two steps are followed for the input query, and the intermediate results are compared against the results generated offline in order to detect visually similar images.

Figure 2 The offline image and video content analysis process (see online version for colours)



To enable meaningful region detection in the available cultural heritage images collections, a segmentation process takes place during the first step using the approach described in Mezaris et al. (2004). There are several advantages in using regions for image retrieval and these are mainly derived from the fact that users usually search for objects displayed in images rather than whole images instead. This is the typical case in the area of cultural heritage as the main interest in retrieval is the item being displayed in an image regardless of any surroundings or background. The applied colour image segmentation

algorithm produces connected regions that correspond to meaningful objects shown in the image by performing the segmentation process in the combined intensity and texture feature space. Results of this segmentation approach on different images from the cultural area are presented in Figure 3.

Figure 3 Cultural heritage images (top) and their corresponding segmentations (bottom) (see online version for colours)



The second step in analysis involves low-level feature extraction from the resulting regions of the segmentation mask and also from the whole image itself. For this purpose, the MPEG-7 features were selected as they represent the state of the art in low-level visual descriptors. For the extraction, the MPEG-7 eXperimentation Model (MPEG-7 XM, 2006) was used as it realises the standardised descriptors and apart from extraction it also utilises methods for similarity based retrieval.

Finally at the third step the feature vectors are indexed and stored in a database in order to be easily accessible during the system runtime.

This procedure could be extended to video analysis. In such a case the video stream is firstly divided into shots using the method described in Kobla et al. (1996). For each detected shot, a keyframe is extracted which is treated as a compact representation of the entire shot. This keyframe is then analysed as in the still image case, i.e., it is segmented into regions and feature extraction using MPEG-7 XM is performed.

2.3 Ontology-based retrieval

This search mode, described by equation (1), is more appropriate for the cases in which the user knows to an adequate degree of confidence the semantic annotation of the material he or she is searching for and for the cases. To submit a query, the user provides constraints on the concepts of the five ontologies of Section 3.2. During search time, the system retrieves the semantically connected content according to users selections. In this way, the system can automatically retrieve items, not necessarily of the same type (e.g., coins and inscriptions), that were made, for instance, of the same material, or in the same period or were found in the same place. In this search mode, the system is capable of handling complex queries that require the combination of multiple concept-based search criteria and thus can retrieve different cultural items that share common data. Other related systems, as that of the State Hermitage Museum, do not allow for this kind of functionality but restrict user choices to a specific item category.

Figure 4 illustrates the application of this technique by presenting a proper interface for ontology based search.

In the example in the figure, a search for the available bibliographic references has been requested by selecting the appropriate class from the ontology and, as filtering predicates, the user has selected items that are exhibited in a specific museum and are referenced in a specific historical

book. The results are displayed as shown in Figure 5. This is an example illustrating that the developed system can handle complex queries that require the combination of multiple concept-based search criteria and thus can retrieve different cultural items that share common data.

3 Search engine and evaluation corpus

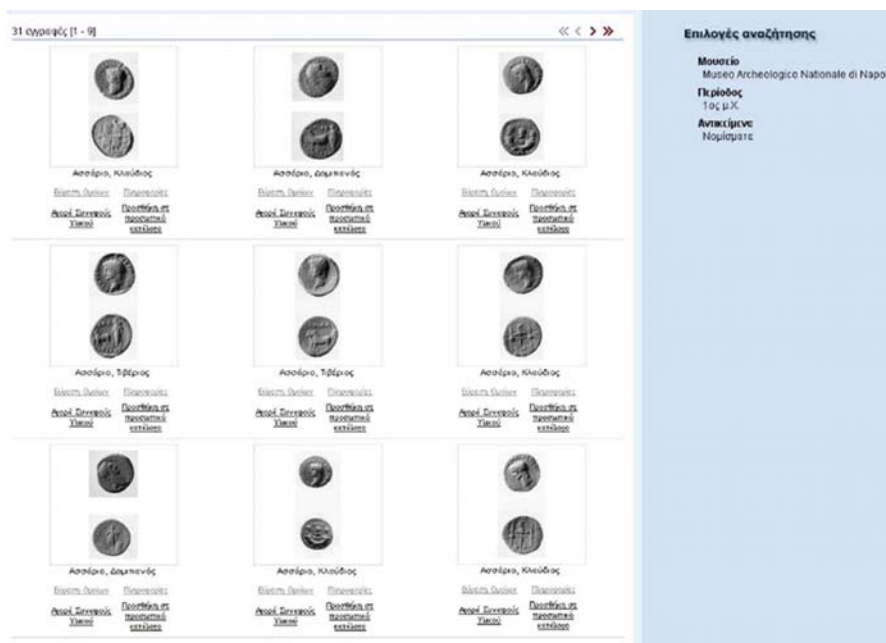
The evaluation procedure took place with the employment of a search engine based on the aforementioned retrieval model and it is capable of retrieving cultural visual content. This engine is used to search through a collection of ancient Greek coins and inscriptions, a collection of paintings and photographs, however its architecture and search approach are actually domain independent. In order to implement the ontology-based retrieval, an ontology-based representation was created with the objective to provide enhanced, unified access to heterogeneous cultural heritage digital databases. The approach followed has been to comply and align with the CIDOC-CRM (Doerr, 2003) core ontology, proposed as an ISO standard for cultural heritage material structure and representation, in order to guarantee the compatibility with other ontology initiatives from the same domain.

Illustrations of the developed search engine are displayed in Figure 4. The design of the GUI and the ontology browser plays a significant role for ontology-based retrieval. As shown in Figure 4, the GUI of our search engine provides a view of the ontologies, enabling the browsing through their structure and hierarchy; selected concepts of each one are automatically organised according to their class hierarchy in a tree-like fashion. Moreover, the stored metadata for each concept are displayed, in order to help the user to formulate the search criteria. Using this design in the system interface

Figure 4 Search engine interface where the user can select one or more ontology instances and concepts to form the search criteria. bottom: the results of the ontology-based search for specific selections (see online version for colours)

The screenshot shows the REACH search engine interface. At the top, it says "REACH - Νέες Μορφές Κατατεμημένης Οργάνωσης και Πρόσβασης σε Υλικό Πολιτιστικής Κληρονομιάς". Below this is a search bar and a "Search Criteria" section. On the left, there are several ontology categories: "Αντικείμενα", "Γεωγραφία", "Χρονολογία", "Υλικό Κατασκευής", and "Βιβλιογραφία". The "Αντικείμενα" category is expanded to show "Επιγραφές" and "Νομίσματα". The "Επιγραφές" category is further expanded to show "Τύπος", "Μαθηματικοπέδιο", "Οπτικά χαρακτηριστικά", "Ενεργημένο κείμενο", "Ατσακόνιστος", "Μονόγραμμα", "Στοιχείο", "Σχέδιο", "Πλευρά", "Εμπροσθεν", and "Όπισθεν". The "Νομίσματα" category is expanded to show "Τύπος", "Μαθηματικοπέδιο", "Οπτικά χαρακτηριστικά", "Ενεργημένο κείμενο", "Ατσακόνιστος", "Μονόγραμμα", "Στοιχείο", "Σχέδιο", "Πλευρά", "Εμπροσθεν", and "Όπισθεν". Below the ontologies, there is a "Class Hierarchy of selected concepts" section showing a tree-like structure of selected concepts. At the bottom, there are search results for "Αναφορές: 245 Εγγραφές" and "Αρχαιολογικών Δελτίων". The results include "Buletin de correspondance hellénique. Ecole Française d' Athènes", "BuletinÉpigraphique. Revue des Études Grecques", "SupplementumÉpigraphicum Graecum", "Γουναροπούλου, Α. – Χατζόπουλος, Μ.Β., Επιγραφές Κέλαι Μακεδονίας (μετά Βασιλίου όρους και του Αβού ποναρού). Τόμος Α'", "Επιγραφές Βιροίας (Αθήνα, 1998)", "M. Ανδρόνικου, Αρχαία Επιγραφεί Βιροίας (δισσαλοβόνη, 1950)", and "Publucéna. Semaines philéropolitaines de l'histoire et de la culture thrace".

Figure 5 The user has asked for artefacts from coins collection (using the ‘Artefacts’ concept) displayed in the “Museo Archeologico Nazionale di Napoli” (using the ‘museum’ concept) and are dated in the 1st century BC (using the ‘period’ concept) (see online version for colours)



has proven to be more user-friendly as all the selections are readily available without introducing any confusion, while in the same time the underlying ontology structure is obvious, letting the user understand the employed knowledge structure. This would not be feasible if a drop-down list GUI architecture had been selected instead. Apart from the ontology-based retrieval, the search engine supports content-based queries in order to produce results depending on visual similarity. Combining the two techniques, the search engine is capable of providing the hybrid searching functionality in order to enhance the initial results and generate recommendations for the user, based on the hybrid retrieval model as was described in detail in Section 2.1.

3.1 Visual content

The main content provider is the Center for Greek and Roman Antiquity (KERA),⁵ which offers a large collection of inscriptions and coins from the Greco-Roman time period, accompanied with detailed documentation. Furthermore, a rich collection of Greek paintings from the 18th to 20th century is provided by the Greek museum: Teloglion Foundation of Art.⁶

In addition to the above, Alinari⁷ Photographic Archives offered a large collection of black and white photographs covering a vast variety of themes mainly from the 19th century.

3.2 Ontology

Cultural heritage collections are accompanied by a rich set of annotations that describe various details related to each item regarding historical data or details regarding

administrative information like, for example, current exhibition location.

However, these annotations are often unstructured or registered in a non-standard form, usually proprietary, for every collection, which renders them unusable for inter-collection searching. To overcome this problem, appropriate ontologies for the cultural heritage domain have been defined. Such ontologies can be used for searching purposes when the search criteria are the collection item annotations rather than the visual appearance. The ontologies are also defined in a way that makes them suitable for administrative purposes of the collection material as during the design process the available metadata were studied and each ontology was defined in a way which was best to represent these data but also taking into account that they should also be usable in a search engine scenario.

The use of an ontological infrastructure was preferred to that of a thesaurus, such as the Art & Architecture Thesaurus (AAT),⁸ mainly because although thesaurus and terminologies can be used to identify concepts, they lack the ability to describe the relationship between them. Thesauri are used to organise a, often large, vocabulary according to linguistic hierarchy and principles (Brewster and Wilks, 2004). While such an organisation could prove useful for classification purposes, in the area of information retrieval, as this work falls into, is proven insufficient. The use of an ontological structure for integration of metadata, organised in a way to reflect human perception, can be employed for more efficient information retrieval as relationships between concepts and data can be exploited and thus provide enhanced retrieval capabilities for the developed search engine (Vallet et al., 2005). Cross-concept relations have been

utilised in the ontology-based and hybrid search, as has been previously described in Section 2.

An additional advantage of using ontologies over thesauri is that the retrieval system could be extended to support rule-based reasoning, in order to augment retrieval capabilities. As an example to show how the search engine would benefit from the use of rules upon the knowledge base, a rule could declare that if a search is employed e.g., for inscriptions based on finding location (place) and this place has mostly uncovered coins of a particular time period then it would be interesting, even for an expert user, to additionally show results of a search based on that particular time period.

Taking into account the content originally available for our use case scenario an ontology infrastructure has been defined to efficiently describe and represent all knowledge related to each collection. The proposed architecture consists of two layers and makes use of five different RDF(S) ontologies, namely *Annotation*, which is generic, and *Coins*, *Inscriptions*, *Paintings* and *Photographs* which are specific to the collection itemsets of our scenario (Figure 6).

A large set of information fields inside the image annotation set is common for each item, regardless of the collection that is part of. As such, it was decided to use a separate, higher-level ontology specifically intended for representing this kind of concepts and relations, which covers information like date and place of creation, current location, construction material, dimensions, etc. Such data is an example of properties that appear in and characterise

every item inside the collection. Consequently, the role of the *Annotation* ontology (Figure 6) is to conceptualise and hold all common data in a structured way, thus forming a representation standard for every collection to be integrated with the search engine. Additionally it is much easier to inter-relate the different items using these concepts as they are all represented in the same ontology while keeping the ontology infrastructure modular.

The properties that are specific to a collection item category are captured by complementary ontologies; more specifically there is a separate ontology for each category, as the particular details that correspond to the collection items can vary greatly for each class. For example, the information that one requires to search through Coin collections, such as monetary subdivision, is significantly different from the information used for Inscriptions searching, such as inscription text. A thorough study of the annotations has shown that this kind of specific information does not overlap across items as is the case with the *Annotation* ontology. As a result, the definition of a *Coins*, an *Inscriptions*, a *Paintings* and a *Photographs* ontology was the most appropriate approach in our case, since it can efficiently handle the data. Moreover, it does not restrict the extensibility of the system as the addition of cultural items of an additional type only requires the definition of a specific domain ontology for that type, and the mapping of its common data to the *Annotation* ontology, i.e., the architecture is extensible and in the future it can cover additional collections like statues, vases, manuscripts, and so on.

Figure 6 A graphical representation of the concepts in the ontology and their relations (see online version for colours)

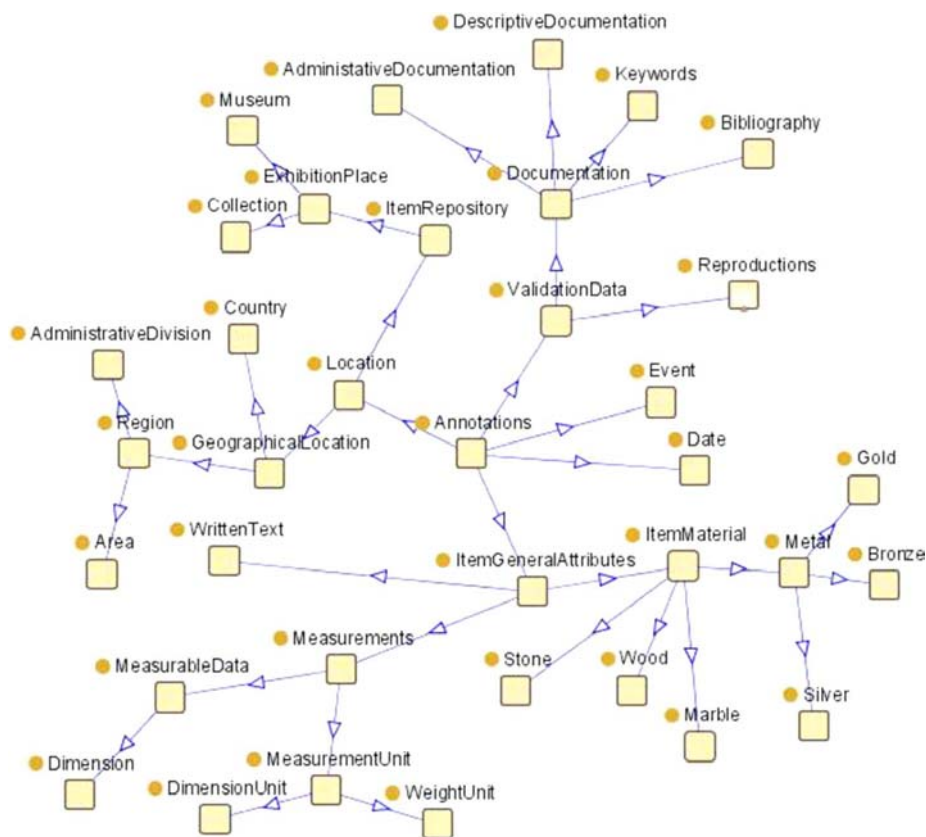
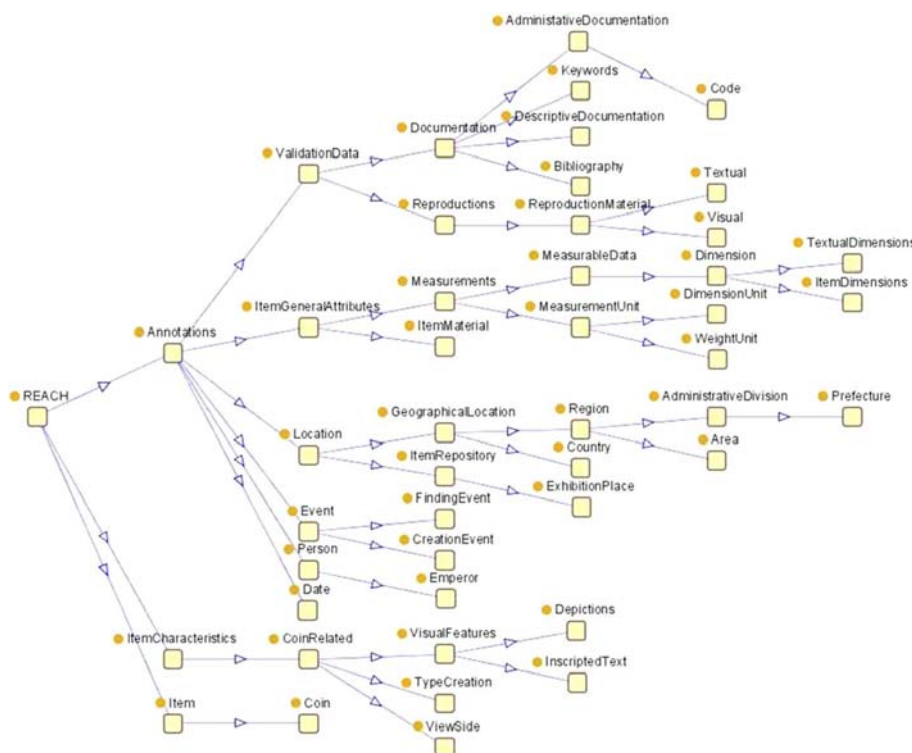


Figure 7 Ontological representation for the Coins (see online version for colours)



A more detailed view of the ontology structure e.g., for the Coin items is shown in Figure 7 where all the concepts required for a proper representation and documentation, are illustrated in a simplified manner for better understanding. In this view, note the extent of the *Annotation* ontology which is used to describe the common concepts for all hosted items. The most important of them are analysed below:

- *Validation Data*: Data used for registering an item. They contain in turn concepts for describing the available documentation regarding the item (the Documentation class). These can be represented by keywords relevant to the item, mainly used for retrieval purposes, bibliographic references, unique identification codes or/and textual descriptions in natural language. Other registration data include the available reproductions of an item, either textual (e.g., a copy of the inscribed text of an inscription) or visual (e.g., photographs).
- *Location*: Concepts that refer to places relative to an item like place of exhibition and finding or creation location.
- *Item General Attributes*: These include concepts that represent measurements for dimensions or weight and the materials used for assembly of an item.
- *Date*: Concept for the temporal registration of item creation or finding events.

Regarding the specific domain related data, i.e., the Coins in the current example, the separate domain ontology contains information like depictions and inscribed text.

As a further step, to support interoperability of the system with other semantic-enabled cultural heritage systems, the aforementioned ontologies were mapped to the CIDOC-CRM (Doerr, 2003) core ontology, which has been proposed as an ISO standard for cultural heritage material structuring and representation. To enable this functionality, appropriate mappings between the concepts of our defined ontologies and the CRM were drawn.

This rather time consuming process gives an advantage to cultural heritage collection management systems, independently of whether they are ontology based or not, regarding the ability to allow for seamless cross-communication with other similar systems as inter-collection relations can be made public using a standardised schema. Some example mappings are presented in Table 1, where on the left column it is illustrated how concepts and relations are represented in our ontology, whereas on the right column the corresponding relation in CIDOC-CRM is shown.

Table 1 Example Mappings of the main General Item Attributes defined in our ontology (left) to the CIDOC-CRM (right)

Original relation	CIDOC-CRM mapping
Item → <i>isMadeOf</i> → ItemMaterial	E22.Man-Made_Object → P45F.consists_of E57.Material
Item → <i>isQuotedIn</i> → Bibliography	E22.Man-Made_Object → P70B.is_documented_in E31.Document

To summarise, the ontology infrastructure of the search engine developed has the following key characteristics:

- It follows a modular design approach. There is a separate module for each collection itemset, which is linked to a generic module, called the *Annotation* ontology. Apart from a clear design, this approach guarantees the extensibility support of the system.
- Instead of using one large and complex ontology, the modular ontologies are kept simple enough, but with no sacrifice in terms of item descriptiveness. This approach makes it easier for visualisation purposes but also enables easier reviewing and editing, e.g., in the development phase.
- The ontologies are mapped to the ISO-standard CIDOC-CRM core ontology, thus allowing for interoperability with other, third-party ontologies from the cultural heritage domain.

3.3 Ground truth definition

The ground truth used in order to evaluate the results of the experiments were different for each retrieval mode.

Regarding the content-based experiments as ground truth was considered the (subjective) visual similarity of the objects (inscriptions, paintings, etc.). More specifically the visual features that were taken into account to prove visual similarity were the shape which can be either rectangular, or oval or even undefined and the colour while the existence of visually related regions between the objects (i.e., human figures, trees, etc.) can also be considered as factor of visual resemblance.

The results from ontology-based queries could be easily evaluated due to the existing annotations i.g., the set of coins that belongs to ‘Domitianos’ period is specific and known.

The recommendations, which are results of the hybrid mode, are considered to be results related to the initial set of the semantic or visually based results. Although the definition of the term recommendation can be subjective, it can be defined as any result which is related to the initial output in terms of visual or semantic similarity. The recommendations serve the purpose of broadening the query by presenting an additional set of results that could be of interest to the user. Hence as ground truth was considered any semantic or visual relation, which can properly justified, that the recommendation set could have with the initial set of results.

4 Results

The content-based and ontology-based modes are complementary to each other, and as such, it is meaningless to compare them directly in terms of metrics like precision and recall.

They actually cover complementary user needs, and consequently, the hybrid proposal is of particular practical

interest. This section serves two purposes. Firstly, to demonstrate the advanced functionality of the hybrid search engine, through some use cases. Secondly, to provide insights into the performance of the two search flavours, mainly content-based and ontology-based. In this section the advanced functionalities of the hybrid search engine are demonstrated through use cases and insights into the performance of the two search flavours are provided.

4.1 Hybrid search: use cases

The hybrid search engine is capable of detecting implicit semantic relationships between visually dissimilar images, and extract the relevant artefacts. To demonstrate this capability, two use cases are presented in this section, which are summarised in Figures 8 and 9.

In the first use case (Figure 8), the input query is the painting “Iera Moni Agiou Pavlou” (“Holy monastery of Saint Paul”) (1933), which was created by Reggos Polykleitos and is displayed in the Teloglion Foundation of Art. During the content-based search visually similar inscriptions are extracted (Figure 8(a)). As shown in the figure, the results have several visual features in common, such as depiction of buildings and mountains. To trigger an ontology-based search, the system retrieves the most common semantic feature, which was found to be the name of the painter. As such, the system is capable of automatically returning the same results as if the user was capable of submitting a query of the type (in natural language)

“Find all the artefacts that look like the painting “Iera Moni Agiou Pavlou” or all the Paintings created by Reggos Polykleitos.” (Figure 8(b))

The first set of results shows visually similar images by illustrating mostly paintings with monasteries and buildings while the recommendations include paintings of the same painter who has created paintings of similar themes (monastery, churches, etc.) and in addition a portrait which cannot be retrieved by visual similarity (Figure 8).

The second use case follows the opposite approach where the initial query is based on ontology fields while the set of recommendations derives from visual similarity. In this scenario, the user searches for inscriptions characterised as anaglyphs. The first set of results includes the two inscriptions which are named as anaglyphs (Figure 9(a)) while the recommendations provide results visually similar with the two inscriptions. As it can be observed in Figure 9(b) the recommendation images include more inscriptions with similar shapes and figures with the two anaglyphs in addition to some images from the photographs collection. However it is remarkable that pictures in the photographs illustrate visually similar regions with the two initial inscriptions as human shapes and sculptures. Consequently the recommendations broaden the initial query as they reveal visually similar images of potential interest that the user may not have been aware of when submitted the query.

Figure 8 The first use case: (a) initial set of results derived from visual similarity search, and (b) set of recommendations based on a complementary semantic query (see online version for colours)



(a)



(b)

Figure 9 The second use case: (a) initial set of results based on a semantic query and (b) set of recommendations includes visually similar images (see online version for colours)



From the aforementioned use cases, two important observations can be made. Firstly, the hybrid search engine is capable of recommending items to the user based on the semantic or the visual similarity of the results which outcomes from the user input query, without any additional effort from the user. Secondly, the user receives feedback on the search criteria of the transparent, ontology-based or visual query. This is a key aspect for the better interpretation of the additional results.

4.2 Performance insights

In the previous section, three different search policies were presented that provide three complementary options to query formulation, so that the users can find their desired content even in the case where the search criteria are rather


complex. In this section, we present a closer inspection on the efficiency and the performance of each of the two basic methods (content- and ontology-based) and draw conclusions on the advantages and disadvantages of each method with respect to precision of retrieval and response times.

The experiments were conducted on a PC, with a P5 3.0GHz Intel CPU and 1GB RAM. The knowledge base containing the ontological metadata is Sesame 1.2 running a MySQL DBMS at the back-end. MySQL is also used to store the actual non-multimedia content and the links to the multimedia files. The dataset is consisted of roughly 4000 images including inscriptions, coins, paintings and photographs along with a complete set of semantic annotations. The visual descriptors are stored in a collocated MPEG-7 XM server. For both content-based

and ontology-based search, five queries, either visual or semantic, were used and the mean times are presented below. To evaluate the content-based search we selected five random images.

Figure 10 shows the Precision-Recall diagram for the content-based retrieval. The curves correspond to the mean precision value that was measured after several retrieval tasks. For the ontology-based search since it is based on selecting available concepts describing the content, the estimation of Precision-Recall diagram is not relevant. The precision of the ontology-based method is, as expected, 100% as we assume that the complete set of the annotations related to the items is manually edited and is precise, whereas the precision of the content-based depends on the efficiency of the feature extraction and distance evaluation algorithm. As such, the content-based method is not as satisfactory, in terms of precision, as the ontology-based search, but as it was explained earlier it serves different user needs.

Figure 10 Precision-recall diagram for the content-based method (see online version for colours)



Also, the precision is degraded because of the available visual content, which is characterised by a rather small variance in terms of structure, i.e., all inscriptions have roughly the same shape apart from those who have sustained damage and have somehow lost their original shape. The average response time for the ontology-based search is 0.163 s, while for the content-based search is 0.773 s. For the content-based search it includes the communication with the MPEG-7 XM server and if we assume that the metadata used (i.e., the visual features) are evaluated and stored in the preprocessing step in the multimedia database instead of the MPEG-7 XM server, then the time cost of content-based search is reduced to 0.068 s.

Comparing the two methods, we should keep in mind that

- these methods work on different representations of the available data and their use is intended to satisfy different needs

- ontology-based search presupposes the manual annotation of collection items.

In summary, ontology-based search aims at making use of the semantic annotations associated to an item, with respect to historical data (e.g., date of creation, place, etc.), while the content-based search aims at making use of lower level characteristics of the multimedia content corresponding to an item, like shape and colour distribution, which can be automatically extracted. Such information is not likely to be found in the metadata accompanying a cultural heritage collection. Someone looking for items that are similar in shape, for instance, will use visual similarity as compared to a user interested in finding items that belong to a certain time period, and thus benefits more from the semantic search engine.

The behaviour of the hybrid search is expected to combine the benefits of the other two approaches providing recommendations to the user in order to broaden the query. Precision-recall graphs for the recommendations are not presented as these strongly depend on the nature of the retrieval task, and on the objective and purpose of the user when submitting a query. A recommendation can be considered highly subjective especially when personalisation of the user is not applied. The personal background of the user is very important when speaking about cultural content as i.e., an archaeologist would have different expectations from a typical visitor of a museum regarding recommendations of a search engine. Consequently, solid measurement method for the hybrid search is difficult to obtain because of the strongly subjective nature of this proposed search option. However the recommendations provided by the system were considered as satisfactory, as they are related either semantically or visually with the initial set of the results and this conclusion was supported by an adequate number of questionnaires which were filled in by users during an evaluation task of the search engine.

Concluding, hybrid search is proposed as a novel way of combining the above two methods to provide results sets that could potentially be of relevance, and are based both on visual features and on the concepts defined in the ontology.

5 Related work

Multimedia search engines have attracted a lot of interest both from the web search engine industry (such as Google, Yahoo!, and so on) and from academia. Also, the emergence of MPEG-7 standard has played a significant role in content-based search becoming a mature technology. For a survey, the reader can refer to Auffret et al. (1999) and Chang et al. (1999). However, as already mentioned, no framework has been proposed on the combined use of ontology- and content-based retrieval.

Several efforts for the efficient management and indexing of cultural heritage collections have been reported, however they all lack the integrated functionality

of the presented hybrid search engine. In MuseoSuomi (Hyvonen et al., 2004) automatic semantic association generation was used through database schemas and metadata using underlying RDF(S) ontologies. The system supports runtime distributed querying through geographically remote collections while at the same time maintaining semantic relations between collection items. The users can search through the collections using a multi-facet interface (Hearst et al., 2002). Although the system can efficiently handle semantic metadata, there is no support for visual search. The SCULPTEUR project (Sinclair et al., 2005b) employs the CIDOC-CRM ontology to enable concept-based browsing of the annotations. Additionally cross collection searching is enabled by contacting web-services, and more specifically an extension to SRW,⁹ to allow for access to different content repositories. CIDOC-CRM is used for mapping each collection's proprietary schema to a standardised structured representation. Queries to the collection are in Common Query Language (CQL) and an mSpace-based (Schraefel et al., 2003) interface is employed to enable searching using the CRM's concepts. Multimedia search is also supported, both 2D and 3D. Nevertheless, SCULPTEUR makes no proposal for coupling semantic data and low-level visual features. The ongoing EU project eChase (Sinclair et al., 2005a) also uses the CIDOC-CRM ontology to map various cultural heritage collections. As the mapping is left to the cultural heritage content providers, there can be different mappings for the same CRM relation thus resulting in loss of semantic links. To tackle this, a reasoner will be employed to determine semantic equivalences. The collections will be queried using the SRW extension developed by SCULPTEUR. In eChase content-based search is also planned to be employed but the available demonstrator does not support such kind of functionality. STITCH, which is part of eChase, tries to extract semantic relations and links with the use of vocabularies. In our case semantic relations are created based on the underlying ontological infrastructure. Semantic links generation through the exploitation of common words and their linguistic association found in the metadata is used in Stuer et al. (2001). In García and Sicilia (2003) the authors try to enhance ontology querying and navigation by taking advantage of the ontological structure and by using different user interface search tactics.

Another cultural heritage collection management system is the State Hermitage museum in St. Petersburg¹⁰ developed by IBM. The online collection includes an ontology for describing the items where the main distinguishing concept between items is item category (paintings, sculpture, furniture, etc.) and they are in their turn divided into more detailed subcategories. However there is no ability for cross-category search based on concepts such as date. Content-based image search is supported through IBM's QBIC (Flickner et al., 1995) but there is no support for coupling together visual- and ontology-based search.

6 Conclusions

In this paper, a novel retrieval model for handling visual and multimedia digital libraries is presented in an efficient and effective manner. The retrieval model proposed model adopts three methods for retrieval: two autonomous (ontology and content-based) and one combinational (hybrid). The ontology-based method makes use of the formal, logic-based representation of semantic mark-up metadata accompanying each collection, while an illustrative user interface is used for graphical query formulation. This method is appropriate when the user is interested in semantically similar results. The content-based method makes use of the low-level visual characteristics of the multimedia material, in order to retrieve items with similar appearance. Although the search engine dealt with 2D images of cultural heritage content, there is the potential of extension based on the proposed model to include video content. A notable feature of this work is its modular and extensible ontology infrastructure, which provides mappings to CIDOC-CRM in order to gain interoperability with other ontologies from the cultural domain. The hybrid method, which is the main contribution of this work, makes a combined use of the aforementioned two methods. Thus, it is capable of offering, as a recommendation, a more complete result set to the user, which comprises both visually and semantically similar items, while the input query remains either solely ontology-based or content-based. In a future work it could be very interesting to introduce personalisation features in order to be able to evaluate in a more proper way the recommendations provided by the system and also enable rule-based reasoning and inference support to further exploit the ontology structured data. The search engine is being used for searching through cultural heritage material.

Acknowledgements

This work was supported by the project REACH "New forms of distributed organisation and access to cultural heritage material" funded by the General Secretariat of Research and Technology of Hellas-Greece and by the project CHORUS "Coordinated approach to the European effort on audiovisual search engines" funded by the European Commission.

References

- Athanasiadis, T., Tzouvaras, V., Petridis, K., Precioso, F., Avrithis, Y. and Kompatsiaris, I. (2005) 'Using a multi-media ontology infrastructure for semantic annotation of multimedia content', *5th International Workshop on Knowledge Markup and Semantic Annotation (SemAnnot 2005) at the 4th International Semantic Web Conference, ISWC 2005*, Galway, Ireland, pp.59–68.

- Auffret, G., Foote, J., Li, C-S., Shahraray, B., Syeda-Mahmood, T. and Zhang, H. (1999) 'Multimedia access and retrieval (panel session): the state of the art and future directions', in Chang, S-F. (Ed.): *MULTIMEDIA '99: Proceedings of the Seventh ACM International Conference on Multimedia (Part 1)*, pp.443–445, ACM Press, New York, NY, USA.
- Brewster, C. and Wilks, Y. (2004) 'Ontologies, taxonomies, thesauri: learning from texts', *Proceedings the use of Computational Linguistics in the Extraction of Keyword Information from Digital Library Content Workshop*, Kings College, London, UK.
- Cai, D., He, X., Li, Z., Ma, W-Y. and Wen, J-R. (2004) 'Hierarchical clustering of WWW image search results using visual, textual and link analysis', *12th ACM International Conference on Multimedia*, New York, USA, October, pp.952–959.
- Chang, S-F., Meng, W.C.H., Sundaram, H. and Zhong, D. (1998) 'A fully automated contentbased video search engine supporting spatiotemporal queries', *IEEE Transactions on Circuits and Systems for Video Technology*, Vol. 8, September, pp.602–615.
- Chang, S-F., Huang, Q., Huang, T., Puri, A. and Shahraray, B. (1999) 'Multimedia search and retrieval', in Puri, A. and Chen, T. (Ed.): *Advances in Multimedia: Systems, Standards, and Networks*, Marcel Dekker, New York, pp.36–55.
- Dasiopoulou, S., Mezaris, S., Kompatsiaris, I., Papastathis, V.K. and Strintzis, M.G. (2005) 'Knowledge-assisted semantic video object detection', *IEEE Transactions on Circuits and Systems for Video Technology, Special Issue on Analysis and Understanding for Video Adaptation*, Vol. 15, No. 10, pp.1210–1224.
- Doerr, M. (2003) 'The CIDOC-CRM an ontological approach to semantic interoperability of metadata', *AI Magazine*, Vol. 24, No. 3, Fall, pp.75–92.
- Flickner, M., Sawhney, H., Niblack, W., Ashley, J., Huang, Q., Dom, B., Gorkani, M., Hafner, J., Lee, D., Petkovic, D., Steele, D. and Yanker, P. (1995) 'Query by image and video content: The QBIC system', *IEEE Computer*, Vol. 28, No. 9, pp.23–32.
- García, E. and Sicilia, M.A. (2003) 'User interface tactics in ontology-based information seeking', *Psychology Journal*, Vol. 1, No. 3, pp.242–255.
- Gruber, T. (1993) 'A translation approach to portable ontology specifications', *Knowledge Acquisition*, Vol. 5, pp.199–220.
- Hearst, M., Elliot, A., English, J., Sinha, R., Swearingen, K. and Lee, K-P. (2002) 'Finding the flow in web site search', *CACM*, Vol. 45, No. 9, pp.42–49.
- Hyvonen, E., Junnila, E., Kettula, S., Makela, E., Saarela, S., Salminen, M., Sreeni, A., Valo, A. and Viljanen, K. (2004) 'Finnish museums on the semantic web. User's perspective on MuseumFinland', *Museums and the Web 2004*, Virginia, USA, March.
- Izquierdo, E., Casas, J., Leonardi, R., Migliorati, P., O'Connor, N., Kompatsiaris, I. and Strintzis, M.G. (2003) 'Advanced content-based semantic scene analysis and information retrieval: the schema project', *Proc. Workshop on Image Analysis For Multimedia Interactive Services*, London, UK, pp.519–529.
- Kobla, V., Doermann, D. and Lin, K. (1996) 'Archiving, indexing, and retrieval of video in the compressed domain', *Proc. SPIE Conference on Multimedia Storage and Archiving Systems*, Vol. 2916, pp.78–89.
- Little, S. and Hunter, J. (2004) 'Rules-by-example – a novel approach to semantic indexing and querying of images', *International Semantic Web Conference*, Vol. 3298, Hiroshima, Japan, pp.534–548.
- Mezaris, V., Kompatsiaris, I. and Strintzis, M.G. (2004) 'Still image segmentation tools for object-based multimedia applications', *International Journal of Pattern Recognition and Artificial Intelligence*, Vol. 18, No. 4, June, pp.701–725.
- Mich, O., Brunelli, R. and Modena, C. (1999) 'A survey on video indexing', *Journal of Visual Communications and Image Representation*, Vol. 10, pp.78–112.
- MPEG-7 eXperimentation Model (MPEG-7 XM) (2006) URL:<http://www.lis.ei.tum.de/research/bv/topics/mmdb/empeg7.html>
- Naphade, M., Yeung, M. and Yeo, B. (2000) 'A novel scheme for fast and efficient video sequence matching using compact signatures', *Proc. SPIE Storage and Retrieval for Multimedia Databases*, Vol. 3972, pp.564–572.
- Petrakis, E., Kontis, K., Voutsakis, E. and Milios, E. (2006) 'Relevance feedback methods for logo and trademark image retrieval on the web', *ACM Symposium on Applied Computing (SAC'2006), Special Track on Information Access and Retrieval (IAR)*, Dijon, France, April, pp.1024–1088.
- Schraefel, M.C., Karam, M. and Zhao, S. (2003) 'mSpace: interaction design for user-determined, adaptable domain exploration in hypermedia', *Workshop on Adaptive Hypermedia and Adaptive Web Based Systems*, Nottingham, UK, pp.217–235.
- Sinclair, P., Addis, M., Fina, D., Lewis, P., Martinez, K., Prideaux, D. and Bormida, G.D. (2005a) 'eCHASE: sustainable exploitation of electronic cultural heritage', *European Conference on the Integration of Knowledge, Semantic and Digital Media Technologies (EWIMT)*, November, pp.411–417.
- Sinclair, P., Goodall, S., Lewis, P., Martinez, K. and Addis, M. (2005b) 'Concept browsing for multimedia retrieval in the SCULPTEUR project', *Multimedia and the Semantic Web, held as Part of the 2nd European Semantic Web Conference*, Heraklion, Crete, Greece, May.
- Smeulders, A., Worring, M., Santini, S., Gupta, A. and Jain, R. (2000) 'Content-based image retrieval at the end of the early years', *IEEE Transactions on Pattern Analysis and Machine Intelligence*, Vol. 22, No. 12, pp.1349–1380.
- Stuer, P., Meersman, R. and Bruyne, S.D. (2001) 'The HyperMuseum theme generator system: ontology-based internet support for the active use of digital museum data for teaching and presentation', *Museums and the Web 2001: Selected Papers. Archives and Museum Informatics*, Pittsburgh, Seattle, USA, pp.127–137.
- Vallet, D., Fernandez, M. and Castells, P. (2005) 'An ontology-based information retrieval model', *The Semantic Web: Research and Applications: Second European Semantic Web Conference, ESWC 2005*, Springer, Heraklion, Greece, pp.455–470.

Notes

¹<http://images.google.com/>

²http://www.chin.gc.ca/English/Artefacts_Canada/

³<http://www.museosuomi.fi/>

⁴<http://www.w3.org/RDF/>

⁵<http://www.eie.gr/nhrf/institutes/igra/index-en.html>

⁶<http://web.auth.gr/teloglion/>

⁷<http://www.alinari.com/>

⁸http://www.getty.edu/research/conducting_research/vocabularies/aat

⁹<http://www.loc.gov/z3950/agency/zing/srw>

¹⁰<http://www.hermitagemuseum.org>