Approaches and Standards for Metadata Interoperability in Distributed Image Search and Retrieval

Ruben Tous, Jordi Nin, Jaime Delgado, and Pere Toran

Universitat Politècnica de Catalunya (UPC BarcelonaTech), Barcelona, Spain
rtous@ac.upc.edu, nin@ac.upc.edu, jaime.delgado@ac.upc.edu, ptoran@ac.upc.edu

Abstract. This paper addresses the general problem of how building a distributed image search&retrieval system that copes with metadata heterogeneity. Firstly, We analyze the usage of metadata in current image repositories, the different metadata schemas for image description and their semantic relationships. After, we provide a general classification for the different approaches which provide a unified interface to search images hosted in different systems without degrading query expressiveness. This paper analyzes how these approaches can be implemented on top of the latest standards in the area, ISO/IEC 15938-12 (MPEG Query Format) for the query interface interoperability and ISO/IEC 24800 (JPSearch) for the definition and management of translations between metadata schemas. Finally, we provide insights into an example of a real distributed image search&retrieval system which provides real time access to Flickr, Picassa and Panoramio.

Keywords: metadata, interoperability, image, information retrieval, standards, jpsearch, mpqf.

1 Introduction

Nowadays, digital images are being generated, distributed and stored worldwide at an ever increasing rate. Consequently, in the recent years, image Search&Retrieval tasks arise as an important issue. There are multiple systems, however, almost every one provides a different search interface and multimedia metadata description format. This fact prevents users from experiencing a unified access to the repositories. Systems aiming to provide a unified query interface to search images hosted in different systems without degrading query expressiveness need to address several questions which include but are not limited to the following:

- Is the system going to harvest all the metadata and store them locally? How frequently do the data change and how is the system going to cope with data volatility?
- Which metadata schema or schemas are going to be exposed to user queries? Is the system going to expose a mediated/pivot schema?

A. Hameurlain et al. (Eds.): DEXA 2011, Part II, LNCS 6861, pp. 234–248, 2011.

- How the mappings among the underlying target metadata schemas are going to be generated?
- Which formalism is going to be used to describe the mappings?
- How is the system going to use the mappings during querying?

Currently many standardization efforts are trying to provide answers to some of these questions. Two of the most relevant initiatives are the ISO/IEC 15938-12:2008 (MPEG Query Format or simply MPQF) [5,10,13] and ISO/IEC 24800 (JPEG's JPSearch framework) [14,15]. While MPQF offers a solution for the query interface interoperability, JPSearch (whose Part 3 makes use of MPQF) faces the difficult challenge to provide an interoperable architecture for images metadata management.

This paper provides an analysis of the current approaches and standards for metadata interoperability in distributed image search&retrieval systems. In order to better feature the problem being faced, firstly the paper examines the main metadata schemas currently used for image description and identifies their semantic relationships. Secondly, the paper provides a classification of the different approaches and evaluates their advantages and drawbacks. Thirdly, the paper analyzes how distributed image search&retrieval systems can be implemented on top of the ISO/IEC 15938-12 and ISO/IEC 24800 standards. Finally, the paper provides insights into an example real distributed image search&retrieval system which provides real time access to Flickr, Picassa and Panoramio.

2 Topology of Digital Image Description Metadata Models

In order to better feature the image metadata heterogeneity problem, let's first study which are the main metadata schemas used for image description and which are their relationships. We have selected the following list of metadata schemas, which is not comprehensive, but allows the reader to obtain some conclusions that can be easily extrapolated:

- *Flickr*. Metadata schema [19] representing the image description model of Flickr, the most popular image hosting and video hosting website.
- *Picasa*. Metadata schema [22] representing the image description model of Picasa, Google's image organizer, image viewer an integrated photo-sharing website.
- *Panoramio*. Metadata schema [20] representing the image description model of Panoramio. A geolocation-oriented photo sharing website.
- *Photobucket*. Metadata schema [21] representing the image description model of Photobucket, an image hosting, video hosting, slideshow creation and photo sharing website usually used for personal photographic albums, remote storage of avatars displayed on internet forums, and storage of videos.
- *DeviantART*. Metadata schema [18] representing the image description model of DeviantART, an online community showcasing various forms of user-made artwork.

- *ORDImage*. Metadata schema of the Oracle Multimedia's ORDImage type, which supports the storage, management, and manipulation of image data.
- *DICOM*. Metadata schema of the Digital Imaging and Communication in Medicine (DICOM), a standard for handling, storing, printing, and transmitting information in medical imaging.
- *MPEG-7*. Metadata schema of ISO/IEC 15938 (Multimedia Content Description Interface).
- *Dublin Core*. Metadata schema of Dublin Core (ISO Standard 15836, and NISO Standard Z39.85-2007).
- *JPSearch* Core Schema. Metadata schema of ISO/IEC 24800-2 (JPSearch Core Schema).
- *EXIF*. Exchangeable image file format is a specification for the image file format used by digital cameras (including smartphones) and scanners. it latest version dated April 2010 (2.3) was jointly formulated by JEITA (Japan Electronic Industries Development Association) and CIPA (Camera & Imaging Products Association).

To visualize and compare the different schemas, we have used OpenII [24], a novel schema management tool. OpenII is a suite of open-source tools for information integration (II). The suite includes the following components:

- *Affinity* is a clustering tool for entire schemas. It groups schemas using hierarchical clustering techniques.
- *Harmony* is a semi-automated tool that finds the correspondences across two data schemas using a set of semantic matchers.
- *Proximity* visualizes in a graphical way the inclusion/exclusion relations between a selected reference model and all the others.
- *Unity* semi-automatically produces a common vocabulary, based on aligning the schema elements in a set of source schemas. The vocabulary consists of a list of canonical terms the source schemas agree on.
- *Yggdrasil* is a repository implementing an entity relationship metamodel (called M3) for both schemas and mappings. This repository is implemented on top of a Postgres database. Yggdrasil allows for importing / exporting XML Schemas.

We have used OpenII to find the relationships among the different schemas, and to visualize the topology of the different image description metadata models. Figure 1 shows the results of clustering all the metadata models with the the OpenII's Affinity tool. EXIF and MPEG-7, being the more comprehensive and general schemas, appear clustered together at the bottom of the figure. Schemas from generic image hosting websites, such as the ones from Flickr, Picasa, Panoramio, Photobucket or DevianART, appear clustered together at the top-left of the figure, with ORDImage not far away. These schemas are flat and simple, and they have significant overlappings. DICOM, a specialized schema for medical image tagging, appears at the right of the figure, while the JPSearch Core Schema appears at the top.

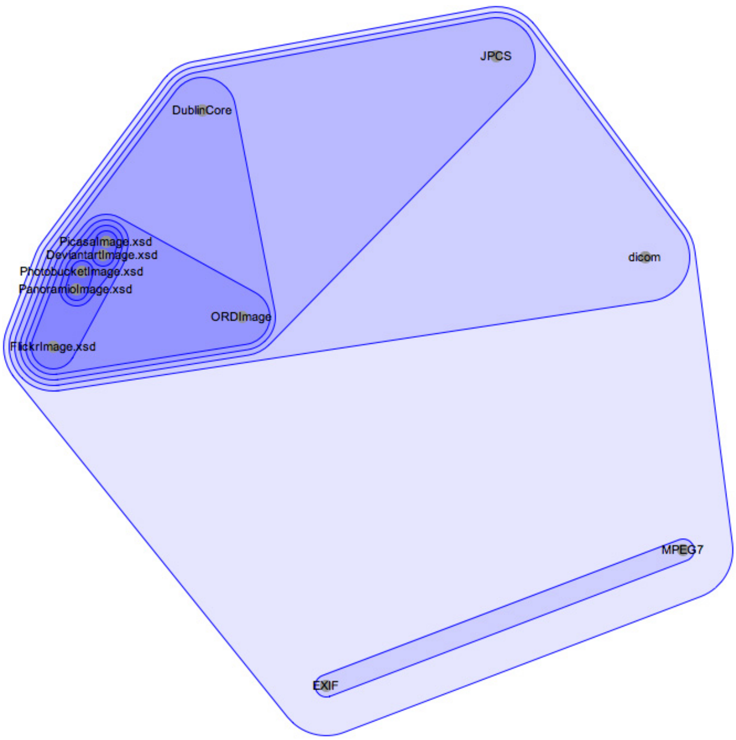

Fig. 1. Relations found for all the schemas using the OpenII's Affinity tool

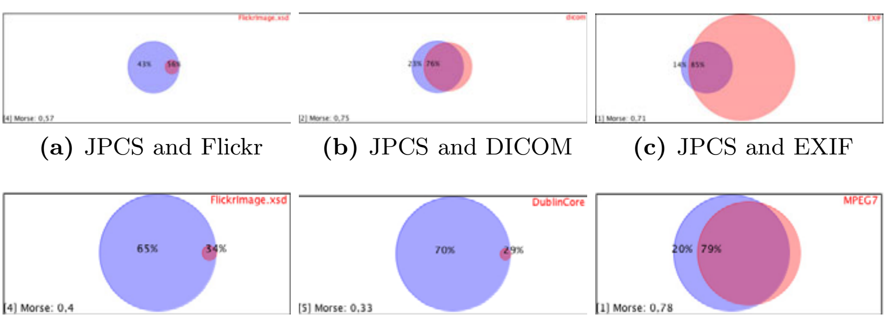

(a) JPCS and Flickr **(b)** JPCS and DICOM **(c)** JPCS and EXIF

(d) EXIF and Flickr **(e)** EXIF and Dublin **(f)** EXIF and MPEG-7

Fig. 2. Relationship between different schemas obtained using the OpenII's Proximity tool

We have also used OpenII to find one-to-one relationships between schemas. The OpenII's Harmony tool allows to automatically obtain the correspondences between two given metadata schemas. Due to the lack of space, we do not include the produced mappings, however, in figure 2 we show the inclusion/exclusion relations between some selected schemas using the harmony tool. Figures 2.(a) and 2.(d) depict that the Flickr metadata information is completely included in the JPSearch Core Schema and EXIF, two general image tagging schemas. Figure 2.(c) shows that EXIF metadata is more general and includes more information than JPCS schema. Additionally, this figure shows that most of the JPCS information (85%) is also included in EXIF. Similar conclusions can be extracted from the Figure 2.(f) where we compare EXIF and MPEG-7 schemas, in this case we also see that MPEG-7 includes more information than JPCS (by comparing the relative sizes of both schemas in Figures 2.(c) and 2.(f)).

Figure 2.(b) illustrates that only the 76% of the DICOM metadata information is shared by JPCS, this has sense because DICOM schema has medical specific information such as patient position or patient orientation. Finally, we would like to comment Figure 2.(e), where we show the relation between EXIF and Dublin core schema. Both schemas are general but the Dublin schema clearly has less information than EXIF, and Dublin information is completely included in EXIF.

3 Approaches to Metadata Interoperability in Distributed Image Search&Retrieval

In this section we provide an overview to the three approaches that we propose to classify the different solutions for the image metadata interoperability problem. Each approach is labelled with a double title. The first part of the title always refers to the most characteristic concept that identifies the architecture of each proposal: Simple Aggregator, Multiple Aggregator and Broker. The second part of the title illustrates the implementation issues implied in each approach. In the Simple Aggregator approach it is necessary to deal with Metadata Conversion issues, while in the Multiple Aggregator approach it is necessary to focus on Inference. Finally, in the Broker approach it is necessary to face Query Rewriting related topics. The architecture in each of the approaches is subdivided into three subsystems: the Querying subsystem, the Central subsystem and the Metadata subsystem. The Querying subsystem manages the queries that the user inputs. The Metadata subsystem represents the different external systems that the Central subsystem makes transparently interoperable. We have focused mainly in the Central and Metadata subsystems.

3.1 Simple Format Metadata Aggregator/Metadata Conversion

The main idea of this approach is that we have a centralized system, called Aggregator, which holds all the metadata needed to satisfy the user queries. The key issue is that we have an intermediate metadata format for the Aggregator.

Thus, all the supported formats need to provide a mapping mechanism to convert from their own metadata format to the Aggregator metadata format. Figure 3 presents the basic Simple Format Metadata Aggregator ($SFMA$) architecture diagram. The Aggregator (that is part of the Central subsystem) receives a query. As we said, in order to answer to that query the Aggregator needs to have all the external metadata formats available and converted into its own one. All supported external Metadata Formats (F_1, F_2, ... , F_n) need to provide some mapping mechanism (F_1 to F_A, F_2 to F_A, ... ,F_n to F_A) to map from their own format to the one supported by the Aggregator (F_A). Thus, assuming that the Aggregator has all the information obtained from the external metadata formats converted to its own metadata format, the system uses this metadata database to obtain the results demanded and send back them to the user.

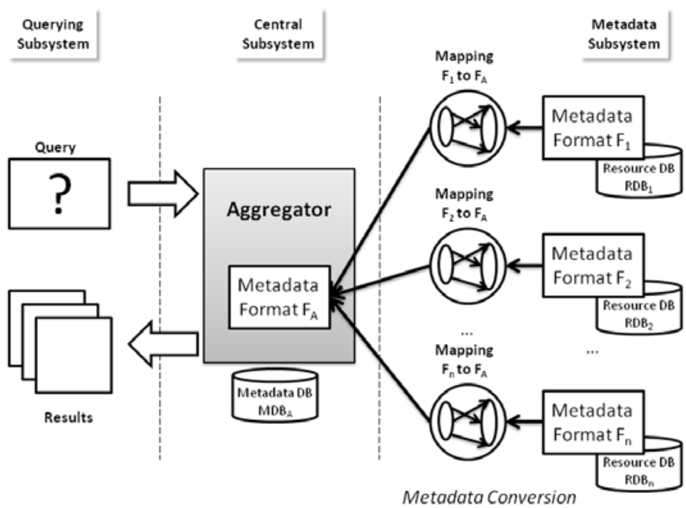

Fig. 3. Basic architecture for the "Simple Format Metadata Aggregator" (SFMA) approach

In Figure 3 we have also added the different databases distribution used by the system. It is interesting to include this information in the figure to compare the different distributions of the databases that every approach implies. Basically, we will have two types of databases: (1) resource databases and (2) metadata databases. In the resource database is where the actual resources we search for are located, i.e. the image file. On the other hand, the metadata database contains all the metadata information associated to the images of the resource database. In the SFMA approach, the metadata database is located at the Central subsystem. It aggregates the different metadata databases into a single one labelled MDB_A, which is the unique metadata database that the system uses. However, the resource databases (RDB_1, RDB_2, ... , RDB_n) may remain at the external systems and the Aggregator obtains the resource directly from there after the query Q has been solved. This approach greatly simplifies the query evaluation process, but does still require a solution for metadata conversion.

3.2 Multiple Format Metadata Aggregator/Inference

In this model, the Aggregator is able to answer queries using previously loaded information (metadata and mapping information) from the external metadata-based systems that the Central subsystem supports.

The main characteristic of the Multiple Format Metadata Aggregator (MFMA) approach (see Figure 4) is that it is entirely based on ontologies and an ontology Reasoner [2]. Thus, the information that the Aggregator expects from every external system (and needs to be preloaded before receiving any query) is an ontology representing each external metadata format (F_1, F_2, ... , F_n) and another ontology that provides the mapping from a given external metadata format to the internal metadata format of the Aggregator (F_1 to F_A, F_2 to F_A, ... ,F_n to F_A). In addition to that, the Aggregator also needs to have the metadata information of the actual content that the user is querying for preloaded. This metadata is stored in a metadata database (MDB_A) in the form of ontology individuals (extracted from or provided by the external systems in the Metadata subsystem).

So, when the Aggregator receives a query, the Reasoner expands all the ontology information using the mappings provided and infers the results before sending them back to the user. Typically, the results will be direct links to the

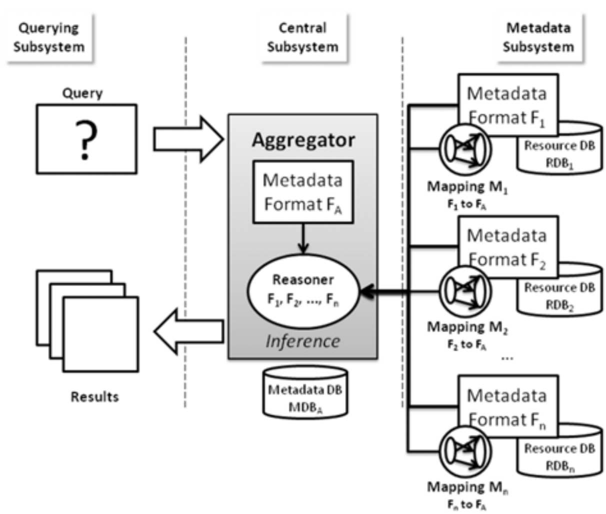

Fig. 4. Basic architecture for the "Multiple Format Metadata Aggregator" (MFMA) approach

external Resource Databases (RDB_1, RDB_2, ..., RDB_n) of the supported systems, where the actual content (i.e. images, videos, music...) is stored.

Examples of the MFMA approach include all the projects providing semantic search capabilities over Linked Data [7] related datasets. Linked Data is a way of exposing RDF metadata and interlinking them. In the years significant amounts of data have been generated, increasingly forming a globally connected, distributed data space. DBpedia [8] is an example for such a source in the Web of Data. In the case of multimedia there are examples such as the The Linked Movie DataBase (LinkedMDB) [16], which contains semantic metadata from several movie-related data sources, or DBtune [4], a collection of music-related data sets.

such as DBpedia [8]

3.3 Broker/Query Rewriting

The Broker acts as an intermediary agent communicating the system user to several external systems. The Broker approach is the only one that does not rely on the local storage of significant amounts of metadata (even though it needs to temporally store certain metadata for query reprocessing). It is based on the definition of a mediated schema, F_B, query rewritting and query reprocessing. This approach can be seen as a reformulation of some issues already faced by the Data Integration community, such as the classic *global-as-view* approach or GAV [1](TSIMMIS)[3][11] and the classic *local-as-view* approach or LAV [17][12][6]. However, there's a more recent definition of the Broker approach in [25]. Figure 4 shows the basic Broker architecture diagram.

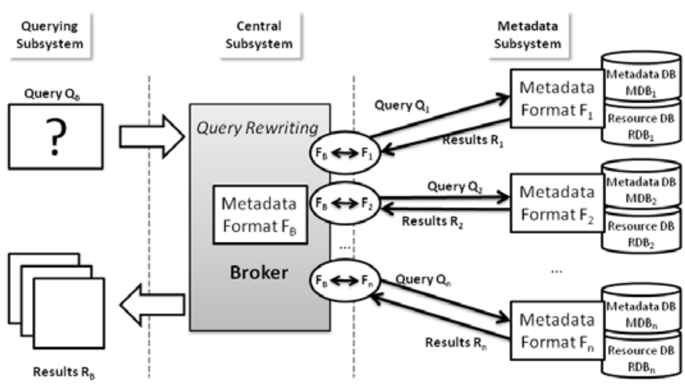

Fig. 5. Basic architecture for the "Broker" approach

The Broker receives a query (Q_B) in the metadata format F_B, it rewrites it once for every metadata format that is supported (F_1, F_2, ..., F_n) via the corresponding interface ($F_B \leftrightarrow F1$, $F_B \leftrightarrow F_2$, ..., $F_B \leftrightarrow F_n$). The queries Q_1, Q_2, ..., Q_n are generated and sent to every external system. The resulting data is processed back again by the corresponding interface, mapped back to the metadata format F_B and presented to the user. The local reprocessing of results allows answering queries addressing properties not available at the source's query interfaces but included in the result sets. When a source does not include a metadata field neither in the query interface nor in the result set the source cannot be used for a query addressing this metadata field and the user must be properly informed.

In the Broker approach, we have all the metadata databases (MDB_1, MDB_2, ..., MDB_n) and all the resource databases (RDB_1, RDB_2, .., RDB_n) at the metadata subsystem, so there is no need to maintain any kind of central database at the Central subsystem.

As we can see, from the architectural point of view this approach could be classified as a Broker based model, but from the implementation point of view we could talk about a Query Rewriting scenario [9]. The reason is that the main responsibility of the Broker system is to rewrite every query that receives and propagate it to all supported external systems.

3.4 Approaches Comparison

The Simple Format Metadata Aggregator is strongly related to the Metadata Conversion topic. All the incorporated metadata formats are converted to the internal metadata format of the aggregator before the system is ready to answer to user queries. This approach has probably a more static and deterministic nature in the sense that its behavior is easier to predict because of the absolute control of the metadata information that it offers. $SFMA$ systems are easy to design and implement and offer better performance than the other approaches. However, the need to harvest and locally store all the metadata can become a problem if the amount of metadata or their volatility is too big. This approach would probably be suitable for closed applications such as local multimedia repositories for multimedia players or file indexers.

The Broker approach requires a more complex design and implementation, it avoids the problems related to scalability and volatility. However, its performance is sensible to eventual problems (e.g. delays or unavailability) in the underlying databases. This approach would be suitable for distributing queries among big Internet image hosting services such as Flickr, Picassa or Panoramio.

The Multiple Format Metadata Aggregator is entirely based on the concept of Inference in ontologies through a reasoner. In this approach we also need to build a centralized metadata database in the Central subsystem. However, this time we do not need to perform metadata conversion before including the metadata information into the central database, the semantic correspondences are considered by the inference reasoner on-the-fly. $MFMA$ systems have small design complexity, being the main task the ontology mapping design. Besides, the inclusion of support new formats is straighforward, only requiring to be fed the knowledge base with the corresponding new ontologies and mappings.

4 Standards for Metadata Interoperability in Distributed Image Search&Retrieval

In this section we describe two metadata interoperability standards for image search&retrieval. The first one, MPQF, proposed by the MPEG working group and the second one, JPsearch, proposed by the JPEG committee.

4.1 ISO/IEC 15938-12:2008 Standard (MPEG Query Format or MPQF)

A key element in all the different approaches to distributed image search&retrieval is the interchange of queries and API calls among all the involved parties. The usage of different proprietary interfaces for this task makes extremely difficult the deployment of distributed image search services without degrading the query expressiveness. The progressive adoption of an unified query interface would greatly alleviate this problem. For its features, we conclude that the ISO/IEC 15938-12:2008 standard (MPEG Query Format or MPQF) is the most suited language for this purpose. MPQF is an XML-based query language that defines the format of queries and replies to be interchanged between clients and servers in a distributed multimedia information search&retrieval context. MPQF is an XML-based in the sense that all MPQF instances (queries and responses) must be XML documents. Formally, MPQF is Part 12 of ISO/IEC 15938, "Information Technology - Multimedia Content Description Interface" (MPEG-7 [23]). However, the query format was technically decoupled from MPEG-7 and it is now metadata-neutral. So, MPQF is not coupled to any particular metadata standard.

Example in Code1. shows an input MPQF query asking for JPEG images taken after the *2011/01/15* with the keyword *"Tokyo"* somwhere in their metadata.

4.2 ISO/IEC 24800 Standard (JPSearch)

The selection of a unified query interface is not enough to guarantee interoperability if it is not accompanied with a proper mechanism to manage metadata heterogeneity. The need of dealing with the management of metadata translations is a common factor in all the approaches to distributed image search&retrieval. Currently there is a standard solution to this problem proposed by the JPEG Committee, named JPSearch (ISO/IEC 24800). JPSearch provides a set of standardized interfaces of an abstract image retrieval framework. On one hand, JPSearch specifies the pivot JPSearch's Core Metadata Schema as the main component of the metadata interoperability strategy in ISO/IEC 24800. The core schema contains a set of minimal core terms which serve as metadata basis supporting interoperability during search among multiple image retrieval systems. The core schema is used by clients to formulate, in combination with the MPEG Query Format, search requests to JPSearch compliant search systems. In addition to the definition of JPSearch Core Metadata Schema, ISO/IEC 24800 provides a mechanism which allows a JPSearch

Code 1. Example MPQF input query

```
<MpegQuery>
  <Query>
    <Input>
      <OutputDescription>
        <ReqField>title</ReqField>
        <ReqField>date</ReqField>
      </OutputDescription>
      <QueryCondition>
        <TargetMediaType>image/jpg
        </TargetMediaType>
        <Condition xsi:type="AND">
          <Condition xsi:type="QueryByFreeText">
            <FreeText>Tokyo</FreeText>
          </Condition>
          <Condition xsi:type="GreaterThanEqual">
            <DateTimeField>date</DateTimeField>
            <DateValue>2011-01-15</DateValue>
          </Condition>
        </Condition>
      </QueryCondition>
    </Input>
  </Query>
</MpegQuery>
```

compliant system taking profit from proprietary or community-specific metadata schemas. A translation rules language is defined, allowing the publication of machine-readable translations between metadata terms belonging to proprietary metadata schemas and metadata terms in the JPSearch Core Metadata Schema. Users can choose which metadata language to use in a JPSearch-based interaction (annotation, querying, etc.) if the proper translations are available.

On the other hand, JPSearch specifies JPSearch Translation Rules Declaration Language (JPTRDL). JPTRDL allows the publication of machine-readable translations between metadata terms belonging to proprietary metadata schemas and metadata terms in the JPSearch Core Metadata Schema. Users can choose which metadata language to use in a JPSearch-based interaction if the proper translations are available. Code 2. shows a one-to-many translation rule which maps the JPSearch Core Schema date element into three fields.

4.3 JPSearch Registration Authority

According to the JPSearch specification, ISO/IEC 24800 compliant systems can manage multiple proprietary or community-specific metadata schemas, besides the JPSearch Core Metadata Schema. The multiplicity of schemas is solved by allowing the publication of machine-readable translations between metadata terms belonging to proprietary metadata schemas and metadata terms in the

Code 2. Example JPSearch translation rule

```
<?xml version="1.0" encoding="iso-8859-1"?>
<TranslationRules>
  <TranslationRule xsi:type="OneToManyFieldTranslationType">
    <FromField xsi:type="FilteredSourceFieldType">
      <XPathExpression>date</XPathExpression>
      <FilterWithRegExpr>(\d\d)/(\d\d)/(\d\d\d\d)</FilterWithRegExpr>
    </FromField>
    <ToField xsi:type="FormattedTargetFieldType">
      <XPathExpression>day</XPathExpression>
      <ReplaceWithRegExpr>£1</ReplaceWithRegExpr>
    </ToField>
    <ToField xsi:type="FormattedTargetFieldType">
        <XPathExpression>month</XPathExpression>
        <ReplaceWithRegExpr>£2</ReplaceWithRegExpr>
    </ToField>
    <ToField xsi:type="FormattedTargetFieldType">
        <XPathExpression>year</XPathExpression>
        <ReplaceWithRegExpr>£3</ReplaceWithRegExpr>
    </ToField>
  </TranslationRule>
</TranslationRules>
```

JPSearch Core Metadata Schema. In order to rationalize the usage of schemas and translation rules across different JPSearch systems, Subclause 3.3.3 of Part 2 of ISO/IEC 24800-2 specifies that a global authority for schemas and their translation rules will be established where all JPSearch compliant retrieval applications can obtain the information needed.

The establishment of a JPSearch Registration Authority (JPSearch RA) was formally approved during the 54th JPEG meeting in Tokyo, Japan, in February 2011, and will be operative in July 2011. The JPSearch RA will maintain a list of Metadata Schemas together with their related Translation Rules, if any. Those schemas and rules will be directly stored in the JPSearch RA web site or the JPSearch RA web site will provide a link to an external organization in charge of keeping that information updated. Registration forms will be available from the Registration Authority. Any person or organization will be eligible to apply. More information about the JPSearch RA can be obtained at www.iso.org/iso/maintenance_agencies/or directly at the JPEG home page (www.jpeg.org).

5 Example Real Distributed Image Search&Retrieval System

As a proof of concept, now we will describe an example real system that we have developed at the Distributed Multimedia Applications Group (DMAG, UPC BarcelonaTech). This experience is of special interest to the scope of this paper because it consists on a real implementation of the most complex approach, the only one which suits the requirements of a large scale web image hosting services aggregator, i.e. the Broker approach. The software we have developed provides a centralized place for searching images from Panoramio, Picasa and Flickr. Our system is compliant with the current version of the ISO/IEC 15938-12:2008 and ISO/IEC 24800 standards. Because the system follows the Broker architectural approach, there is a subsystem that receives MPQF queries addressing metadata in JPsearch format (extended with some image specific fields and with some EXIF fields, like the camera make and model) and rewrites them once for every metadata format that is supported (Panoramio, Picasa and Flickr). The system is split into two modules: a central application that implements a broker-based metadata aggregator and a web portal as the user front-end. Modules communicate each other by using MPQF queries.

Figure 6 shows a screenshot of the system's web front-end. In order to enable the maximum query expressiveness, the search form can be duplicated, as it is shown in this figure, allowing launching multi-queries. In fact, they will be all assembled with an OR operand and sent to the application using a single MPQF query, but the applications MPQF interpreter will split it into several searches, one by each form, and the broker will launch them in parallel.

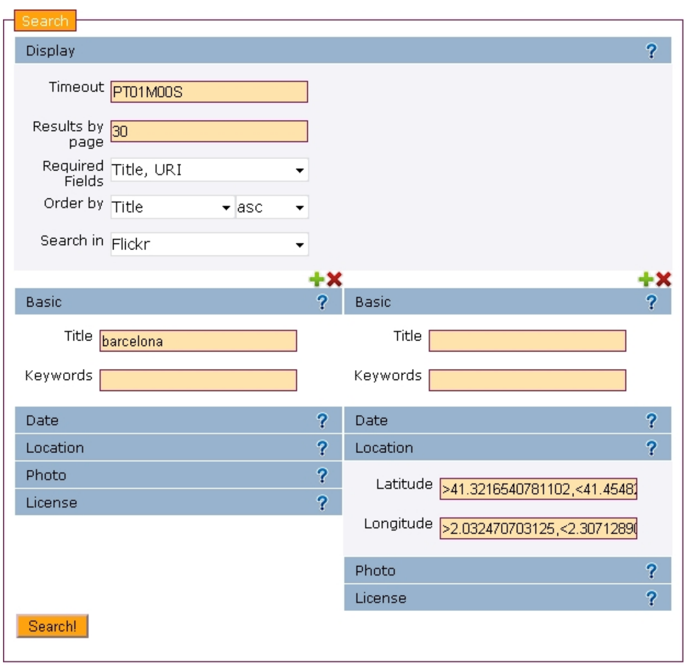

Fig. 6. Example web form

6 Conclusions

This paper has analyzed the current approaches and standards for metadata interoperability in distributed image search&retrieval. From the analysis of the main metadata schemas currently used for image description we conclude that they contain sigificant overlappings, but their specificity and constant evolution disallows any approach relying on a unified model. We have classified the several solutions that face this problem into three different approaches, the Simple Format Metadata Aggregator, the Multiple Format Metadata Aggregator and the Broker. We conclude that only the Broker approach suits the requirements of a large scale web image hosting services aggregator, while implying more complexity and performance constraints. Both the Simple Format Metadata Aggregator and the Multiple Format Metadata Aggregator suit situations with less challenges in terms of scalability and volatility, but offer better performance and are more simple to design and manage. We have also analyzed how a distributed image search&retrieval system, independently of the approach it uses, can be implemented on top of the ISO/IEC 15938-12 and ISO/IEC 24800 standards. Finally, we have described an example real distributed image search&retrieval system which provides real time access to Flickr, Picassa and Panoramio.

Acknowledgments. This work has been partly supported by the Spanish government (TEC2008-06692-C02-01).

References

1. Adali, S., Candan, K.S., Papakonstantinou, Y., Subrahmanian, V.S.: Query caching and optimization in distributed mediator systems. In: SIGMOD Conference, pp. 137–148 (1996)
2. Amann, B., Beeri, C., Fundulaki, I., Scholl, M.O.: Ontology-Based Integration of XML Web Resources. In: Horrocks, I., Hendler, J. (eds.) ISWC 2002. LNCS, vol. 2342, pp. 117–131. Springer, Heidelberg (2002)
3. Sudarshan, S., Chawathe, S.S., Hector, G.-M., Hammer, J., Ireland, K., Papakonstantinou, Y., Ullman, J.D., Widom, J.: The tsimmis project: Integration of heterogeneous information sources. In: IPSJ, pp. 7–18 (1994)
4. Dbtune - serving music-related rdf since (2007), `http://dbtune.org/`
5. Döller, M., Tous, R., Gruhne, M., Yoon, K., Sano, M., Burnett, I.S.: The MPEG Query Format: On the Way to Unify the Access to Multimedia Retrieval Systems. IEEE Multimedia 15(4) (2008); ISSN: 1070-986X
6. Duschka, O.M., Genesereth, M.R.: Answering recursive queries using views. In: PODS 1997: Proceedings of the Sixteenth ACM SIGACT-SIGMOD-SIGART Symposium on Principles of Database Systems, pp. 109–116. ACM Press, New York (1997)
7. Schandl, B., et al.: Linked data and multimedia: The state of affairs. Multimedia Tools and Applications, 1–34 (2011)
8. Bizer, C., et al.: Dbpedia - a crystallization point for the web of data. J. Web Sem. 7(3), 154–165 (2009)

9. Gergatsoulis, M., Bountouri, L., Gaitanou, P., Papatheodorou, C.: Query Transformation in a CIDOC CRM Based Cultural Metadata Integration Environment. In: Lalmas, M., Jose, J., Rauber, A., Sebastiani, F., Frommholz, I. (eds.) ECDL 2010. LNCS, vol. 6273, pp. 38–45. Springer, Heidelberg (2010)
10. Gruhne, M., Tous, R., Döller, M., Delgado, J., Kosch, H.: MP7QF: An MPEG-7 Query Format. In: Proceedings of the 3rd International Conference on Automated Production of Cross Media Content for Multi-Channel Distribution (AXMEDIS 2007), Barcelona, Spain, pp. 15–18 (November 2007)
11. Haas, L.M., Kossmann, D., Wimmers, E.L., Yang, J.: Optimizing queries across diverse data sources. In: Proceedings of the Twenty-third International Conference on Very Large Databases, pp. 276–285. VLDB Endowment, Athens (1997), http://citeseer.ist.psu.edu/article/haas97optimizing.html
12. Halevy, A.Y., Rajaraman, A., Ordille, J.J.: Querying heterogeneous information sources using source descriptions. In: Proceedings of the Twenty-Second International Conference on Very Large Databases, Bombay, India, pp. 251–262. VLDB Endowment, Saratoga (1996), http://citeseer.ist.psu.edu/levy96querying.html
13. ISO/IEC 15938-12:2008 Information Technology - Multimedia Content Description Interface - Part 12: Query Format (2008)
14. ISO/IEC 24800-3:2010 Information technology - JPSearch - Part 3: JPSearch Query format (2010)
15. ISO/IEC 24800-2:2011 Information technology - JPSearch - Part 2: Registration, identification and management of schema and ontology (2011)
16. The linked movie database (linkedmdb), http://www.linkedmdb.org/
17. Manolescu, I., Florescu, D., Kossmann, D.K.: Answering XML queries over heterogeneous data sources. In: Proceedings of the 27th International Conference on Very Large Data Bases, pp. 241–250 (2001)
18. Deviantart webpage, http://www.deviantart.com/
19. Flickr webpage, http://www.flickr.com/
20. Panoramio webpage, http://www.panoramio.com/
21. Photobucket webpage, http://photobucket.com/
22. Picasa webpage, http://picasa.google.com/
23. ISO/IEC 15938 Version 2. Information Technology - Multimedia Content Description Interface, MPEG-7 (2004)
24. Seligman, L., Mork, P., Halevy, A., Smith, K., Carey, M., Chen, K., Wolf, C., Madhavan, J., Kannan, A., Burdick, D.: Openii: an open source information integration toolkit. In: ACM Int. Conf. on Management of data (SIGMOD), pp. 1057–1059 (2010)
25. Tous, R., Delgado, J.: Advanced Meta-Search of News in the Web. In: 6th International ICCC/IFIP Conference on Electronic Publishing (elPub 2002), Karlovy Vary, Czech Republic (November 2002)