

# Interoperable Multimedia Annotation and Retrieval for the Tourism Sector

Antonios Chatzitoulousis, Pavlos S. Efraimidis, and Ioannis N. Athanasiadis<sup>(✉)</sup>

Electrical and Computer Engineering Department,  
Democritus University of Thrace, Xanthi, Greece  
achatzit@ee.duth.gr, pefraimi@ee.duth.gr, ioannis@athanasiadis.info

**Abstract.** The Atlas Metadata System (AMS) employs semantic web annotation techniques in order to create an interoperable information annotation and retrieval platform for the tourism sector. AMS adopts state-of-the-art metadata vocabularies, annotation techniques and semantic web technologies. Interoperability is achieved by reusing several vocabularies and ontologies, including Dublin Core, PROV-O, FOAF, Geonames, Creative commons, SKOS, and CiTO, each of which provides with orthogonal views for annotating different aspects of digital assets. Our system invests a great deal in managing geospatial and temporal metadata, as they are extremely relevant for tourism-related applications. AMS has been implemented as a graph database using Neo4j, and is demonstrated with a dataset of more than 160000 images downloaded from Flickr. The system provides with online recommendations, via queries that exploit social networks, spatiotemporal references, and user rankings. AMS is offered via service-oriented endpoints using public vocabularies to ensure reusability.

## 1 Introduction

Data shared on the Internet are typically not accompanied with useful metadata to enable efficient search and discovery. Even in the cases meta-information is present, it comes in custom formats, not using standard dictionaries, nomenclatures or organization. Semantic Web [9] and Linked Data [10] attempt to resolve such problems by adding rich meta information for both machines and humans, while creating and adopting standards and organization in sharing data on the Web. In the past years we have observed the creation and adoption of several templates for sharing data that include dictionaries and ontologies for describing information about data, along with semantic services that allow for the search and discovery of data on the web, such as DAML-S/UDDI Matchmaker [20], KIM [12], or OWLS-MX [13].

This work contributes to the realization of the Semantic Web by designing and implementing a metadata annotation system for tourism-related applications, called Atlas Metadata System (AMS for short), which consists of a metadata schema for multimedia content, that reuses common terms from existing vocabularies and ontologies. The system deployed enables interoperability not

only by reusing existing terminology, but also by adopting services for search and discovery on the Semantic Web. On top of that, we demonstrate that by the annotation of user network and preferences, we can exploit *social network* dynamics for improved recommendations. This work aims to investigate interoperability in sharing multimedia content in the case of tourism-related applications, but also contributes to the demonstration of systems for open, reusable data on the web.

### 1.1 Problem Definition and Challenges

With this work, we try to overcome several general open issues in sharing content online along with particular shortfalls related to the tourism sector. Below we outline those issues in brief.

**Lost in the Data Ocean:** We are overflowed with an abundance of data that do not come together with useful meta-information. There are so many photographs and videos on the Web that could be potentially interesting for tourism related applications. However, most of the time they come along with low-level metadata that are concerned with the file name, size or the date of creation. They lack information about the content itself, as a title, a description, or any kind of classification or tags. There are several tools developed to automatically tagging multimedia content related to tourism, as in [15]. At the same time, information about the creator, licensing and sharing principles are also typically absent. Geospatial information is sometimes available, but not always in ways that makes the content reusable. For example, to make multimedia content directly relevant to tourism applications, relationships to *Points of Interest* are needed.

**Limited Standard Adoption:** There is an anarchy on which metadata dictionaries to adopt. It is common practice to develop new metadata schemas in isolation, instead of looking for relevant existing nomenclatures. This is not just a matter of bad practice that which results into problems of ontology alignment. It involves several difficulties in drawing boundaries between sectors, and lots of community work to reach consensus. W3C and OGC communities have produced relevant dictionaries and practices for sharing Web and Geospatial Information, but there are several other overlapping domain specific efforts.

**Discovery for Machines, not Humans:** Still most search and discovery tools are made for humans and require interpretation. Service oriented technologies allow for machine to machine transactions that together with public Application Programming Interfaces (API) may enable for advanced tourism contextual recommendations.

In this work we address those challenges in the context of tourism recommendation systems. The Atlas Metadata System (AMS) aims to be an interoperable platform that reuses standard vocabularies for annotating multimedia metadata and allow for search and discovery via open services.

## 1.2 Related Work

There have been several metadata models proposed for annotating multimedia content related to mobile photography and video, or more specifically tourism related applications. Already in 2002, Maedche and Staab envisioned scenarios on how Semantic Web technologies can be used for next-generation tourism information systems [14]. The issue of designing an appropriate metadata schema for tourism related applications is recurring. A straightforward schema has been proposed in [18], which consists of three main entities people, events, and places. In the same work, annotation is considered a side product of some other user activity, with her mobile device, like photo sharing or social discourse. In 2006, Torniai *et al.* presented a system which employs RDF for the description of photographs, including location and compass heading information, and is used to discover geo-related pictures from other users [22]. Most of the metadata schema follows Dublin Core. PhotoMap [23] focuses on tours/itineraries and encompasses five context dimensions for a metadata schema in OWL-DL: spatial, temporal, social, computational and spatiotemporal. These concepts correspond to the major elements for describing a photo (i.e., where, when, who, with what). Kanellopoulos and Panagopoulos in [11] propose a metadata model in RDF to annotate information about tourist destinations and peers. Last, but not least, a methodology that allows for the automatic population of ontologies on the basis of natural language processing is proposed in [17]. In particular, with their approach, a given ontology can be enriched by adding instances gathered from natural language texts.

Here we build upon previous work, while concentrating mostly on existing image repositories and how their metadata can be exploited for tourism related applications. Our results show that reuse of existing vocabularies together with service oriented technologies are the key for interoperability.

## 2 Atlas Metadata System

### 2.1 Goals and Systems Specification

The Atlas Metadata System (AMS) employs semantic web annotation techniques in order to create an interoperable information annotation and retrieval platform for the tourism sector. The Atlas Metadata System is equipped with a metadata schema that combines both anagraphic user information and tourism related classifications, spatial and temporal references, social networks and user preferences. More importantly, the schema adopted in AMS reuses to the greatest extend existing vocabularies for expressing such relationships.

Specifically, the proposed system aims to:

- Fuse existing content from public repositories, as Flickr and Panoramio, and enhance their metadata, with data from other sources, to become more relevant and useful for tourism recommendations.

- Allow for improved recommendations by exploiting social networks and spatial and temporal relationships.
- Make multimedia metadata available as web services using standard protocols and vocabularies.

Using AMS services, users (either machines or humans) can search and discover for tourism-related metadata by employing appropriate ontologies; e.g. Geonames dictionary for spatial information, Dublin Core for general information, FOAF for social relationships, etc.

AMS is structured in two modules. The first module operates offline and is responsible to extract data from existing APIs and databases to populate AMS repository. This involves scripting against legacy systems or popular services such as Flickr. At the end of this step, nodes have been created in the graph database; identities of content, users and Points of Interest have been resolved; and spatial indexes have been encoded.

The second component is the online system that serves the metadata through various interfaces, such as SPARQL, RESTfull services, etc. Users (humans or machines) though the online component ask for recommendations on tourism-related multimedia, either based on their profiles, history, current or intended location.

Both components share a common metadata schema that has been designed with reusing existing vocabularies.

## 2.2 AMS Metadata Schema

We identify three main entities in our system:

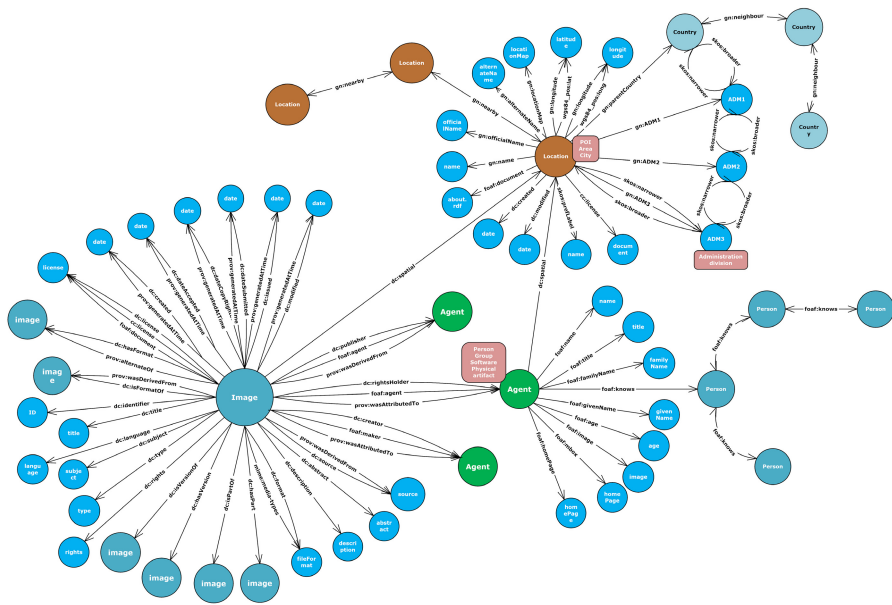
- **Image:** These are photographs taken by users. It is the main entity we want to annotate.
- **Agent:** An entity that owns, created, likes or published an Image. An agent may be a person, an organization, or even an artifact. Points of Interests may be agents too.
- **Location:** This is the place where a photograph was taken, or an agent resides. It is related with other geographical entities through spatial relations.

Relations between those entities and their attributes are expressed using existing vocabularies, as discussed below. An overview of the attributes of each entity is illustrated in Figure 1.

## 2.3 Vocabularies Reused

In order to maximize the potential for interoperability, we reused several existing terms for designing the AMS schema. Below we introduce the various dictionaries employed, and list some of the terms we reused from each one.

Dublin Core is a generic purpose metadata schema for annotating digital artifacts [2]. We employed several terms from this scheme to represent the



**Fig. 1.** The metadata schema proposed. The main entities and their attributes are shown.

information about our items, their creation history, licensing, versions spatiotemporal coverage and theme classifications, using `dc:abstract`, `dc:created`, `dc:creator`, `dc:dateAccepted`, `dc:dateCopyrighted`, `dc:dateSubmitted`, `dc:description`, `dc:format`, `dc:hasFormat`, `dc:hasPart`, `dc:hasVersion`, `dc:identifier`, `dc:isFormatOf`, `dc:isPartOf`, `dc:isVersionOf`, `dc:issued`, `dc:language`, `dc:license`, `dc:modified`, `dc:publisher`, `dc:rights`, `dc:rightsHolder`, `dc:source`, `dc:spatial`, `dc:subject`, `dc:temporal`, `dc:title`, `dc:type`.

PROV-O is a W3C recommendation [7] for representing and interchanging provenance information generated in different systems and under different contexts. We reused the following terms `prov:generatedAtTime`, `prov:wasAttributedTo`, `prov:alternateOf`, and `prov:wasDerivedFrom`.

FOAF is vocabulary for describing and linking persons and their relations [3]. We used it to represent the social relationships between our users, with `foaf:agent`, `foaf:person`, `foaf:knows`.

Geonames [4] is an ontology that provides unique identifiers for over 8.3 million toponyms listed in geonames.org database and their spatial relations. In Atlas we employed Geonames to attribute geographical entities to our media using: `gn:name`, `gn:alternateName`, `gn:officialName`, `gn:parentCountry`, `gn:ADM1`, `gn:ADM2`, `gn:ADM3`, `gn:locationMap`, `gn:latitude`.

SKOS, the Simple Knowledge Organization System, is a W3C recommended common data schema for knowledge organization systems such as thesauri, classification schemes, subject heading systems and taxonomies [6]. We employed SKOS for representing the narrower/broader relationships of classification terms, with `skos:narrower`, `skos:broader`, `skos:preferredLabel`

CiTO, originally an ontology for scholar citations [19], has been reused in our work for its capacity to express *like* relationships, via the property `cito:likes`.

Creative Commons [8] also maintains a machine readable dictionary for content licensing which we adopted, via `cc:license`.

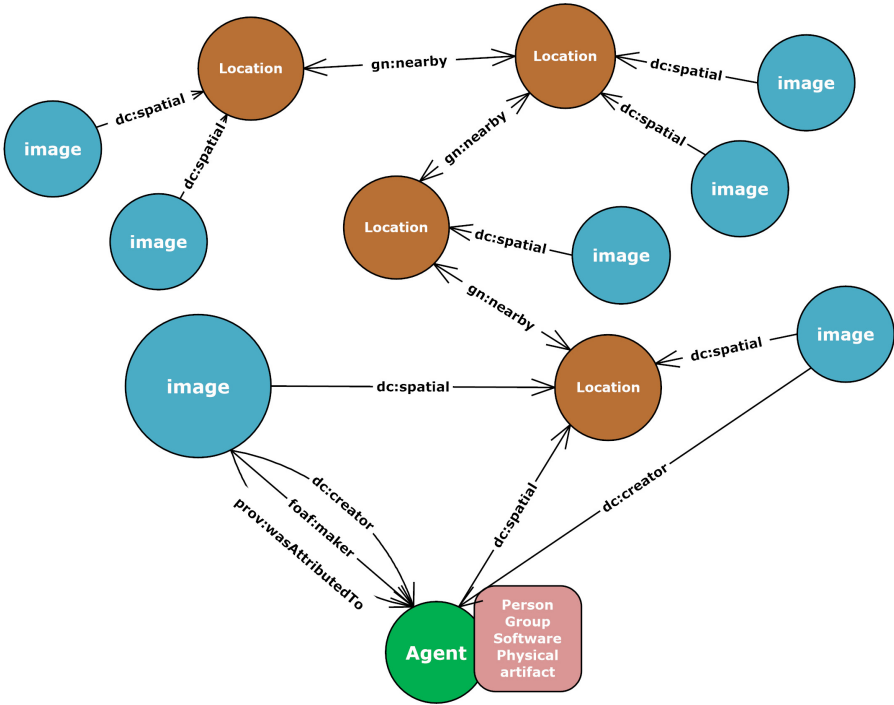


Fig. 2. Relations between entities from various dictionaries

Table 1 illustrates the synonym relations that are provided by different vocabularies for the image nodes, and Figure 2 illustrates some of those relations as a graph. Several different public dictionaries can be used to traverse the graph and find relationships in Atlas.

### 2.4 Graph Relationships

A feature of the AMS is that entities may be related to entities of the same class, creating linked graphs. This is the case of *Locations* that are linked to each other with the `gn:nearby` property, creating a spatial graph. The same

**Table 1.** Some of the synonym relations for `Image` nodes.

Domain Properties	Range
<code>Image</code> <code>dc:creator</code> , <code>prov:wasAttributedTo</code> , <code>foaf:maker</code>	<code>Person</code>
<code>Image</code> <code>dc:created</code> , <code>prov:generatedAtTime</code>	<code>date</code>
<code>Image</code> <code>dc:license</code> , <code>cc:license</code>	<code>license</code>

holds for Agent nodes that are linked via `foaf:knows` to other agents. These relations eventually create a social graph that can be exploited in our prototype for retrieving higher order relationships between locations or agents, by using graph databases principles [16].

## 2.5 Implementation

AMS is implemented with several state-of-the-art semantic web technologies. Specifically, we employed Neo4j [5] to deploy AMS metadata storage. Neo4j allows for rapid traversals of the AMS metadata registry, and supports for RDF/SPARQL endpoints. This seems as a natural choice for tourism systems where recommendations will come from spatial or social graphs via a graph querying language as Cypher, rather than ontological reasoning or geoprocessing computations. At the same time, AMS metadata can become available as REST-full web services using various end points of Neo4j, including Cypher query language, SPARQL, and Neo4j spatial plugin.

## 3 Demonstration

### 3.1 Experimental Setup

To demonstrate AMS we used the Image Collection used in the ATLAS project [1]. The collection consists of 162'583 images geolocated in Greece, originally published on Flickr. Through the Flickr API we extracted metadata about the images and their 28'358 contributors, including their social network on flickr. We enhanced Flickr metadata by developing and applying some scripts which identified entities, added extra relationships from tags, created spatial and temporal relationships, resolved toponyms with Geonames, and added Creative Commons licenses. The AMS graph database is powered with Neo4j, occupies more than 600 MB of metadata, includes enhanced features, and publishes metadata using standard vocabularies.

### 3.2 Simple Queries with Cypher

Using the Cypher query language we can express queries using alternative, equivalent vocabularies, to demonstrate the AMS capacity for interoperability. In the example below we retrieve ten images of a certain user, using the `dc:creator` relationship. The same query could be expressed via the FOAF ontology using `foaf:maker`.

```
START n=node:owner(owner_id = "10179878@N03")
MATCH (n)-[r:'dc:creator']-(m) RETURN n,r,m LIMIT 10
```

Graph databases are best suited for graph traversals. In the following example, the indirect connections are extracted from a social graph. Specifically, we ask for higher order friends-of-a-friend of a certain user, and sort them by the number of the different paths between them.

```
START n=node:owner(owner_id = "10179878@N03")
MATCH (n)-[:'foaf:knows'*2..2]-(friend_of_friend)
WHERE (NOT (n)-[]-(friend_of_friend))
RETURN friend_of_friend, COUNT(*)
ORDER BY COUNT(*) DESC
```

“Like” relationships may be exploited for discovering users with similar interests. In the example below we ask for users that share similar ratings in common photographs. Here we exploit `cito:likes` and `foaf:knows` semantics to represent the underlying relationships.

```
MATCH (me:Owner { owner_id: '12337376@N03' }) -
      [ml:'cito:likes']->(image:Image)<-[pl:'cito:likes'
      - (person:Owner)
WHERE NOT (me)-[:'foaf:knows']->(person) and
      abs(pl.rating - ml.rating) < 2
RETURN person.username,image.title,ml.rating,pl.rating
```

### 3.3 Open Querying Endpoints via Plugins

An important features of Neo4j is that it comes with a handful set of plugins, that support for machine interoperability with several protocols, which maximize our system capacity for reuse. First is the Neo4j REST-full API that allows to directly interact with the database via the http protocol. All the example queries of the previous section may be directly executed via the transactional Neo4j API. A simple query to retrieve all properties of a certain node may look like this:

```
GET http://ams.example.org/db/data/node/415508
```

The SPARQL plugin allows for semantic web interoperability. SPARQL is a W3C standardized RDF query language, and is recognized as one of the key technologies for Linked Data and the Semantic Web. A sample query in AMS is the following one that requests for subjects and objects of the relations `foaf:knows`, as:

```
POST http://ams.example.org/db/data/ext/SPARQLPlugin/graphdb/
      execute_sparql
{
  "query" : "SELECT ?x ?y
            WHERE { ?x <http://xmlns.com/foaf/spec/#term_knows> ?y . }"
}
```



Last but not least, the Spatial plugin allows for geospatial queries. For example, one may retrieve all images within a distance of 0.5km around a certain coordinate via the http protocol, as:

```
POST http://ams.example.org/db/data/ext/SpatialPlugin/graphdb/
      findGeometriesWithinDistance
Content-Type: application/json
{
  "layer" : "images",
  "pointY" : 40.626340,
  "pointX" : 22.948362,
  "distanceInKm" : 0.5
}
```

### 3.4 Advanced Recommendations

AMS is capable to respond efficiently to more advanced queries, which can be used for contextual suggestions to its users. For example, we want to make image suggestions based on the ratings of users with similar interests. For this we employ the cosine similarity measure [21] with the  $k$ -nearest neighbors algorithm to make image recommendations, similar to [24]. Cosine similarity is a distance measure for estimating the similarities in image ratings, and with the  $k$ -nearest neighbors algorithm, the  $k$  closest images are selected according to this distance. An implementation in AMS system would look like:

```
MATCH (b:Owner)-[r:'cito:likes']->(i:Image), (b)-[s:SIMILARITY]
      -(a:Owner {username:'Antonis Chatzitoulousis'})
WHERE NOT((a)-[:'cito:likes']->(i))
WITH i, s.similarity AS similarity, r.rating AS rating
ORDER BY i.title, similarity DESC
WITH i.title AS image, COLLECT(rating)[0..3] AS ratings
WITH image, REDUCE(s = 0, j IN ratings | s + j)*1.0 /
      LENGTH(ratings) AS reco
ORDER BY reco DESC
RETURN image AS Image, reco AS Recommendation
```

The corresponding results screen of Neo4j graph database is shown in Figure 3. The same method could be followed for suggestions on various Points of Interests, including attractions, cities, categories, etc, by recommending POIs that have been rated by users with similar profiles and have not been visited yet, by exploiting user profiles. Recommendations can be improved by including the spatiotemporal aspects, by considering the locations where images have been taken, or periods of time. For example, search for photographs of POIs close to Parthenon in December based on the ratings of users with similar interests. Temporal aspects may consider museum/office hours, weekends, periods of festivities etc. Spatial aspects may consider geographical regions for focal search. The social network can be explored further for giving suggestions by limiting the  $k$ -nearest neighbors algorithm to the friends-of-a-friend network.

Query: MATCH (b:Owner)-[:likes]->(i:Image), (b)-[:SIMILARITY]-(:a:Owner {username:'Katerina Mal...})

Image	Recommendation
Θεσσαλονίκη-Παραλία	7.333333333333333
Καρνάγιο , Θεσσαλονίκη .	5.666666666666667
Καστοριά - Παγκάκι δίπλα στην λίμνη / Kastoria by the lake	5.666666666666667
Καλή Πρωτομαγιά! - Have a fantastic May Day!	5
Κνωσός - Knossos, one of the oldest street of Europe	4.666666666666667
Θαλάσσια χελώνα - Sea Turtle Caretta Caretta	2.666666666666667

Returned 6 rows in 755 ms.

Fig. 3. Image suggestions that were highly-rated from users with similar interests

## 4 Conclusions and Future Work

This work illustrated how AMS, an advanced metadata system for the tourism sector may be implemented to overcome common problems in sharing Open Linked Data on the Semantic Web. Interoperability is achieved through multiple technologies and interfaces deployed, allowing for option in both human and machine interaction. Reuse of common dictionaries and ontologies ensures that a set of options are available to the various clients, that do not need to be adapted to a custom metadata schema. Finally, the use of social network metadata enables advanced recommendations that exploit user profiles and preference. Another lesson learned by this exercise is that graph databases are both mature and robust technology that can be used effectively for metadata management, esp in the case of larger repositories. Semantic Web technologies are adequately supported and the service-oriented paradigm allows for easy integration, adaptive management and fast responses.

Future extensions for this work may exploit geospatial services for visualizing results using Google Earth, OpenLayers, and similar tools. Another line of action could consider matching user profiles with other social platforms, through Twitter or Facebook social graphs.

**Acknowledgments.** Research in this study has received funding by European and national funds from NSRF 2007-2013, OP Competitiveness and Entrepreneurship, Cooperation 2011, in the context of the *Advanced Tourism Planning System (ATLAS)* project, grant number 11SYN-10-1730.

## References

1. Advanced tourism planning system (ATLAS). <http://atlas.web.auth.gr/>
2. Dublin Core Metadata Initiative (DCMI) Metadata Terms. <http://dublincore.org/documents/dcmi-terms/>
3. FOAF Vocabulary Specification. <http://xmlns.com/foaf/spec/20140114.html>
4. GeoNames Ontology. <http://www.geonames.org/ontology/>
5. Neo4j, the World's Leading Graph Database. <http://neo4j.com/>
6. SKOS Simple Knowledge Organization System Reference. W3C Recommendation (2009). <http://www.w3.org/TR/2009/REC-skos-reference-20090818/>
7. PROV-O: The PROV Ontology. W3C Recommendation (2013). <http://www.w3.org/TR/prov-o/>
8. Abelson, H., Adida, B., Linksvayer, M., Yergler, N.: CC REL: the creative commons rights expression language. In: de Rosnay, M.D., Martin, J.C.D. (eds.) *The Digital Public Domain - Foundations for an Open Culture*, pp. 149–187. OpenBook Publishers (2012)
9. Berners-Lee, T., Hendler, J., Lassila, O., et al.: The semantic web. *Scientific American* **284**(5), 28–37 (2001)
10. Bizer, C., Heath, T., Idehen, K., Berners-Lee, T.: Linked data on the web (LDOW2008). In: *Proceedings of the 17th International Conference on World Wide Web*, pp. 1265–1266. ACM (2008)
11. Kanellopoulos, D.N., Panagopoulos, A.A.: Exploiting tourism destinations' knowledge in an RDF-based P2P network. *Journal of Network and Computer Applications* **31**(2), 179–200 (2008)
12. Kiryakov, A., Popov, B., Terziev, I., Manov, D., Ognyanoff, D.: Semantic annotation, indexing, and retrieval. *Web Semantics: Science, Services and Agents on the World Wide Web* **2**(1), 49–79 (2004)
13. Klusch, M., Fries, B., Sycara, K.: OWLS-MX: A hybrid semantic web service matchmaker for OWL-S services. *Web Semantics: Science, Services and Agents on the World Wide Web* **7**(2), 121–133 (2009)
14. Maedche, A., Staab, S.: Applying semantic web technologies for tourism information systems. In: *Information and Communication Technologies in Tourism (ENTER 2002)*. Springer (2002)
15. Pliakos, K., Kotropoulos, C.: PLSA-driven image annotation, classification, and tourism recommendations. In: *Proceedings of IEEE International Conference on Image Processing (ICIP)*, Paris, France, October 2014
16. Robinson, I., Webber, J., Eifrem, E.: *Graph Databases*. O'Reilly (2013)
17. Ruiz-Martinez, J.M., Minarro-Gimenez, J.A., Castellanos-Nieves, D., Garcia-Sanchez, F., Valencia-Garcia, R.: Ontology population: an application for the e-tourism domain. *International Journal of Innovative Computing, Information and Control* **7**(11), 6115–6134 (2011)
18. Sarvas, R.: User-centric metadata for mobile photos. In: *Pervasive Image Capture and Sharing Workshop at Ubicomp 2005*. Citeseer (2005)
19. Shotton, D.: CiTO, the citation typing ontology. *J. Biomedical Semantics* **1**(S-1), S6 (2010). <http://vocab.ox.ac.uk/cito>
20. Sycara, K., Paolucci, M., Ankolekar, A., Srinivasan, N.: Automated discovery, interaction and composition of semantic web services. *Web Semantics: Science, Services and Agents on the World Wide Web* **1**(1), 27–46 (2003)
21. Tan, P.N., Steinbach, M., Kumar, V.: *Introduction to Data Mining*. Addison-Wesley (2005)

22. Torniai, C., Battle, S., Cayzer, S.: Sharing, discovering and browsing photo collections through RDF geo-metadata. In: Proceedings of the 3rd Italian Semantic Web Workshop, SWAP 2006. CEUR Workshop Proceedings, vol. 201 (2006)
23. Viana, W., Filho, J.B., Gensel, J., Villanova Oliver, M., Martin, H.: PhotoMap – automatic spatiotemporal annotation for mobile photos. In: Ware, J.M., Taylor, G.E. (eds.) W2GIS 2007. LNCS, vol. 4857, pp. 187–201. Springer, Heidelberg (2007)
24. White, N.: Movie recommendations with k-nearest neighbors and cosine similarity, Graph Gist winter challenge winners. <http://gist.neo4j.org/?8173017>