

# Mobile cultural heritage guide: location-aware semantic search

Chris van Aart<sup>1</sup>, Bob Wielinga<sup>1,2</sup> and Willem Robert van Hage<sup>1</sup>

<sup>1</sup> VU University Amsterdam, the Netherlands

<sup>2</sup> Universiteit van Amsterdam, the Netherlands

cj@vanaart.com, b.j.wielinga@uva.nl, wrvhage@few.vu.nl

**Abstract.** In this paper we explore the use of location aware mobile devices for searching and browsing a large number of general and cultural heritage information repositories. Based on GPS positioning we can determine a user's location and context, composed of physical nearby locations, historic events that have taken place there, artworks that were created at or inspired by those locations and artists that have lived or worked there. Based on a geolocation, the user has three levels of refinement: pointing to a specific heading and selection and facets and sub-facets of cultural heritage objects. In our approach two types of knowledge are combined: general knowledge about geolocations and points of interest and specialized knowledge about a particular domain, i.e. cultural heritage. We use a number of Linked Open Data sources and a number of general sources from the cultural heritage domain (including Art and Architecture Thesaurus, Union List of Artist Names) as well as data from several Dutch cultural institutions. We show three concrete scenarios where a tourist accesses localized information on his iPhone about the current environment, events, artworks or persons, which are enriched by Linked Open Data sources. We show that Linked Open Data sources in isolation are currently too limited to provide interesting semantic information but combined with each other and with a number of other sources a really informative location-based service can be created.

## 1 Introduction

In this paper we explore the use of location aware mobile devices for searching and browsing large collections of general and cultural heritage information repositories using minimal interaction. Given a particular geolocation we provide cultural heritage resources for an end user. The material origins from the MultimediaN E-Culture project which deployed large virtual collections of cultural-heritage resources [7]. These resources are imbedded in the Linked Open Data (LOD) cloud [6].

Current smart phones such as the iPhone, Blackberry, HTC or Android have continuous access to internet, know about their geographic location and even know what direction the user is looking at. These capabilities are being used for a number of applications that show the user a map of his/her current location

with possible places of interest marked on the map (Linked Geo Data browser, Google Maps, Layar, WikiTude, Mobile DBpedia [4]) or they provide the user with detailed information about a particular aspect of the current location, such as interesting architectural structures to be seen. These applications use two categories of knowledge: general knowledge about geolocations and points of interest (POIs), and/or specialized knowledge about a particular domain. The first category of knowledge is present in the LOD cloud, the second category of knowledge may be available from sources not represented in the LOD cloud. Google Maps (and applications based on Google Maps) particularly show POIs and links to a website on a map, but does not provide related specialized knowledge related to a POI.

In our approach these two types of knowledge are combined: general knowledge about geolocations and points of interest (as represented in GeoNames, LinkedGeodata, Freebase and DBpedia) and specialized knowledge about the cultural heritage domain (Art and Architecture Thesaurus, Union List of Artist Names, Thesaurus of Geographic Names) as well as data from several cultural institutions (Netherlands Institute for Art History and Rijksmuseum Amsterdam).

The challenges that we have to resolve to reach the goal are: to enrich location data by constructing an “enriched local map” of nearby Points of Interest enriched with additional information, such as e.g. events, persons and artworks. Next we find a way to present this information to a user on a mobile device, taking into account the constraints of a mobile device and the limited span of attention of that user. We show three concrete scenarios where a tourist can access localized information on his iPhone about locations, artworks, events and persons.

## 2 Domain: Tourist Guides and Cultural Heritage

The profession of tourist guide is almost as old as tourism and is defined as: *a person who guides visitors in the language of their choice and interprets the cultural and natural heritage of an area.* . . [13]. Those who cannot afford a guide, or those who want to explore on their own can make use of guide books, such as provided by the companies: “Lonely Planet” and “Rough Guides”. The self-made tourist or the active tourist finds satisfaction in the process of composing his/her own program for the day [5]. Guide books include information about hotels, restaurants, travel, city life (e.g. culture, economy, environment, etc.), arts (literature, theater, music, cinema, etc.), architecture (e.g. building styles), history and walking tours. When actually visiting a foreign place, the active tourist has questions such as: “What do I see?”, “How did artists look at this location?”, “What is the history?”, “What kind of stories are related?”, “Which events have taken place?”, “Which persons were involved in this place?”, “What is my next stop?”, etc.

Current smart phone and internet technology has the power of providing answers on these questions in the form of digital tourists guides. A lot of these

applications deal with finding locations of interest nearby and guide navigation, e.g. TomTom, Garmin and Navico. Most of these applications rely on own proprietary maps or on public sources such as Google Maps or Open Street Map.

At this moment there are a few mobile applications that make use of the Semantic Web, e.g. DBpedia mobile [4]. However, for enriched storytelling, one needs fast searching mechanism for selecting information and presentation format relevant to the user, based on his/her preferences and the current context [5]. An example is Google Goggles for searching the web, based on pictures from a unified picture library [10], but this application mainly provides names of places of interest, not background information.

The MultimediaN E-Culture project has harvested 200,000 objects from six collections (including Netherlands Institute for Art History and Rijksmuseum Amsterdam) about the cultural heritage of Amsterdam [7]. These collections include digital representations of oil paintings, photographs, artists styles and artists information. These object are annotated with a range of thesauri and proprietary controlled keyword lists adding up to 20 million triples. Several Semantic Web technologies (such as lexical analysis, several conversions, enrichments, alignments) and ontologies (AAT, ULAN, TGN) are applied to convert all this data in to a consistent RDF representation. This is stored in the RDF store of the Semantic Web search engine *ClioPatria* [7]. *ClioPatria* can be accessed via SPARQL and a JSON-REST API. The aim of the current paper is to show how a combination of data from the LOD cloud combined with the E-Culture data can provide interesting, in-depth information about a certain location. A comparable project is SMARTMUSEUM (<http://smartmuseum.eu/>).

### 3 Concepts: Mobile Tourist Guide

Day trips and walking tours described in printed sources, such as the Lonely Planet and Rough Guides, are rather static. We envision a mobile Tourist Guide application able to dynamically combine navigation, information provision and a form of entertainment: *navitainment*. Based on a geolocation and filtering criteria given by the tourist, the app can constructs a dynamic walking tour [1]. Typical cultural filtering criteria are: architecture, paintings (how are artists inspired by a geolocation), photographs (capturing of historical moments), historical locations, etc. The idea is that the tourist starts with an initial criterion, e.g. paintings and can alter his criterion during the tour.

In order to construct dynamic tour guides, we need semantic annotated geolocations. In order to navigate the tourist we need intuitive ways of representing navigation data. Typically for a overview of POIs we can use a table, where each rows describes the name (e.g. Van Gogh Museum), type (e.g. Museum) and distance (e.g. 350m). To navigate we can use a map, showing the current location of the tourist and the path to the selected POIs, see Fig. 1. For this we have selected facets of cultural heritage objects (location, event, artwork or people) and subfacets (e.g.: painting, photograph, book, artist, musician, politician, sport



Fig. 1. Impressions of table, map and augmented reality-based interfaces.

and conflict). Next we need techniques to present the annotated data, such as background descriptions, representations of paintings, art and photographs.

#### 4 Approach: from Geolocation to Real-world Annotation

Our approach is composed of a number of steps. Figure 2 shows the task structure implemented in the system. In the first subtask (“harvest locations”) we gather data about locations nearby the user’s current location. This results in a set of RDF triples about nearby locations. The second step (“merge and align”) is to identify sets of triples that describe the same location. The result is a reduced set of unique locations. Next, for each unique location a semantic enrichment task is performed that searches various sources (a.o. the Dutch Wikipedia and the Eculture data cloud) to find additional information such as events, persons, artworks etc associated with the location. Finally each enriched location is classified in terms of the facet hierarchy. The result is a set of RDF statements that can be sent to a mobile device. Below we will describe the four subtasks in more detail.

Besides the Eculture data cloud and an RDF database about Dutch historical buildings, the system uses the ontologies of the Linkedgeodata initiative (LGDV), the ontology of DBpedia and a set of mapping rules. In total the database consists of almost 12M triples. RDF statements from LOD sources, Wikimapia and Wikipedia are retrieved on line using various server API’s. This process can be somewhat slow (for the Spui, a square in Amsterdam, the entire process takes some 50 seconds). This is not a big problem when the user sends a request to the server when approaching the location of interest. Furthermore,




Fig. 2. Task structure

intermediate results can directly be shown, while processing happens in the background. The performance of the communication between back-end and the iPhone is related to the quality of the Internet connection.

#### 4.1 Harvesting Nearby Locations




Fig. 3. Method: Harvest Locations

Figure 3 shows the reasoning process performed to harvest nearby locations. We start with a geolocation, represented by a point received from a mobile device:  $s = \langle lat, long \rangle$  (e.g.  $s = \langle 52.3638611, 4.88944 \rangle$  for the Spui square in Amsterdam). Using  $s$ , we determine the ontological characterization of the surroundings of the user's current location, such as the features found in geo-knowledgeable LOD repositories such as GeoNames and LinkedGeoData, while using relations such as `owl:sameAs`, `skos:exactMatch` and `skos:closeMatch` properties in the gathered RDF to obtain information about the entire equivalence classes of the nearby features. This includes crawling DBpedia entries. The crawling is done with the space package's `space_crawl_url` predicate [3]. In addition to the data harvested from the LOD sources, we use WikiMapia<sup>3</sup> which not only offers point coordinates, but also polygon and line information about locations such as buildings and streets. Wikimapia also provides links to Wikipedia pages in various languages. These links are followed using the crawling engine.

<sup>3</sup> <http://wikimapia.org/>

This process results in an RDF database of locations and points of interest near the user with additional information, such as names, descriptions and type. Typical results for a user located at the Spui square in Amsterdam from Linked Geo Data include *Spui25*, *Het Lieverdje*, *Nieuwezijds Voorburgwal*, from DBpedia: *Spui (Amsterdam)*, *Universiteit van Amsterdam*, from GeoNames: *Lutherse Kerk*, *Begijnhof*. In addition quite a few historical buildings are found. For the Spui we find 304 URIs related to that square while searching within a 150 meter radius. These 304 URIs are associated with 2467 RDF triples and 678 geographical shape definitions. Due to the crawling process the system will also find places that are further removed than the search radius. Only 103 URIs (with 973 RDF triples and 264 shape descriptions) represent locations that have an actual distance from the user which is less than 150 meters.

#### 4.2 Merge and Align Locations

The URIs that were gathered in the harvesting process by no means correspond to unique locations. Many points of interest have several locations associated with them. In the “Merge and Align” process we try to combine the different results into an “aligned local map”, see Fig. 4. This process involves both spatial reasoning and alignment techniques.



Fig. 4. Method: Merge and Align Locations

We encountered typical Semantic Web challenges, such as different schemas, different labeling conventions, different geodata (e.g. square *Spui* in Amsterdam has at least 5 different coordinates in LOD), errors in geodata and in human annotation and conflicts in typing (e.g. *Begijnhof* `rdf:type way`, *Begijnhof* `rdf:type area`, and *Begijnhof* `rdf:type building`).

We developed a number of mapping rules to align the different vocabularies and schema’s. First, vocabularies such as Wikimapia tags and Wikipedia categories were mapped onto the LGDV ontology, which was slightly extended with a number of relevant concepts. Second, the LGDV top level concepts were mapped onto the facet ontology. In total some 200 mapping rules were created by hand.

The first step in the merging process is to find candidate URIs that could possibly refer to the same physical location. From the list of candidates we select a root URI, preferably one that has a spatial description in the form of a polygon. Using the `space_nearest` predicate in the spatial reasoning package ([3]), we retrieve those URIs that are within a small distance from the URI we are investigating. We have found that the inaccuracy of the geodata requires a range of at least 35 meters in order to find all possible candidates. Subsequently the

candidate locations will be matched with the root location in terms of type and name. The type matching requires some ontological mappings since the URIs come from different sources which have different schema's. The name matching requires a normalization of labels, since many sources have conventions to qualify labels with tags like language, city or even more specific qualifications (e.g. "Maagdenhuis", "Maagdenhuis (nl)", "Universiteit van Amsterdam: Maagdenhuis"). Normalizing labels is not a guarantee that different names of the same object will be mapped onto the same location. In our example dataset the URI "[http://dbpedia.org/resource/University\\_of\\_Amsterdam](http://dbpedia.org/resource/University_of_Amsterdam)" falls within the location of the Maagdenhuis (the administrative centre of the University of Amsterdam), but the URI describes the University of Amsterdam in general and name matching fails. In such cases a "skos:relatedTo" relation will be added.


A second step concerns the alignment of resources. When a number of URIs have been identified as pointing to equivalent locations the information of each URI will have to be integrated. A new (unique) URI will be generated with a type that conforms to the LGDV ontology with our own extensions. The new URI will contain provenance information about its sources, the normalized label will be used as new skos:prefLabel, original labels will be used as skos:altLabel and scope notes will be copied from all sources. Integrating the spatial information is a bit more difficult. A set of URIs may have associated points, lines or polygons. Our current alignment algorithm takes the largest polygon that encompasses the most points in the locations and discards points that are outside this preferred polygon. In addition the centroid of the polygon is added as the point coordinates of the location. More sophisticated spatial reasoning could be employed here, for example we could use the fact that crowd-sourced coordinate data may be subject to a discrepancy between a camera location and the actual location of the object being photographed. In addition we could use type and location information to constrain certain location interpretations, e.g. it is unlikely that a pub is located within the administrative centre of a university. The current system does not implement these constraints. The result of this subtask is a set of URIs that represent unique physical locations with their integrated and aligned properties.

### 4.3 Semantic Enrichment

In this subtask, we start a "semantic crawling" process by using the labels found in the previous subtask as key for several search engine queries.

Figure 5 shows the reasoning steps that will enrich the data acquired in the previous processes. We use two sources for semantic enrichment: the Dutch Wikipedia server and the ECulture data cloud server. Both servers are queried with keywords derived from the label fields of the locations combined with background knowledge. For example, the label "Het Lieverdje" is converted into a query (**spui+lieverdje**) to the Cliopatria search engine to find artworks relevant to the location.

The results can yield new keywords (such as the name of a person) for further crawling. Where DBpedia does not give any results, Wikipedia pages are



**Fig. 5. Method: Semantic Enrichment**

retrieved and basic information is extracted from the HTML source, such as geo-coordinates, category information and (href) links to other topics. Since this crawling process can –in principle– continue indefinitely, we put a pragmatic limit to the length of the link paths followed. This limit depends on the facet selection that the user has made, and is usually set to 3.

Data that have been retrieved in the semantic enrichment process are in general not annotated with a type that can be related to the facets. Dutch Wikipedia pages have a category that essentially is a string. We use mapping rules to classify the Wikipedia categories to WordNet classes. For example, the Dutch string “Nederlands architect” (Dutch architect) is mapped to the concept architect in WordNet. The data from the cultural institutions generally use literal terms to describe subjects of art works. Using simple lexical matching and some mapping rules we map the subject terms to WordNet concepts. For example the Dutch word “bezetting” (occupation –of a building–) will be mapped to the WordNet concept `occupation-3`.

#### 4.4 Classification of URIs

The URIs collected in the previous steps come from many different schema’s and use different ontologies. For example, for the Spui square the enriched location set of URIs contains 43 different values for the `rdf:type` property (a.o. restaurant, shop, building, church, place\_of\_worship, university, way, bequinage, market, marketplace). Each of these types has to be classified in terms of the facets and subfacets. Most of these types occur in the (extended) linked geo data ontology (LGDV). The hierarchy of the facets and LGDV are mapped onto each other such that each type maps to a facet-subfacet pair. In addition to location types, the RDF database contains URIs pointing to persons, organisations, artworks, events etc. We use the WordNet hierarchies to construct a mapping between these types and the facet hierarchy.

#### 4.5 Interaction with the Mobile Device

The moment the user opens the application we already know the geolocation. The mobile device can then send a request to the server to create an RDF



database, which is subsequently send back. After the mobile device has recieved the RDF triples that were collected at the server, the results have to be presented to the user. From there the user can use three levels of refinement: (1) pointing to a specific heading, where  $h = [0..359]$ , (2) select facets of resources relevant to the current geolocation and heading, where cultural related facets are location, event, artwork and people, and (3) select subfacets of the selected facet, e.g. painting, photograph, book, artist, musician, politician, sport or conflict.

## 5 Architecture: a Light Weight Client with a Heavy Endpoint

In order to find an intuitive way to present the enriched data, we apply a number of constraints. We already know a lot about the users: they are mobile, they want to be able to see useful content immediately without too much configuration and they need to be able to accomplish things with just a few taps [9]. Furthermore, a mobile device, such as an iPhone is limited by bandwidth, computing and power capacity. Therefore we need to develop a light weight client for user interaction. The GUI of this device is limited:  $7 \pm 2$  items is about what a smart phone can display and be controlled by Fingertip or stylus-based touching. The 7 Fingertip-Size Targets is similar to the Magical Seven defined by the psychologist George Miller. He stated that human short-term memory has a short-term memory span of approximately seven items plus or minus two [2]. End-user interaction is handled by a mobile device, in our case an iPhone 3GS (with GPS capabilities, a digital compass and assuming an internet subscription).

Most iPhone applications uses the **UIView Controller**: one of the basic packages to display content and handle user interaction. To make an intuitive location selection, we use augmented reality<sup>4</sup>, for which we adopted the open source **ARKit** package, which is able to display real world vision via the phone's camera and put labels and controls over this [11]. The **CLLocation** package tells the application the geolocation expressed in WGS 84 and heading in degrees [12]. Finally, facet selection show the user two layers of selections: the main facets and subfacets. When choosing a main facet, the sub facets will adapt accordingly, see Figs 7,8 a and b. The iPhone communicates via a REST interface with the back-end server.

We also know a lot about the cultural heritage domain. There are several sources, such as the Dutch Art History resource and Rijksmuseum Amsterdam resource, centrally accessible via **ClioPatria** [7]. Diverse LOD resources are accessible via SPARQL or via our web services which we access with semantic crawling method described in section 4. We used several existing Prolog packages, able to access LOD resources and perform graph search [8,3]. This resulted in a three tier architecture: user interaction, reasoning and LOD resources, see Fig. 6.

<sup>4</sup> A live direct or indirect view of a physical real-world environment whose elements are merged with (or augmented by) virtual computer-generated imagery - creating a mixed reality, see [http://en.wikipedia.org/wiki/Augmented\\_reality](http://en.wikipedia.org/wiki/Augmented_reality)

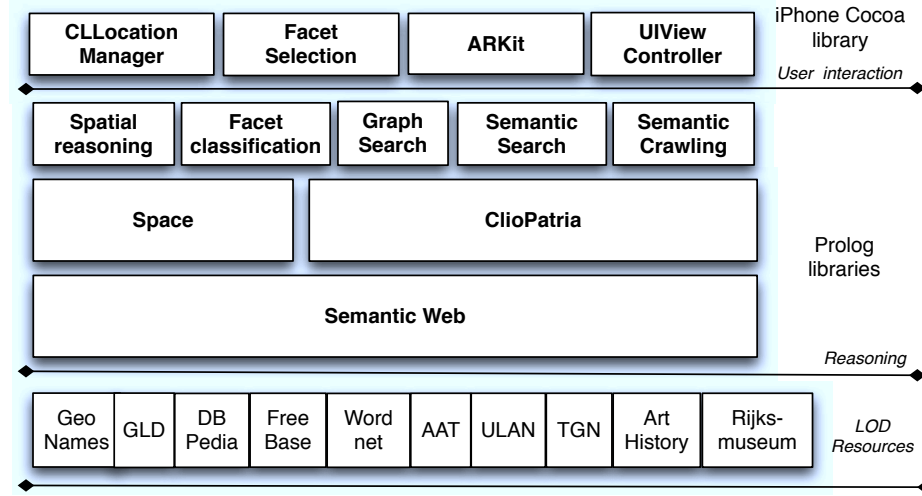


Fig. 6. 3-Tier layer architecture: user interaction on the iPhone, reasoning with Prolog on the Back-end server and LOD resources in the Semantic Web.

## 6 Use Cases: Displaying POIs

In this section we present three use cases, where the user is visiting the famous “Spui” square in Amsterdam. The user will walk over the square and point with his/her iPhone to three touristic hotspots: the “Lutherse Kerk” (a church), “het Maagdenhuis” and the “Helios Building”. For some cases we show the Dutch language information, because the metadata is only available in the Dutch Language. The metadata can be found here: “<http://eculture2.cs.vu.nl/spuittest>”.

### 6.1 Scenario: the “Spui” Square and the “Lutherse Kerk”

A tourist is standing on the Spui square in Amsterdam and opens our iPhone app. The application sends the geolocation  $s = \langle 52.2237, 4.5333 \rangle$  and heading  $h = 182.23$  to the server which starts to retrieve information. The iPhone app receives an RDF dataset from the server relevant to locations and objects within a 150m range of the user. Using the place facet a Google Maps like representation of the area and points of interest could be displayed.

The next step is to use the heading of the user to determine what object the user’s iPhone is directed at. This turns out to be the “Oude Lutherse kerk”, a church. Assuming that the user has selected the artwork/painting facet, the system will launch a search request (**spui+lutherse+kerk**) to the ClioPatria engine, which returns a set of pointers to paintings relevant to the place. One of the paintings is selected and additional information about the painting is retrieved. The results are projected on the screen of the iPhone, see Fig. 7.



Fig. 7. On the left: Augmented reality view on “Lutherse Kerk” (church) with selection: [artwork/painting], combined with annotation (in Dutch) and the facet-based selection [artwork/painting]. On the right explanation of the components of the GUI.

## 6.2 Scenario: the “Maagdenhuis” Building

The “Maagdenhuis” was built in 1783 and served as an orphanage for girls until 1953. Since then it is the administrative centre of the University of Amsterdam. In 1969 the “Maagdenhuis” became famous and an icon for student protest: it was occupied 5 days by students demanding influence in university affairs. Since then it has been occupied around ten times. Searching the ECulture engine with the key “maagdenhuis+amsterdam” results in about 100 hits of objects (paintings, ceramics, other types of objects) about this place. When we filter these results on the “event” facet, 5 photographs remain that are part of the collection of the Amsterdam Historical Museum and depict the student occupation of 1969 (Fig. 8 a).

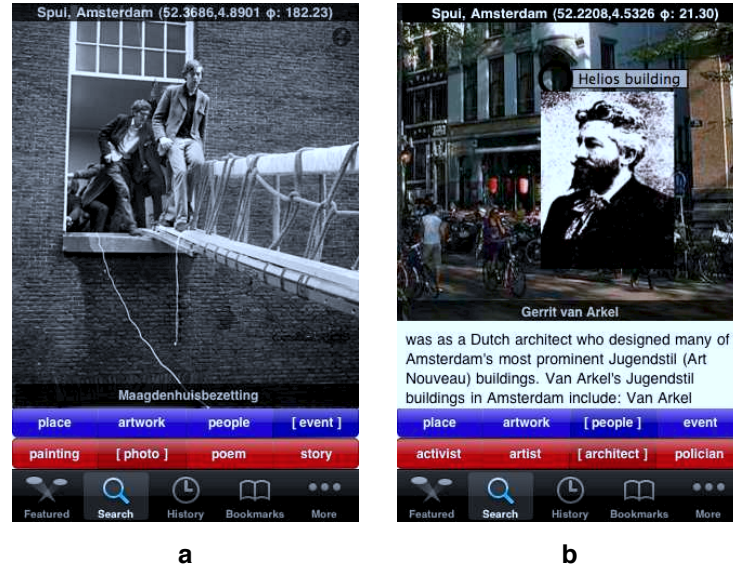


Fig. 8. (a): Photograph of the ending of the Maagdenhuis occupation. (b): Augmented reality view on the “Helios building”, showing the architect with selection [people/artist].

### 6.3 Scenario: the “Helios Building”

When the user chooses the selection [people/artist] the system will attempt to find relevant persons, for example architects. In this case, this results in a description of “Gerrit van Arkel”, the architect of the famous “Helios Building” at the Spui square (Fig. 8 b). The Helios building and its architect could also have been found on the basis of user coordinates and bearing. Data about this building are also found using the location data and the semantic enrichment process, resulting in the retrieval of the Wikipedia page of “Gerrit van Arkel”.

## 7 Discussion

There are countless ways to encode location on the web. There is GML, KML, GeoRSS, the vCard and hCard microformats, etc. We have found that GeoRSS is the most promising of these. Both the Open Geospatial Consortium and the World Wide Web Consortium support GeoRSS and it allows a gradual dumbing down from (partial) GML shape support to simple points (see the Geospatial Vocabulary<sup>5</sup>). The periodically updated World Geodetic System is the only viable coordinate system that works in a uniform way throughout the world. The

<sup>5</sup> <http://www.w3.org/2005/Incubator/geo/>

accuracy might not be sufficient for many indoor augmented reality application, but for outdoor guides like the one presented in this paper it is more than sufficient. We have found that if you want to reason about geospatial concepts, it is important to represent shapes as first-class citizens. This makes conversion between various geospatial formats on the web much easier, as well as allowing you to add support for new types of shapes (e.g. polygons, polygons with holes, geometry collections) in the future if they eventually turn out to be relevant for your project. Also, it is important to draw the boundary between the representation of geospatial and semantic objects at the URI of the geofeature, i.e., e.g. not to represent shapes using RDF triples or `rdfs:subClassOf` relations in GML. This way you can benefit the most from the current standards provided by the OGC and W3C.

While using various sources with geo data information we found that significant discrepancies exist between coordinates for the same location. In many cases these discrepancies exceed a distance of 20m, in some cases even hundreds of meters. Our system could be used to identify such discrepancies and point crowd-sourcing users to possible corrections to be made in the open source data.

Schema	URIs	RDF statements
dbpedia.org	3	153
linkedgeodata.org	21	162
nl.wikipedia.org	6	53
rdf.freebase.com	1	15
rijksmonumenten.wikia.com	227	1619
sws.geonames.org	10	117
wikimapia.org	42	401

**Table 1.** Numbers of URIs and RDF statements for different schema's of the location data for the Spui

Tabel 1 shows the statistics for the location data of the Spui location. The major part of the data comes from non-LOD sources. In addition a significant amount of other data comes from the Eculture sources. Therefore we conclude that the Linked Open Data sources in isolation are currently too limited to provide interesting semantic information but combined with each other and with a number of other sources a really informative location-based service can be created.

Matching the many “synonymous” geofeatures and their types on the web is a challenge for the near future. In both the semantic web and geospatial community this is current research, respectively named ontology alignment or conflation. Another challenge for the future is to provide guided tours through the city-based on the semantics of the surroundings. For example, if you are struck by a building with an interesting style of architecture, it would be great if

your mobile device could route you through town along related buildings, telling the story behind their commonality along the way.

## 8 Conclusions

In this paper we explored the use of location aware mobile devices for acquiring knowledge from, searching and browsing large collections of general and cultural heritage information repositories using minimal interaction. We showed that given a particular geolocation, current Semantic Web data and technology and the constraints of a mobile device, we can find interesting material for an active tourist, providing dynamic information in favor of a classical travel guide.

We presented a novel user interface design, where a combination of location, heading and facet-based filtering provides a user with a dedicated smart phone application. The challenges that we solved in its development are: determining the ontological characterization of the current location, by mapping a geolocation represented by a point to the ontological characterization of that location. We constructed a ‘mental map’ of nearby points of interest with their direction, by taking a range and finding objects of interest within a circular shape. Next we crawled for other information relevant to these locations, using semantic crawling. It turns out that the interplay of sources from the LOD cloud, Wikipedia and cultural heritage data can provide a very rich knowledge base about a certain topic that is machine processable. Semantic crawling resembles the process of a human using Google to find information, using a cycle of key word selection, inspection of results, interpreting and (possibly generating new queries on the basis of this interpretation. Finally, we use augmented reality in combination with facet selection to present this information to a user on a mobile device.

Next to mobile devices, there are also a number of other common devices that become connected to the Web, such as televisions, cars, and other devices in houses (or domotics). All these devices have a form of limitation, such as a remote control for a television or a dashboard, but also an advantage for determining a user’s context, for example watching a certain movie or driving in a certain direction. Semantic crawling can be applied to find background information about movies and actors or locations on the road.

We found that the Linked Open Data sources in isolation are currently too limited to provide much interesting semantic information, but combined with each other and with a number of other sources (for example sources from the cultural heritage domain) a really informative location-based service can be created. Semantic crawling is a major improvement over the current state of the art applications such as Google Maps), which only show labels of resources near a given location, instead of the background knowledge associated with the location. We feel that the power of the Semantic Web concept has clearly been demonstrated in the application we have described. In isolation the currently available repositories provide limited knowledge, but combining a large number of sources and using a semantic crawling approach that accesses many of the Semantic Web services that have become available, yields a reality that is ap-

proaching the original Semantic Web vision. In some ways, knowledge acquisition has moved from acquiring knowledge from human experts to the enterprise of acquiring and integrating knowledge from the rich sources of knowledge on the World Wide Web.

## Acknowledgments

This work has been partially supported by the EU project NoTube (ICT-231761), the Poseidon project (supported by the Dutch Ministry of Economic Affairs under the BSIK03021 program) and the Agora project (funded by NWO in the CATCH programme, grant 640.004.801).

## References

1. Bloem, J., van Aart, C. (2010). On Digital (Mobile) Humanism: helping understand it, engineer its future and ensure its social benefit. WebScience 2010.
2. Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review* 63 (2): 81-97.
3. van Hage, W. R., Wielemaker, J., Schreiber, A. Th. (2010). The Space package: Tight Integration Between Space and Semantics. *Transactions in Geographical Information Systems* 14(20): 131-146.
4. Becker, C., Bizer C. (2008). DBpedia Mobile: A Location-Enabled Linked Data Browser. *Linked Data on the Web (LDOW2008)*.
5. Evjemo, B., Akselsen, S., Schurmann, A. (2008). User Acceptance of Digital Tourist Guides Lessons Learnt from Two Field Studies. In *Human-Computer Interaction, Part I, HCI 2007, LNCS 4550*, pp. 746755.
6. Bizer, C., Heath, T., Berners-Lee, T. (2009) *Linked Data - The Story So Far*. *International Journal on Semantic Web and Information Systems*, pp. 1-22.
7. Schreiber, G., Amin, A., Aroyo, L. Van Assem, M., De Boer, V. Hardman, L., Hildebrand, M., Omelayenko, B., Van Osenbruggen, J., Tordai, A., Wielemaker, J., Wielinga, B. (2008). Semantic annotation and search of cultural-heritage collections: The MultimediaN E-Culture demonstrator. *Journal of Web Semantics*, vol 6., pp. 243-249.
8. Van Osenbruggen, J., Amin, A., Hardman, L., Hildebrand, M., van Assem, M., Omelayenko, B., Schreiber, G., Tordai, A., De Boer, V., Wielinga, B. Wielemaker, J. De Niet, M., Taekema, J., Van Orsouw, M., Teasing A (2007). Searching and Annotating Virtual Heritage Collections with Semantic-Web Techniques. In: *Museums and the Web 2007*, April 11-14.
9. iPhone Human Interface Guidelines, <http://developer.apple.com/iphone/library/documentation/UserExperience/Conceptual/MobileHIG/Introduction/Introduction.html>
10. Google goggles. <http://www.google.com/mobile/goggles/#landmark>
11. Open source augmented reality package: ARKit: <http://code.google.com/p/iphonearkit>
12. CLLocation Class Reference:incorporates the geographical coordinates and heading of a devices location: [http://developer.apple.com/iphone/library/documentation/CoreLocation/Reference/CLLocation\\_Class/CLLocation/CLLocation.html](http://developer.apple.com/iphone/library/documentation/CoreLocation/Reference/CLLocation_Class/CLLocation/CLLocation.html)
13. CEN: European Committee for Standardization, <http://www.cen.eu>.