

DBpedia Mobile: A Location-Aware Semantic Web Client

Christian Becker and Christian Bizer

Freie Universität Berlin, Germany
chris@beckr.org, chris@bizer.de

Abstract. DBpedia Mobile is a location-aware client for the Semantic Web that can be used on an iPhone and other mobile devices. Based on the current GPS position of a mobile device, DBpedia Mobile renders a map indicating nearby locations from the DBpedia dataset. Starting from this map, the user can explore background information about his surroundings by navigating along data links into other Web data sources. DBpedia Mobile has been designed for the use case of a tourist exploring a city. As the application is not restricted to a fixed set of data sources but can retrieve and display data from arbitrary Web data sources, DBpedia Mobile can also be employed within other use cases, including ones unforeseen by its developers. Besides accessing Web data, DBpedia Mobile also enables users to publish their current location, pictures and reviews to the Semantic Web so that they can be used by other Semantic Web applications. Instead of simply being tagged with geographical coordinates, published content is interlinked with a nearby DBpedia resource and thus contributes to the overall richness of the Geospatial Semantic Web.

Key words: Semantic Web, Linked Data, Geospatial Web, DBpedia, Location-Aware Applications

1 Introduction

Mobile phones feature increasingly powerful hardware, software and data connectivity, and more and more phones such as the iPhone 3G are shipped with built-in GPS receivers, whose positioning capabilities are exposed to third party applications. In parallel, the Semantic Web is growing rapidly, and contains a large amount of location-related data.

A Semantic Web data source which provides information that could be useful for a tourist exploring a city is DBpedia [1]. The DBpedia dataset has been extracted from Wikipedia. For currently more than 2.49 million “things”, it features labels and short abstracts in 14 different languages, 588,000 links to images and 3,150,000 links to external web pages. The DBpedia dataset contains information about more than 300,000 locations. DBpedia data about locations is interlinked with various other location-related datasets. Examples include the GeoNames dataset, which classifies locations in a feature hierarchy; the flickrTM

wrappr, which provides images depicting locations; the EuroStat, Riese, US Census and CIA Factbook datasets, which provide statistical information about locations; Revyu, which might provide reviews about a location; the YAGO dataset, which provides type information for locations such as museums, train stations or monuments; and FOAF profiles by individuals that state that they are based near, work at or were born at DBpedia locations. Altogether there are around 2,180,000 data links pointing from DBpedia into other RDF data sources on the Semantic Web. Figure 1 gives an overview of all Web data sources that are currently interlinked with DBpedia. The data sources that can be reached directly from DBpedia by following outgoing data links are highlighted in the figure.

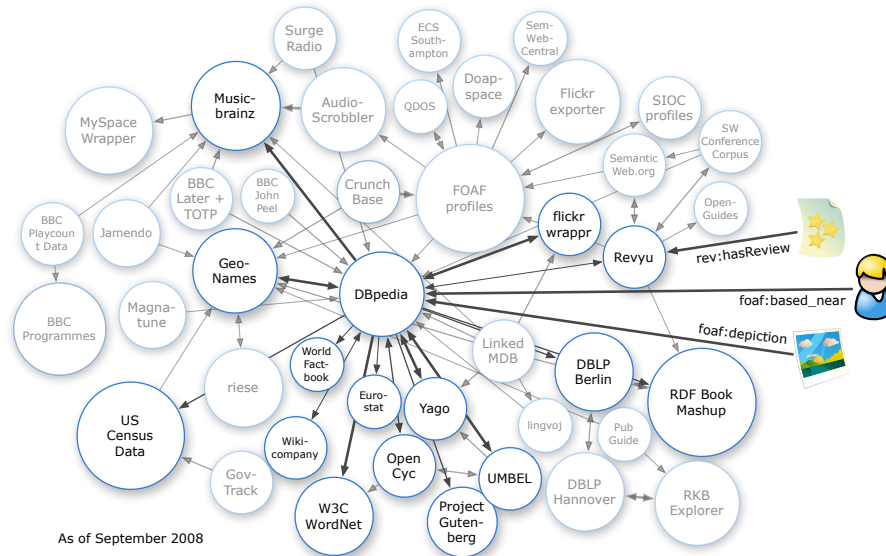


Fig. 1. Data links pointing from DBpedia into other data sources on the Semantic Web and links provided by DBpedia Mobile.

2 DBpedia Mobile

*DBpedia Mobile*¹ allows users to discover, search and publish Linked Data pertaining to their current physical environment using their mobile phones as well as standard web browsers. The application consists of a map view and a Fresnel [2]-based Linked Data browser. Starting from this map, users can explore background information about their surroundings by navigating along data links into

¹ <http://beckr.org/DBpediaMobile>

other Web data sources. Besides accessing Web data, DBpedia Mobile also enables users to publish their current location, pictures and reviews to the Semantic Web. Instead of simply being tagged with geographical coordinates, published data is interlinked with a nearby DBpedia location. In Figure 1, user-generated content that is published using DBpedia Mobile is depicted by the three outer right items.

2.1 Browsing the Geospatial Semantic Web

DBpedia Mobile's initial view is a map display that indicates the user's position and nearby DBpedia resources, using appropriate labels and icons. Figure 2 shows DBpedia Mobile's initial view on an iPhone. The map can be moved by dragging it on the touch screen. It can be zoomed using a pinch gesture or the provided controls.



Fig. 2. DBpedia Mobile running on an iPhone 3G and showing a map view of resources in the user's proximity.

Locations are depicted with adequate icons based on a mapping of selected YAGO categories [3]. DBpedia Mobile allows users to switch between 14 languages to be used for labels and summary texts (English, German, French, Dutch, Polish, Italian, Spanish, Japanese, Portuguese, Swedish, and Chinese).

Clicking on a resource brings up a *summary view* of the selected item. A summary view includes a short text describing the resource and optionally an image, which is often provided by DBpedia or the flickrTM wrapper, a link to the resource's `foaf:homepage` as well as reviews of the resource from Revyu. It also shows persons that have indicated that they are `foaf:based_near` the resource in their FOAF profile or using DBpedia Mobile. At the bottom of the page, links are provided to access a *photo view* and a *full view* for the resource.

Figure 3 shows a *summary view* for the *Brandenburg Gate* which includes a review obtained from Revyu as well as a photo obtained from the flickrTM

wrapp, and lists nearby DBpedia Mobile users. Specific site icons (such as a blue star for Revyu) are used to indicate the provenance of displayed data.

The *photo view* displays further depictions of the resource, obtained from DBpedia and the flickrTM wrapp.

The *full view* displays all known properties of the resource. Figure 4 shows a full view of the Brandenburg Gate's district *Tiergarten*, which incorporates Linked Data from GeoNames, the flickrTM wrapp, and DBpedia. If the displayed data contains RDF links into other datasets, the user may click them to obtain a full view of the referenced resource. In this manner, he can navigate from the DBpedia dataset into other interlinked data sources. DBpedia Mobile is not limited to a fixed set of data sources but may be used to access all data sources that are or will in the future be interlinked with DBpedia or with other data sources that are reachable from DBpedia. This allows interesting navigation paths: For instance, a user may navigate into GeoNames and traverse its `parentFeature` hierarchy to find out more about the city, state and country in which a resource is located. From a location, he may navigate to a person within the DBpedia dataset that was born, died or worked at the location. If the person is an author, he may then follow data links into the RDF Book Mashup or the Project Gutenberg data sources and explore information about the author's books². If the tourist is interested in local bands, he may navigate from DBpedia into Musicbrainz and find out more about albums of the bands.

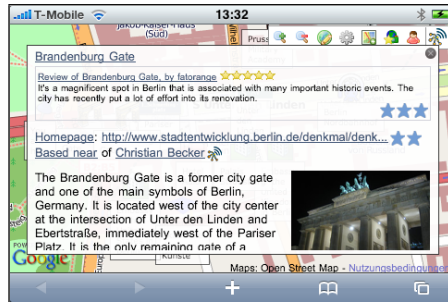


Fig. 3. A *summary view* of the *Brandenburg Gate* that includes a review obtained from Revyu as well as an abstract text and a photo obtained from DBpedia, and lists a nearby DBpedia Mobile user.

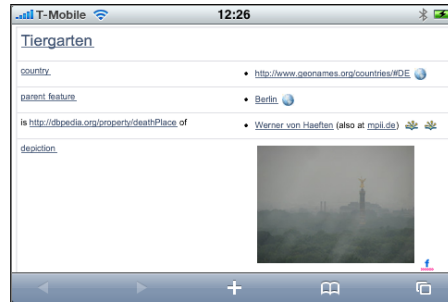


Fig. 4. A *full view* of the *Brandenburg Gate's district Tiergarten*, which incorporates Linked Data from GeoNames, the flickrTM wrapp, and DBpedia. From here, the user may navigate into other interlinked datasets.

² For example, this works for the navigation paths *Bedford* → *John Bunyan* and then to his publications on Project Gutenberg, or *University of Southampton* → *Tim Berners-Lee* and then to his publications on DBLP.

DBpedia Mobile generates the different views using Fresnel [2] lenses and formats on the server side. Prior to rendering a view for a resource, DBpedia Mobile performs data augmentation, whereby it retrieves interlinked data from the Web and caches retrieved data in a server-side RDF store. This involves dereferencing the resource URI and querying the Sindice Semantic Web search engine³ for related information, as well as Revyu for reviews. In a similar manner as the Semantic Web Client Library⁴, specific predicates found in retrieved data such as `owl:sameAs` and `rdfs:seeAlso` are then followed for up to two levels in order to gain more information about the resource, and to obtain human-friendly resource labels. Because of this approach, there is no inherent restriction on which data sources are discovered. In fact, anyone may make statements about a DBpedia resource, and once this data link has been picked up by Sindice, it will be found by DBpedia Mobile and integrated into its output.

Different data sources may contain statements about the same resource, but may each use their own URIs to identify the resource. In the context of Linked Data, the predicate `owl:sameAs` is commonly used to connect such *URI Aliases* between distinct data sources [4]. In order to also incorporate statements that pertain to URI Aliases, an `owl:sameAs` inferencer is employed before a view is rendered.

2.2 Filtering Information

The user may filter the map for resources that match specific constraints. DBpedia Mobile's RDF store initially includes only those DBpedia datasets that are required to generate the map display, i.e. geo-coordinates, labels and YAGO classes. To build filters that involve data from other sources, a user may request the currently visible area to be crawled. The crawling process determines all DBpedia resources in the area and then queries DBpedia and Sindice for the URIs of related resources. These URIs are then dereferenced and the retrieved data is loaded into DBpedia Mobile's RDF store, where it can be used for filtering and display.

Filters are constructed using a Filter Builder that supports the creation of *Simple* and *SPARQL Filters*: *Simple Filters* consist of one or multiple conditions that are applied in conjunction. Conditions can pertain to *resource types*, *ratings* that are associated with resources, and for demonstration purposes, the *population* of inhabited areas and the *number of stories* of buildings. Resource types may be chosen from pre-defined sets that encompass one or more YAGO categories, such as *Museums* or *Train Stations*, whereas ratings and numeric conditions are specified using an operator drop-down and an associated operand input field. In Figure 5, the Simple Filter Builder interface is used to limit the display to resources of type *Sightseeing* which have received a rating of at least three stars.

³ <http://sindice.com/>

⁴ <http://sites.wiwiss.fu-berlin.de/suhl/bizer/ng4j/semwebclient/>

SPARQL Filters are SPARQL `CONSTRUCT` queries which are evaluated against the RDF store. The resulting RDF graph is then used to generate the map display. For example, the filter “Stations of Berlin’s U7 train line” restricts the display to resources that have the DBpedia infobox properties `system` and `line` equal to “BVG”, Berlin’s transportation company, and “U7”, respectively. Figure 6 shows the resulting map display.

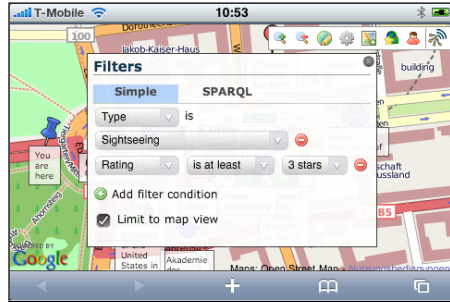


Fig. 5. The Simple Filter Builder Interface.

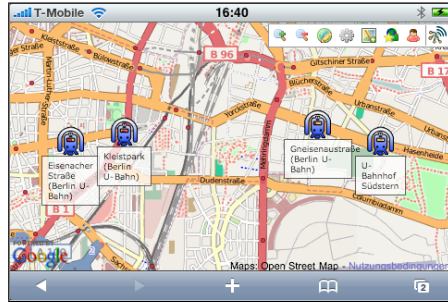


Fig. 6. Specific train lines plotted using the filter “Stations of Berlin’s U7 train line”.

2.3 Publishing Linked Data to the Web

DBpedia Mobile facilitates the publication of content with attached location information to the Semantic Web, directly from the mobile device. A user may publish his current location as well as photos and reviews, interlinked with DBpedia resources.

Tim Berners-Lee et al. discuss the publication of RDF statements by users in [5]. Specifically, they introduce graphs that anyone may edit. Contrarily, DBpedia Mobile assigns each user an individual RDF graph that only he can update and which is Web-accessible via the URI <http://beckr.org/DBpediaMobile/graphs/username>.

DBpedia Mobile wraps the triple generation process in user interface dialogs, which dramatically reduces the technical expertise required to publish content. To publish content, the user logs in and opens the Content Creation Panel (depicted in Figure 7). He then selects the type of information to publish (i.e. position, photo or review) and chooses a DBpedia location to be used to tag the published content. Based on his location, the nearest DBpedia resource is automatically suggested, as shown in Figure 8.

A location update is automatically reflected in his user graph as a triple stating that he is `foaf:based_near` the selected resource. In the case of a photo, he is prompted to upload the image, which is then stored on the server and is published in his User Graph as a `foaf:depiction` of the selected resource. If

the user chose to publish a review, he is presented with a review form which will be directly submitted to Revyu.

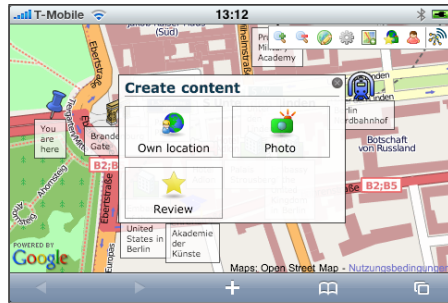


Fig. 7. The Content Creation Panel.



Fig. 8. The nearest DBpedia resource is suggested based on the user's position.

A user may wish to interlink the location updates and photos he publishes using DBpedia Mobile with his FOAF profile. To do so, he states that the resource URI he uses to refer to himself in his FOAF profile is identical (`owl:sameAs`) to his personal resource URI with DBpedia Mobile. This enables Semantic Web crawlers and Linked Data browsers to find and merge information from both sources.

3 Technical Background

DBpedia Mobile is realized as a client-server application. Queries, data retrieval and storage, as well as formatting activities are performed by the server-side Marbles engine⁵. This provides the application with high bandwidth, processing and storage resources and allows search requests as well as the Fresnel-based view generation to touch on large amounts of data. The client application is written in JavaScript and can be accessed with web browsers that feature adequate Document Object Model (DOM) support to host the underlying Google Maps API, such as Safari on the iPhone or Opera Mobile 8 on Windows Mobile, and all major desktop browsers.

DBpedia Mobile may be initialized with the user's GPS position using a supplemental launcher application, or via the Yahoo! FireEagle⁶ web service. Alternatively, the user's position may be estimated based on his IP address location.

The server application is implemented as a Java Servlet which uses the Sesame RDF framework. Triple storage is realized in conjunction with a MySQL

⁵ <http://beckr.org/marbles>

⁶ <http://fireeagle.yahoo.net/>

database, whereby Sesame’s RDBMS storage capabilities were enhanced to support inference, as well as storage and filtering of geospatial points using MySQL Spatial Extensions. The SIMILE Fresnel Engine and the Saxon XSLT processor are used for the generation of resource views.

4 Conclusions

DBpedia Mobile allows users to discover, search and publish Linked Data on the Semantic Web using their mobile devices. The application is not restricted to a fixed set of data sources but may retrieve and display data from arbitrary Web data sources. This enables DBpedia Mobile to be employed within unforeseen use cases. Realized and proposed use cases include tourism⁷, nearby events⁸ and personal recommendations⁹. DBpedia Mobile publishes user-generated content with attached location information to the Semantic Web. This advances “geo-tagging” to “resource-tagging” and allows users to contribute to the overall richness of the Geospatial Semantic Web.

More information about DBpedia Mobile is available at <http://wiki.dbpedia.org/DBpediaMobile>.

References

1. S. Auer, C. Bizer, J. Lehmann, G. Kobilarov, R. Cyganiak, and Z. Ives. DBpedia: A nucleus for a web of open data. In *Proceedings of the 6th International Semantic Web Conference and 2nd Asian Semantic Web Conference (ISWC/ASWC2007)*, volume 4825 of *LNCS*, pages 715–728, Springer, 2007.
2. E. Pietriga, C. Bizer, D. Karger, and R. Lee. Fresnel: A browser-independent presentation vocabulary for RDF. In *International Semantic Web Conference (ISWC2006)*, volume 4273 of *LNCS*, pages 158–171, Springer, 2006.
3. F. M. Suchanek, G. Kasneci, and G. Weikum. Yago: A core of semantic knowledge. In *WWW ’07: Proceedings of the 16th international conference on World Wide Web*, pages 697–706, ACM, 2007.
4. Bizer, C., Cyganiak, R., Heath, T.: How to Publish Linked Data on the Web. URL: <http://www4.wiwi.fu-berlin.de/bizer/pub/LinkedDataTutorial/02/18/2008>.
5. Berners-Lee, T., Hollenbach, J., Lu, K., Presbrey, J., Prud’ommeaux, E., Schraefel, M.: Tabulator Redux: Browsing and Writing Linked Data. In: *1st Workshop about Linked Data on the Web (LDOW 2008)*, April 2008. URL: <http://events.linkedata.org/ldow2008/papers/11-berners-lee-hollenbach-tabulator-redux.pdf>, 03/02/2008.

⁷ <http://tomheath.com/slides/2008-09-graz-triplei.pdf>

⁸ <http://blog.dbtune.org/post/2008/07/30/Lastfm-events-and-DBpedia-mobile>

⁹ <http://apassant.net/blog/2008/04/24/social-mobile-semantic/>