

INVESTIGATING THE USE OF SEMANTIC TECHNOLOGIES IN SPATIAL MAPPING APPLICATIONS

Taha Osman, Luke Shires
School of Science and Technology,
Nottingham Trent University,
Nottingham NG11 8NS
{taha.osman, luke.shires}@ntu.ac.uk

Tope Omitola, Nigel Shadbolt
School of Electronics and Computer Science
University of Southampton
Southampton SO17 1BJ
Email: {t.omitola, nrs}@ecs.soton.ac.uk

Jeremy Hague
Business Innovation Team,
Nottingham Trent University,
Nottingham NG1 4BU
jeremy.hague@ntu.ac.uk

KEYWORDS

Semantic Web, Spatial Mapping, Context Awareness, Linked Data, Information Retrieval, Semantic Knowledgebase.

ABSTRACT

Semantic Web Technologies are ideally suited to build context-aware information retrieval applications. However, the geospatial aspect of context awareness presents unique challenges such as the semantic modelling of geographical references for efficient handling of spatial queries, the reconciliation of the heterogeneity at the semantic and geo-representation levels, maintaining the quality of service and scalability of communicating, and the efficient rendering of the spatial queries' results. In this paper, we describe the modelling decisions taken to solve these challenges by analysing our implementation of an intelligent planning and recommendation tool that provides location-aware advice for a specific application domain. This paper contributes to the methodology of integrating heterogeneous geo-referenced data into semantic knowledgebases, and also proposes mechanisms for efficient spatial interrogation of the semantic knowledgebase and optimising the rendering of the dynamically retrieved context-relevant information on a web frontend.

INTRODUCTION AND MOTIVATION

Ubiquitous and affordable access to interactive mapping solutions such as Google Maps and OpenStreetMap has encouraged their use beyond simple navigation to develop powerful tools that are useful in a multiplicity of applications from reporting and monitoring local crime to tourism.

The research reported in this paper was inspired by work on a multidisciplinary project at Nottingham Trent University that aimed to develop an interactive mapping tool that can visually map a multidimensional (social, economic, recreational, etc.) sense of real places and

communities. The tool should, for instance, visually aid local authorities in planning for public services provisioning in specific geographic regions by correlating petty crime statistics to economic affluence and to the residents' sense of belonging to the place.

The specification of the proposed interactive mapping tool lends itself naturally to context-aware modelling, as it has to analyse the situational condition associated with a geographical space and map them in accordance to the user needs. We argue that Semantic Web technologies are best placed to implement the proposed tool because of the following:

- Semantic Web technologies allow the modelling of the contexts using formal, machine comprehensible representations that can be interpreted by software agents in accordance with the user query parameters.
- Formal semantic modelling also allows leveraging the increasingly rich amount of relevant datasets available in the Linked Open Data (LOD) cloud; of particular interest are the large volumes of semantically-annotated data released by the UK government (on health, crime, housing, etc.) as part of the open government project (data.gov.uk).
- Semantic models are inherently extensible. New domains of interest and respective datasets can seamlessly be added and integrated into existing models using object and data properties (Alemang and Hendler 2011). The semantic knowledgebase developed for this project was adapted in another project (Brown et al. 2010), quite seamlessly, to power location-aware assistive technology aiding people with disabilities.

Nevertheless, there are a number of challenges associated with the geospatial representation and of the semantic knowledge and the efficient utilisation of Semantic Web technology to deliver good quality of service to the interactive mapping application users. In particular, in this paper we present a host of design decisions and development solutions to aid in the uniform interrogation of heterogeneous geo-tagged

datasets, optimise the utilisation of semantic repository in context-aware processing of geospatial queries, and render the retrieved information on web-based interactive maps.

The remainder of the paper is organised as follows: section 2 surveys related work, section 3 introduces the specification of the semantic interactive mapping tool, section 4 discusses the construction of the semantic knowledgebase, section 5 discusses the challenges of managing the geospatial aspects in semantic technologies, and section 6 concludes the paper and presents our plans for further work.

RELATED WORK

While the use of Semantic Web technologies in driving mapping applications is becoming increasingly popular, a large body of the reported work, while valuable in its domain of interest, either focuses on the usability of the semantic of geo-spatial applications (Cloug et al. 2011)(Becker and Bizer 2009), or predominantly focuses on geo-spatial problem solving aspects (Miron et al. 2007)(Brodaric 2007). Research that is more related to our work investigates more explicitly the processes of exploiting the Semantic Web technology.

There is a large body of work that endeavours to contribute to the middleware supporting semantic geo-spatial indexing. For instance, the effort in (Brodt et al. 2010) reports on the implementation of a triple store that natively integrates spatial query processing (deep integration). In their implementation, spatial query predicates of the OpenGIS specification are supported via SPARQL filtering. Another important contribution at the same middleware level is the work in (Cantador et al. 2008), which reports on the experience of integrating Geographic Information Systems (GIS) and Semantic Web technologies. They integrated geospatial features as a plugin to the LarKC platform for building web-scale semantic computing. A mapping tool was used to convert the GIS data into SPARQL-accessible virtual RDF format, resulting in the loss of the geospatial features, which were partially restored by populating the street and intersection coordinates into the AllegroGraph triplestore. A pathfinder workflow is also developed that implements a reasoner level plugin to find the shortest path in a semantic graph.

In our investigation, we focus on the utilisation of the geo-spatially enabled semantic technologies rather than their implementation. We adopted OWLIM (Geo-Spatial 2013) to host our semantic knowledgebase as it offers non-restricted and well-supported version for academic use. OWLIM also scored very well in terms of response-time and recall in a benchmarking investigation we performed in partnership with the Press Association (Thakker et al. 2010).

The focus of the LinkedGeoData project (Auer et al. 2009) is on the transformation and publication of the OpenStreetMap (OSM) data as semantically tagged

spatial information. They adopt a mixed approach for storing part of the OSM data to reduce the amount of data and increase performance. The reported work details the huge effort put into dynamic interlinking with other spatial datasets by using the Triplify approach (Auer et al. 2009b) to publish linked data from relational databases, and learning algorithms for mapping semantic datasets at the Description Logic level.

The class of applications that our use-case represents focuses on consuming (localising) external datasets rather than integration with and republication as Linked Open Data, assuming black-box use of geo-enabled semantic technology.

REQUIREMENT SPECIFICATION OF THE SEMANTIC INTERACTIVE MAPPING TOOL

Figure 1 below shows a screenshot of the PPlace Market mapping tool interface; the tool's name originates from the Lace Market, a historic area in Nottingham, which was the centre of the world's lace industry during the days of the British Empire. The inspiration behind the project was to produce a mapping tool that can aid the multidimensional regeneration of the Lace Market area in the centre of Nottingham by promoting business investment, tourism, and increasing local awareness of the area's heritage. The demonstrated use-case in Figure 1 aids entrepreneurs in selecting a business startup location by visually mapping relevant information about competitors, suppliers, product outlets, local talent, population demographics, as well as multimedia of cultural aspects of the area.

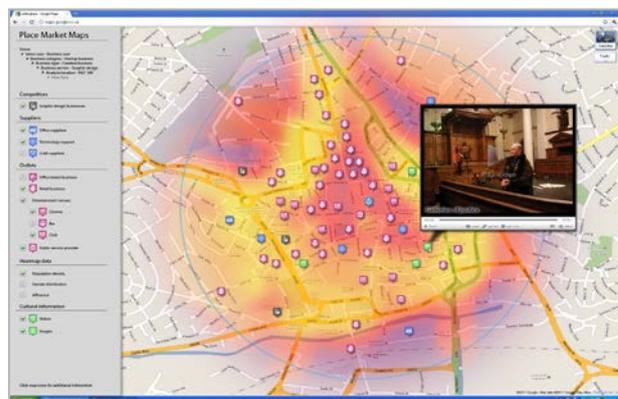


Figure 1: The PPlace Market interactive mapping tool

The tool is intended for public use; hence, naturally, a web interface is required. Several web-based mapping applications are available including Google Maps, Bing Maps and OpenStreetMap. We required a mapping application that provided a programming API, allowing developers to create an application using the mapping service as a base, which alleviates the complex task of generating the required set of map tiles and also the requirement of setting up a hosting server for the map tiles that provides a real-time streaming service for map tiles. We adopted Google Maps to build our user

interface as it provides a convenient and well-documented JavaScript API to support mapping application development.

Given the requirement to maintain cross-device and cross-platform deployment, the user interface was built using HTML5 as it supports advanced JavaScript functionality allowing applications to be developed entirely within a web page.

The application interface provides the user with a facet-based hierarchical menu that can be navigated to compile a profile that can be used to generate semantic queries tailored to the individual user's needs; for instance, a business user, looking to startup a new venture specialising in Graphic Design, may be interested in competitor and outlet information in the area around the NG1 3AY postcode.

In passive mode, the semantic repository is then queried to retrieve sets of relevant geo-tagged information to be visualised on the map, which the user can toggle off or on. The retrieved information is displayed as marker icons or multimedia HTML elements. The tool can also display data that is continuous across regions using heatmapping. Specifically, heatmaps were used to illustrate information about population density and demography.

POPULATING THE MAPPING TOOL'S SEMANTIC KNOWLEDGEBASE

The mapping tool ontology is primarily focused around the concepts of a place, user, spatial and temporal descriptors, and a service (public or private). It also contains hierarchical classifications and interrelationships between these entities, such as built spaces, businesses, employees, user opinions, media, outlets, and suppliers.

The Business Innovation Team at NTU provided the base data for our use-case. It provided information about the creative industries in the area such as location, turnover, number of employees, etc., and the tool needed to respond to complex queries with regard to business start-up, supply-outlet chain, local skills, market potential, etc.

We further enriched the base data with two public datasets, data.gov.uk's population statistics and OpenStreetMap/LinkedGeoData (LinkedGeoData 2013) along with geo-referenced information about available services (retail, health, food, tourism, etc.). The interactivity requirement of the mapping tool necessitates minimizing the query response time, hence the taxonomy of the public datasets was mapped into our ontology and they were locally cached to optimise their use.

Despite the 'openness' promise of linked data, such data is not always available in RDF, which would have significantly encouraged their reusability and/or localisation for internal consumption. With respect to data.gov.uk datasets, Prof Shadbolt and Sir Tim Berners

Lee were instrumental in the process of opening-up government data (Omitola et al. 2010). However, only a small percentage of the data is accessible as semantic datasets. In our use-case, we demonstrated the benefit of utilising geo-referenced population data to aid the decision-making process in business setup and relocation. Unfortunately, at the time of implementing the tool, the population data was available only in CSV format, necessitating a cumbersome process of manual extraction of data and geo-tagging its elements before mapping it into our semantic ontology.

Similarly, to localise the OpenStreetMap (Haklay and Weber 2008) dataset, we initially converted a MySQL dump of the dataset into N3 format using the D2RQ server (Casanave 2013), and then mapped it into our semantic ontology using SPARQL CONSTRUCT queries as illustrated below.

```
CONSTRUCT {
  ?inst a services:GPSCordinate.
  ?gpsid a services:label.
  ?inst services:label ?gpsid.
  ?lat a services:hasLatitude.
  ?inst services:hasLatitude ?lat.
  ?lon a services:hasLongitude.
  ?inst services:hasLongitude ?lon.
  ?created a services:created.
  ?inst services:created ?created .
}
WHERE {
  ?inst a vocab:gpsName.
  ?inst vocab:gpsid ?gpsid.
  ?inst vocab:lat ?lat .
  ?inst vocab:lon ?lon .
  ?inst vocab:created ?created .
};
```

Once the semantic knowledgebase was modelled, the challenge was to utilise and adapt the appropriate Semantic Web technologies to deliver the required QoS by the interactive mapping tool.

MANAGING THE GEOSPATIAL FACTORS OF SEMANTIC TECHNOLOGIES

Reconciling the Coordinate Systems

The semantic repository was populated with data from two main sources. The creative industry dataset provided rich information detailing many attributes of creative businesses including data such as turnover, employment and services provided. Data from OpenStreetMap was used to provide crowd-sourced information regarding relevant points-of-interest such as restaurants, hospitals, museums, etc. The third dataset was sourced from Data.Gov.UK to provide demographic information about the UK population. The OpenStreetMap data is gathered with the intention of mapping and provides accurate GPS based location.

The data used in our use-case was provided by a survey about describing creative businesses in the UK midlands area and referred only to UK postcodes. The postcode geographic representation was obtained from the Ordnance Survey (OS), which provides the coordinates in the OSGB36 coordinate reference system (Ordnance Survey 2013). However, the underlying spatial indexing system used by the semantic repository (OWLIM) and the mapping API (Google Maps) uses the World Geodetic System WGS84 (World Geodetic 2013). Therefore, we first deployed a third party library (Convert Coordinates 2013) to perform the complex conversion of the postcode polygon points from OSGB32 to WGS84 in order to reconcile the drastically different geo-referencing systems. Then we computed the centre location of the postcode area to use it as an approximate position for the creative businesses within that postcode.

A further challenge is that the postcode geo-reference is not accurate, so when two creative businesses lay in the same postcode, the marker position on the map is identical. The results in one marker being completely occluded by the other, and the problem increases as the number of creative businesses in a postcode increases. More precise GPS location of individual postal addresses can be commercially obtained, but this was not considered a cost-effective solution.

This inconsistency in location resolution between different businesses in the repository caused problems when trying to implement geospatial queries to find nearby businesses. The businesses gathered from OpenStreetMap have a direct link to their GPS location, while to resolve the creative industry locations the applications must traverse the relationship to the attached postcode, then to the postcode's centre point geo position. To allow both types of location specification to be queried simultaneously, the non geospatial part of a query was separated and two versions of the query were produced, one that uses the postcode location resolution and the other that looks for direct GPS location. These are then joined into a single SPARQL query using the UNION keyword.

```
SELECT ?x ?y ?bus {
  {
    ?bus rdf:type place:Business .
    ?bus place:hasLocation ?loc .
    ?loc ontoGeo:nearby(lat long dist) .
    ?loc place:hasLongitude ?x .
    ?loc place:hasLatitude ?y .
  }
  UNION {
    ?bus rdf:type place:Business .
    ?bus place:hasPostalCode ?pcode .
    ?pcode place:hasCenterPoint ?loc .
    ?loc ontoGeo:nearby(lat long dist) .
    ?loc place:hasLongitude ?x .
    ?loc place:hasLatitude ?y .
  }
}
```

Efficient Retrieval of Geo-Tagged Semantic Information

By and large, data for mapping applications is gathered for a large geographic area or region even though the user is typically interested in viewing data from a smaller geographic area of interest. Hence, the interactive mapping tool requires a scalable solution that can cope with the linear increase in the size of retrieved data with the size of the queried geographic region while maintaining a good user experience.

Semantic Spatial Indexing to Improve Query Performance

The interaction requirement of the mapping tool's visual interface imposes real-time constraints on the spatial information retrieval process from the underlying semantic knowledgebase.

Due to the geospatial nature of the application, the data was tagged with geo positioning information. Therefore, it was straightforward to determine if the data is in the user's area of interest by examining the longitude and latitude of its geo points. Using SPARQL, the standard FILTER operation can be used to select geo point results based on longitude and latitude values.

```
...
?place NtuSpatial:hasLat ?lat.
?place NtuSpatial:hasLng ?lng.
FILTER (?lat>40.0 && ?lat<42.1 &&
?lng>20.34 && ?lng<23.7)
}
```

This code segment shows how to limit the ?place geo point to geo points that are within the defined longitude and latitude points. This filter creates a bounding box aligned to longitude and latitude lines around the earth. More complex filter operations using more mathematical operations could provide different shaped areas of interest.

However, using this filtering method requires locating every geo point, retrieving its longitude and latitude positions, and then evaluating them against the filter condition. This means that every geo position stored in the repository must be evaluated to determine if it is in the area of interest. Therefore, adding more geo places to the repository would negatively impact on the query response time regardless of their spatial proximity to the query area of interest. This presents problems for a typical mapping application as it could contain a large number of geo points but only be interested in a small subset of these when looking at a specific geographic region.

The OWLIM semantic repository provides a solution by providing SPARQL extensions to allow for indexing of URIs based on longitude and latitude values. It divides the globe into a finite set of bounding boxes along the global longitude and latitude lines. The extension creates an index based on the longitude and latitude position of geographic points classified as the RDF type SpatialThing, which was provided through an

ontology that was integrated into the OWLIM repository. Our implementation of the geospatial semantic repository uses SpatialThing as the base-type for all other spatial types in our geo-spatial ontology, thus allowing all geospatial information to benefit from OWLIM's scalable indexing method. Using the base SpatialType as the root spatial property, several abstract basic spatial types were identified that could be used and combined to represent real spatial geographic information. GeoPoints were created as a simple node with latitude and longitude decimal literal properties with direct mapping to the underlying SpatialThing. Axis-aligned BoundingBoxes were a collection of two GeoPoints marking the upper left and the lower right coordinates. Concave polygons were created as an ordered list of GeoPoints.

Queries can then make use of the geospatial SPARQL extensions to filter results based on bounding box, radius from a geo point, or a polygon define by geo points. These extensions are able to determine what indexing regions are close to the user's geographic area of interest, thus resulting in more efficient geospatial information retrieval. This can significantly reduce the number of geo points that need to be evaluated, and as a result dramatically speed up query times.

Using a subset of the OpenStreetMap dataset containing the locations of businesses relevant to this application, two benchmarks were performed to evaluate the spatial query response time based on standard SPARQL filter methods in comparison to OWLIM's spatial indexing technology.

The first benchmark limited the OpenStreetMap dataset to business locations within Nottingham (1275 places), as it was geared to evaluate the technology's performance with respect to the use-case and its class of applications. The second benchmark used the OpenStreetMap data for all business locations (world-wide available data – 760,706 places) to examine scalability.

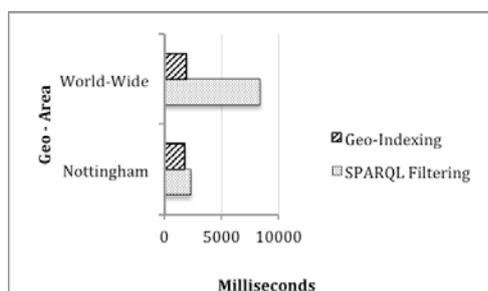


Figure 2: Query response time for the filtering and geo-indexing spatial query mechanisms

The test results in Figure 2 above showed little difference in the query response time between the two mechanisms when used to retrieve data from the Nottingham locality. However, for the large volume of geo-positions evaluated in the second test, the test results show that the OWLIM spatial indexing is clearly

outperforming the standard SPARQL filtering methods. Another important remark is that, for the size of evaluated geo-positions, spatial indexing does not seem to be significantly affected by the increase in data size. However, the SPARQL filtering method is clearly not scalable and could only be applied in scenarios with small datasets.

Optimising the Utilisation Of the Geo-Enabled Semantic Repository

OWLIM supports SPARQL query extensions that allow the retrieval of points positioned within a user-specified concave polygon. Therefore, it would be beneficial if our spatial data representation allowed querying over polygons stored in the semantic repository. However, this proved challenging as a limitation in the OWLIM polygon querying method is that the number of points within the polygon must be specified at the time of composing the query; consequently, any polygons used to geo-reference data in the semantic repository must have the same number of points as specified in the SPARQL spatial query.

In order to avoid utilising an inefficient multistage query workaround, a solution was to model all polygons with a fixed number of points. The polygon datasets acquired from the Ordnance Survey were found to be highly detailed, with some containing over a thousand points. To accommodate this, the vast majority of polygons that were smaller would have to contain hundreds of duplicated points. This was clearly a wasteful solution and would enforce an upper limit on the number of points any polygon could contain. A waypoint structure similar to the polygon was considered, but the use-case data provided did not require it and was omitted from the prototype system.

Deciding that it was not feasible to perform queries on polygons directly inside the repository, we developed a comma separated values structure that stored the list of polygon points as a single string literal. This significantly reduced the number of relationships that repository had to manage. For situations where the exact polygon area was not required for spatial computation, each polygon was preprocessed to provide its own bounding box giving a rough indication of the polygon area that could potentially be used for early escape intersection tests.

The aforementioned basic spatial types were used to represent real world spatial information such as postcode and Super Output Areas (SOA), which is a set of geographical areas developed following the UK 2001 census to facilitate the calculation of the Indices of deprivation, crime, employability, etc. The postcode and SOA data were compiled from data gathered from the Ordnance Survey and the Office for National Statistics, respectively. When the repository was populated with basic spatial representations, additional data could be added and mapped to the repository using these spatial

instances to provide a relationship to a geographic position that could be used to position the data on a map. This use-case described in this paper involved mapping geo-referenced data about creative businesses, public and entertainment services, and the local population/demography.

OWLIM's spatial indexing method breaks the world down into finite spatial regions based on longitude and latitude. Each SpatialThing is assigned to one spatial region based on where their longitude and latitude attributes lie, which allows fast retrieval of all SpatialThings that belong to a given region. The built in OWLIM SPARQL spatial query extensions can take advantage of this fact and ignore triples not associated with relevant spatial indexes, allowing a significant number of triples to be ignored during the query process. This is a clear advantage over non spatially-enabled semantic repositories as the alternative involves discarding results outside of the area of interest through the SPARQL FILTER method. This method must evaluate every possible query match then discard results during the filtering stage. For instance, when querying results in a small geographic region, all SpatialThings in the repository must be considered and then filtered out. Hence, this method does not scale well as the entire repository must be considered during a query. The spatial indexing technology provides a scalable solution as it limits the data that must be evaluated to that which is geographically close to the target area, allowing most data to be safely ignored.

To benchmark the scalability of OWLIM's geospatial indexing, a query to find all business in a Nottingham area that provided a food service was run over a repository created by extracting all businesses from the OpenStreetMap dataset. In the small dataset, only businesses from within the Greater Nottingham area were included. In the larger set, the entire UK area was used.

The smaller set contained 1135 total businesses, of which 170 were known to provide food service. The large set contained 286,236 total businesses, of which 39,567 provided a food service. For both datasets, the query looked for food services within 2.1km of the latitude longitude point of 52.9647027748771 - 1.1715030670066016 (Nottingham city centre).

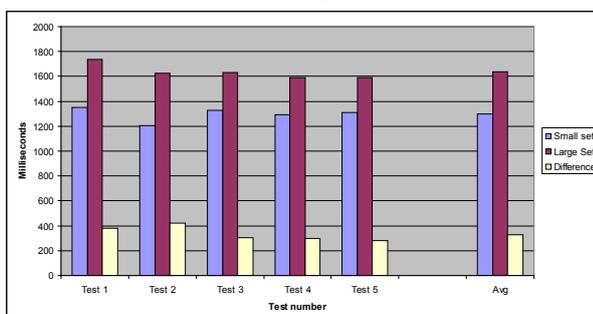


Figure 3: Evaluating the scalability of semantic geo-indexing

As Figure 3 shows the difference in query time between the two data sets is around 300 milliseconds longer for the larger data set. However, the large set is over 250 times larger than the small set. Although the query time has increased, it is sub-linear and is still within an acceptable query time for an interactive application.

Coping with Demanding Visualisation Elements - Heatmaps

Our mapping tool uses heatmaps to display data that is continuous across regions, specifically to illustrate information about population density and demography, which can be useful, for instance, to assess customer-base. A heatmap is usually shown as an overlay layer to the map the uses colour to represent the changing values of space. In most mapping applications, heatmaps can be preprocessed to display a specific dataset that is already determined before the application is run, but owing to the dynamic nature of our context-aware application, the data that needs to be conveyed to the user is context-specific and cannot be pre-computed. To overcome this problem, we leveraged the facility of the HTML5 canvas element that allows real time rendering of image inside a web page, thus allowing heatmaps to be generated at runtime and dynamically overlaid in accordance with the user context. The client-side processing of the heatmap allows the map user interface to react instantly to any changes in the user context without the need for server-side reprocessing to rebuild large heatmap tile sets.

The disadvantage of real-time processing of heatmaps is that it can incur considerable overhead on the response time of mapping applications, which is exacerbated by longer retrieval times because of the processing of the semantic metadata. We developed a technique to optimise the heatmap rendering to bring response times to within a few seconds of the users request.

The HTML5 canvas element provides a virtual image to the web application. Since heatmapping images are generally low resolution, we purposely utilised a small canvas element of 200x200 pixels to render the heatmap image while still capturing most of the important image detail. Once rendered, the heatmap can then be stretched to fit the user's display. Stretching the heatmap by several factors adds significant blur to the image, which is actually desirable as continuous spatial data rarely conforms to statistical boundaries and allowing this blurring shows a more realistic view of the statistics.

In contrast to low-resolution heatmaps, the high accuracy geo polygons captured by the Ordnance Survey contained sub-pixel resolution data. Rendering polygons with a large number of points (often around 800 per polygon) took significant time with the low-resolution heatmap unable to show this level of detail. A simple polygon reduction algorithm was used to greatly reduce

the number of points within the polygon while maintaining its general shape. Although the heatmap generation now involved an additional polygon reduction stage, the reduced rendering time was shortened significantly to bring the overall generation time to within a second.

We tested the performance gain from the polygon-reduction mechanism on a variety of HTML5-enabled web browsers. The tests were carried out on 8GB RAM Intel Core 2 Quad CPU at 3 GHz running Windows XP. JavaScript does not allow multi threading so there is no advantage from the multiple CPU cores. The browsers were Chrome version 12.0, FireFox version 3.6.13 and Safari version 5.0.5.

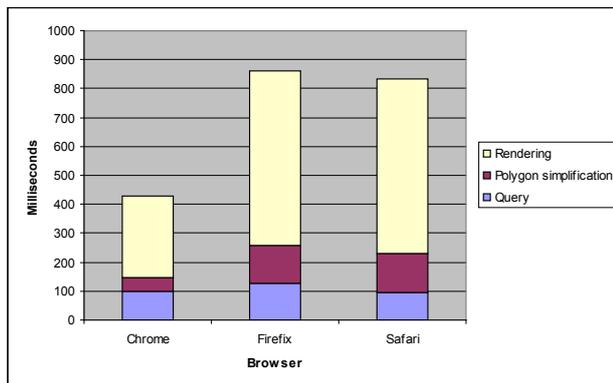


Figure 4: Evaluating the performance of the heatmapting mechanism

CONCLUSIONS AND FURTHER WORK

This paper draws on our experience in developing an interactive mapping tool in order to contribute to the efforts in tackling the challenges associated with using Semantic Web technologies in the development of context-sensitive geospatial applications. The first challenge was developing a method for the reconciliation and uniform interrogation of heterogeneous geo-tagged datasets. The next challenge was the efficient retrieval of context-specific information from the semantic knowledgebase, where we deployed a semantic repository with geospatial reasoning capabilities and optimised its utilisation by readapting its polygon query technology and introduced an efficient query transfer method with the web front-end before evaluating the scalability of the semantic repository in processing geospatial datasets. The paper also presents an efficient mechanism for real-time rendering of heatmaps, which is important for maintaining an acceptable response time in dynamic context-aware applications.

Our plans for further work include investigating the deployment of clustering algorithms for improved visualisation of highly condensed interactive maps while retaining user specific accessibility to the semantically tagged data. We also intend to experiment with the utilisation of specialised semantic geographic data

representations such as GeoRDF (GeoRDF 2013) and evaluate the use of Linked Data mapping applications such as map4rdf (map4rdf 2013) to streamline the consumption and support dynamic integration of Linked Geo Data.

REFERENCES

- Allemang, D. and Hendler, J. 2011. "Semantic Web for the Working Ontologist: Effective Modeling in RDFS and OWL." Morgan Kaufmann- 2nd edition.
- Auer, S., Dietzold, S., Lehmann, J., Hellmann, S., Aumueller, D. 2009b "Triplify - lightweight linked data publication from relational databases." In Proceedings of the 17th International Conference on World Wide Web, WWW 2009, pages 621–630, Madrid, Spain.
- Auer, S., Lehmann, J., Hellmann, S. 2009. "LinkedGeoData: Adding a spatial dimension to the Web of data." In Proceedings of Eighth International Semantic Web Conference (ISWC 2009), pp. 731-746, Chantilly, Virginia.
- Becker, B. and Bizer, C. 2009. "Exploring the Geospatial Semantic Web with DBpedia Mobile." Published in: Journal Web Semantics: Science, Services and Agents on the World Wide Web archive. Volume 7 Issue 4, December, 2009.
- Brodaric, B. 2007. "Geo-Pragmatics for the Geospatial Semantic Web." Transactions in GIS Volume 11, Issue 3, pages 453–477, June 2007.
- Brod, A., Nicklas, D., Ossietzky, C., Mitschang, B. 2010. "Deep integration of spatial query pro-cessing into native RDF triple stores." In: Proceedings of the 18th SIGSPATIAL International Conference on Advances in Geographic Information Systems (GIS '10), pp. 33-42 , San Jose, California.
- Brown, D.; Standen, P., Evett, L., Battersby, S., Shopland, N. 2010. "Designing Serious Games for People with Dual Diagnosis: Learning Disabilities and Sensory Impairments." In P. Zemliansky, & D. Wilcox (Eds.), Design and Implementation of Educational Games: Theoretical and Practical Perspectives, pp. 424-439.
- Cantador, I., Bellogin, A., Castells, A. 2008. "News@hand: A Semantic Web Approach to Rec-ommending News." ADAPTIVE HYPERMEDIA AND ADAPTIVE WEB-BASED SYSTEMS. Lecture Notes in Computer Science, 2008, Volume 5149/2008, pp. 279-283.
- Clough, P., Tang, J., Hall, M., Warne, A. 2011 "Linking archival data to location: a case study at the UK National Archives." Aslib Proceedings: New Information Perspectives, 2011, Vol.63(2/3), p.127-14.
- Convert Coordinates between WGS-84 and OSGB36. 2013 (last accessed Feb 2013). <http://www.movable-type.co.uk/scripts/latlong-convert-coords.html>.
- Cory Casanave, C. 2013 (last accessed Feb 2013). "Designing a semantic repository." <http://www.w3.org/2007/06/eGovdc/papers/SemanticRepository.pdf>
- Geo-Spatial Indexing in Owl. 2013 (last accessed Feb 2013). <http://www.ontotext.com/owlim/geo-spatial>
- GeoRDF - An RDF compatible profile for geo information. 2013. (last accessed Feb 2013) <http://www.w3.org/wiki/GeoRDF>
- Haklay, M. and Weber, P. 2008. "OpenStreetMap: User-Generated Street Maps." IEEE Pervasive Computing Magazine, Volume: 7 , Issue: 4, pp. 12-18 Oct.-Dec. 2008.

- LinkedGeoData. 2013 (last accessed Feb 2013). <http://linkedgeodata.org>
- map4rdf - Map viewer of RDF with Geometrical Information. 2013 (last accessed Feb 2013). <http://oegdev.dia.fi.upm.es/map4rdf/>.
- Miron, A., Gensel, J., Villanova-Oliver, M., Martin, H. 2007. "Towards the Geo-spatial Querying of the Semantic Web with ONTOAST." Web and Wireless Geographical Information Systems. Lecture Notes in Computer Science Volume 4857, 2007, pp. 121-136.
- Omitola, T., Koumenides, T., Popov, I., Yang, Y., Salvadores, M., Szomszor, M., Berners-Lee, T., Gibbins, N., Hall, W., Schraefel, M., and Shadbolt, N. 2010. "Put in Your Postcode, Out Comes the Data: A Case Study". In: ESWC. Ed. by L. Aroyo, G. Antoniou, E. Hyvönen, A. ten Teije, H. Stuckenschmidt, L. Cabral, and T. Tudorache. Lecture Notes in Computer Science 6088. Springer, June 2010, pp. 318-332.
- Ordnance Survey osgb36 coordinate system. 2013 (last accessed Feb 2013) <http://www.ordnancesurvey.co.uk/oswebsite/gps/information/coordinatesystemsinfo/guidecontents/guide5.html>. Last accessed.
- Thakker, D., Osman, T., Gohil, S., Lakin, P. 2010. "A Pragmatic Approach To Semantic Repositories Benchmarking". ESWC'10 Proceedings of the 7th international conference on The Semantic Web: research and Applications - Volume Part I. pp 379-393.
- World Geodetic System 1984. 2013 (last accessed Feb 2013). http://earth-info.nga.mil/GandG/publications/tr8350.2/tr8350_2.html.

AUTHOR BIOGRAPHIES



Taha Osman is a principle Lecturer at the College of Science and Technology, Nottingham Trent University. He gained his PhD in Fault-Tolerance of Distributed Computing Systems from the same University in 1998 and immediately joined its academic staff. Dr Osman leads the Semantic Web research network at the department of computing and informatics and his research interests include Semantic Web, Knowledge Engineering and Intelligent Information Retrieval, Intelligent Agent Systems.



Luke Shires studied at Nottingham Trent University for a BSc in Computer Science graduating in 2011. He also received the Vice Chancellors award for his involvement in founding and maintaining the Nottingham Trent Developers Society, which creates work and placement opportunities for students by enabling them to learn skills in game development. He is primarily interested researching how games and game

technologies can be applied to solve real world problems. Luke is currently working towards a Ph.D. looking at the application of markerless tracking technologies to improve games for upper extremity rehabilitation following stroke. Prior to beginning his Ph.D. he has worked on several projects including the use of applying semantic knowledge bases to enable intelligent mapping solutions for helping cognitively impaired users during route learning and using games technologies to create a real time simulation of visual impairments.



Tope Omitola is a Research Fellow at the Web and Internet Science Group in electronics and computer science at the University of Southampton. His research interests include the Linked Data and Semantic Web, Web science, and knowledge technologies. Omitola is a graduate of King's College London and Jesus College Cambridge.



Nigel Shadbolt is a professor of artificial intelligence and head of the Web and Internet Science Group in electronics and computer science at the University of Southampton, and is the co-director of the Open Data Institute in London. He is also an adviser to the UK government on open data and a member of the UK Public Sector Transparency Board. His research interests include the Semantic Web, Web science, and knowledge technologies. Shadbolt has a PhD in AI from the University of Edinburgh.



Following a degree in clothing and textile management, **Dr. Jeremy Hague** worked for nine years within a multi-disciplinary action research team at Nottingham Trent University. Funded solely through consultancy and research income, the team specialised in promoting innovation through employee involvement and new forms of working; particularly team-based manufacturing. During this time Jeremy pioneered the use of forum theatre in business settings, which stimulated creative development across all levels of an organisation. Jeremy currently leads the Business Innovation Team within NTU's Commercial Directorate. The team supports the development of commercial products, relationships and income through company networks and funded programmes. Jeremy also represents the University in the Officers Group of D2N2, the Local Enterprise Partnership for Derbyshire and Nottinghamshire