

Automating Content Generation for Large-scale Virtual Learning Environments using Semantic Web Services

Ian Dunwell¹, Panagiotis Petridis¹, Aristos Protopsaltis¹, Sara de Freitas¹, David Panzoli¹ and Peter Samuels²

¹ Serious Games Institute, Coventry University, UK
{idunwell, ppetridis, aprotopsaltis, sdefreitas, dpanzoli}@coventry.ac.uk
² Faculty of Engineering and Computing, Coventry University, UK
{psamuels@coventry.ac.uk}

Abstract. The integration of semantic web services with three-dimensional virtual worlds offers many potential avenues for the creation of dynamic, content-rich environments which can be used to entertain, educate, and inform. One such avenue is the fusion of the large volumes of data from Wiki-based sources with virtual representations of historic locations, using semantics to filter and present data to users in effective and personalisable ways. This paper explores the potential for such integration, addressing challenges ranging from accurately transposing virtual world locales to semantically-linked real world data, to integrating diverse ranges of semantic information sources in a user-centric and seamless fashion. A demonstrated proof-of-concept, using the Rome Reborn model, a detailed 3D representation of Ancient Rome within the Aurelian Walls, shows several advantages that can be gained through the use of existing Wiki and semantic web services to rapidly and automatically annotate content, as well as demonstrating the increasing need for Wiki content to be represented in a semantically-rich form. Such an approach has applications in a range of different contexts, including education, training, and cultural heritage.

Keywords: semantic web applications, virtual learning environments, information systems applications

1 Introduction

Increasingly, web content is represented using semantic metadata formats which support the compilation and interlinking of information. One of the key advantages to such approaches is the ability to query and search this information using novel methods, such as relating 'geocoded' data to other web-based information repositories. Geocoding (for a comprehensive summary: see Goldberg, 2007), the process of transcribing named places to absolute geographic coordinate systems, has allowed

information to be queried in a host of different ways in various application areas; including public health (Rushton et al., 2006) and epidemiology (Krieger, 2003). Through the integration of these services with information sources such as Wikipedia, the potential exists to link both semantic and non-semantic Wiki content to real world locales.

In this paper, we explore an application of this combination of services to virtual learning environments. Whilst many learning environments currently rely on subject matter expertise for content generation and validation, such an approach is time-consuming, costly, and often involves the duplication by-hand of information already available from other sources to suit the format and context of the learning environment. Therefore, we consider a potential solution to be the use of geocoding data to identify an article held on a Wiki, and hence rapidly and autonomously annotate large environments, which can mirror real-world locales based in the past or present. This combination of a virtual world with a dynamic, editable, and peer-reviewable Wiki-based data source has immediate advantages in being able to support exploratory, peer-based learning models without requiring substantial input and guidance from subject matter experts. The source of the information driving the annotation in the proof-of-concept we describe in Section 4 is Wikipedia; However besides providing a demonstration of how semantic services can bridge into non-semantic data sources, this proof-of-concept highlights the long-term benefits that could be achieved by using fully semantic representations of information in these services.

Following an introduction to the state-of-the-art in Section 2, we go on to describe in Section 3 several concepts which underpin the implementation of systems using geocoding web services to provide content for learning environments which can be fed back to users in a range of novel and innovative ways. Section 4 details an implemented proof-of-concept using this approach, which uses the Rome Reborn (www.romereborn.virginia.edu) model alongside the GeoNames service (<http://ws.geonames.org>) to provide information to a user navigating the model in real-time. This proof-of-concept shows a simple approach to feeding information back to the user that can be expanded upon, and to this end we discuss the challenges faced in creating more sophisticated environments and learning experiences as well as the potential for future work in Sections 6 and 7.

2. Background

Many existing approaches towards creating virtual learning environments utilise the knowledge of subject matter experts to annotate content by hand. The integration of Rome Reborn with Google Earth (<http://earth.google.com/rome/>), for example, uses such an approach. Other applications in cultural heritage, such as the ARCO system (White et al., 2004), seek to allow curators or developers to create a dynamic virtual exhibition through the use of XML-based procedural languages, allowing dynamic modelling capabilities to be realised in a virtual scene. This technique enables the development of dynamic, database-driven Virtual Worlds, created by building

parameterised models of virtual scenes based on the model and the data retrieved from the database (White et al., 2009).

The MESMUSES project (Meli, 2003) highlights an interest amongst teaching institutions to provide learners with 'self-learning' environments, providing them with an opportunity to explore various knowledge spaces (i.e. digital information on museum artefacts) in a free-roaming virtual world. To this end, the MESMUSES project demonstrated a system that accesses cultural information through the novel concept of 'knowledge itineraries'. These itineraries represent a series of thematic paths that visitors can choose to follow, and when doing so various resources are offered to them including examples and explanations. Furthermore, the system, with the use of personalization methods, offers different knowledge domains to different categories of visitors. Similarly, the ART-E-FACT project (Marcos, 2005) proposes that the use of the semantic web can enable learning institutions to make cultural content available to researchers, curators or public in an increasingly meaningful and user-centric fashion. Marcos et al. suggest that the use of digital storytelling and mixed reality technologies can also create a new dimension for artistic expression. Within cultural heritage applications, therefore, there are multiple benefits that can be gained from using semantic technologies, such as the potential to gather data from across the web, filter this data using metacontent, and present it to the user in a dynamic and customisable fashion.

Outside of the specific domain of cultural heritage, attempts have also been made to annotate virtual environments to aid user navigation. For example, Van Dijk et al. (2003) demonstrate an approach using geometric and spatial aspects of the virtual worlds and its objects. Within a map of the environment, landmarks are added to identify various locales, supported by a personal agent with knowledge about the current position, visual orientation of the visitor, objects and their properties, geometrics relations between objects and locations, possible paths towards objects and locations, routes to the user and previous communications. By comparison, Pierce and Pausch (2004) present a technique for navigating large virtual worlds using place representation and visible landmarks that scale from town-sized to planet-sized worlds, whereas Kleiner mann and colleagues (2008) propose methods in which the domain expert annotates the virtual world during creation, suggesting that since the world is being created using ontologies, the resultant semantic annotation will be richer. Navigation can then exploit these semantic annotations using a search engine - assuming, however, that the world has been created and annotated using this method.

Fundamentally, the approaches listed in this section primarily rely on direct intervention from either designer or subject matter expert to annotate an environment. Whilst it is undoubtedly the case that such an approach allows for certainty in the accuracy and validity of content, this approach also has drawbacks in requiring human resources for not only creation but also maintenance, since unless content is updated regularly, the experience is unable to retain users for long periods of time as content is gradually consumed. In work being undertaken at the UK Open University through the Luisa project (Mrissa et al., 2009, Dietze et al., 2008), advanced semantic web searching and acquisition techniques are being used to personalise and filter

information dynamically in real-time. This allows for complex reversioning and acquisition of data provided to the user, and can support more complex educational requirements for wider ranges of learner groups. Potentially, in the educational domain, this advance in intelligent querying is supported by service orientated architectures and may support advanced educational scenario developments that may be tailored to individual user requirements. For all these applications, extending the databases at the core of each system to include other web-services and sources of information could enhance both the volume and quality of content in virtual exhibitions and environments, easing user navigation and creating deeper, more compelling learning environments.

Wiki technology offers a basis for supporting such approaches. Since content can be simultaneously generated and peer-reviewed by a large base of users, large volumes of data may be generated with less expense or increased speed when compared to individual subject matter experts. Although validity can pose a concern, semantic representation simplifies identification and comparison between different data sources and can therefore aid designers in addressing this concern, and throughout the background literature an increasing motivation to create virtual worlds which are annotated in increasingly user-oriented fashions can be observed. In the remainder of this paper, we describe an approach to extending and automating annotation for large virtual environments which are based on real-world locales. This approach focuses on the use of geographic information services together with semantic web and Wiki data, to obtain data on points within the world which can form the basis of more complex filtering and mining techniques such as those proposed by the Luisa project (Mrisa et al., 2009). In the next sections, through a demonstrated proof-of-concept, we show the potential of such an approach to quickly and automatically annotate a virtual world with a large volume of information from a Wiki.

3. Automating Content Acquisition via Geocoding

In this Section, we describe in general terms the concepts behind the integration of web services such as Wikis with virtual worlds using geocoding. Whilst the proof-of-concept described in Section 4 focuses on the combination of Wikipedia with historical environment via the GeoNames service, this reflects a more general approach consisting of three fundamental steps. Firstly, coordinates in virtual space must be converted to a form suitable for input into the wide range of web-based GIS systems and databases. Secondly, the information obtained must be filtered to ensure relevance to the locale, time period, and usage scenario. Finally, this information must be presented to the user in an appropriate and coherent fashion. This section discusses these three issues in some detail; although the large number of web services, coupled with their diversity, makes generalisation a challenge, an attempt is made to describe the solution in as general terms as possible.

3.1 Coordinate Conversion

GIS systems commonly take coordinates as longitudes and latitudes. By comparison, virtual worlds contain arbitrary, typically Cartesian, coordinate systems. The conversion of these virtual coordinates to a real-world location can be simplified for relatively small areas with little variation in elevation (e.g. cities) by approximating the longitude and latitude as a Cartesian system. In this case the translation between a point x_0, y_0 on a virtual plane and real-world geographical location at longitude and latitude x_t, y_t can be expanded into simultaneous equations of the form:

$$\begin{aligned}x_t &= \alpha \cos(\theta)x_0 - \alpha \sin(\theta)y_0 + t_x \\y_t &= \alpha \cos(\theta)x_0 + \alpha \sin(\theta)y_0 + t_y\end{aligned}\tag{1}$$

Where α , θ , t_x and t_y describe the rotation, scaling, and translation between the two coordinate systems. This assumes both coordinate systems are aligned along the z-axis; if this is not the case then this can easily be accommodated by introducing the z coordinate in the above equation, although it should be noted the vast majority of 3D models have an immediately apparent vertical axis around which the virtual coordinate system can be defined. Solving these equations simply requires a set of virtual points and their real-world equivalents. The accuracy of the solution is, therefore, predominantly dependent on the accuracy of this point-set, and specifically, how accurately each real-world point identified is mirrored in virtual coordinates (whilst floating-point accuracy will also affect the solution, its impact is negligible in comparison). This is a challenge, since limitations in the fidelity of the virtual space can affect how accurately points can be identified and mapped to real-world points. Similarly, the resolution of GIS data sets limits the accuracy with which locations can be defined, as do difficulties in defining the centre of a building.

There are two potential approaches to increase accuracy. The first is to use GPS hardware to more precisely identify real-world points corresponding to virtual ones, although this is often impractical since it requires real-world presence. Hence, a more desirable solution may often be to increase the number of points sampled and average multiple solutions to Equation (1). It should be noted, though, that a 3-point sample proved adequate for the example given in Section 5, since the level of accuracy required is dependent on how tightly-distributed the information points in queried web services are, and this distribution (in the case of the systems used within the case study) has all points at least 20m apart. Given the capability to rapidly translate points in real space to virtual space and vice-versa, an immediate question is how to best identify the geographic point(s) which best represent the user's interest. It is possible to simply convert the position of the viewpoint in virtual space to a GIS location and provide data on that location, although this is only likely to represent the actual point of interest if the user is above the object looking down. Similarly, getting the user to directly intervene and click the point is a solution, but more seamless integration between information and virtual world could be achieved through saliency mapping and scene analysis, so as to base the selection of objects on the perceptual traits of the user. Studies of related problems in computer science, ranging from interest

management to selective rendering (Dunwell and Whelan, 2008, Sundstedt et al., 2005) have demonstrated that proximity alone is not an accurate measure of salience. A coarse solution is to generate a pick-ray (line in virtual space) from the centre of the viewpoint into the scene and select the first object it intersects, although more sophisticated approaches such as that of Yee et al., (2001) show the potential gains from more accurately modelling how users perceive, and interact with, three-dimensional scenes. Additionally users' historical behaviours can be used to detect motion of the viewpoint around objects and other traits that may indicate interest.

3.2 Filtration

Foremost amongst the advantages of semantic content representation in this context is the ability to filter information on a semantic level. Besides from the spatial filtering achieved through the use of a geocoding service, data can be filtered according to criteria such as date, particularly relevant in the case of Ancient Rome. In our case, we consider both the filtering of non-semantic data held on Wikipedia through a conventional keyword search, and also the use of DBpedia datasets to provide a semantic version of content. Due to the simplicity of the filtering task in this case, given the straightforward application of date and spatial filters, it is possible to use either semantic or non-semantic versions of Wikipedia content; however, in the longer term, more sophisticated applications and increased content volume will benefit from the advantages semantic search techniques bring.

3.3 Presentation

Finally, the filtered information must be presented to the user in a coherent form. This is of particular interest to developers of virtual learning environments, who often seek to present and represent information in new and innovative ways. Simple approaches can include the return of text and images within the application interface; the form in which these are presented to the user can be tailored by the designer to meet practical and pedagogic concerns. It is possible, for example, that this information could be used to create questions, (e.g. requiring the user to name a location then testing against available web data), images showing real-world equivalents of virtual locations, and other learning objects. This area has strong potential for future work: a more advanced method may be the use of virtual characters, for example as Vygotskian learning partners (Rebolledo-Mendez et al., 2009). In this context, the information returned could form the knowledge base of the partner. The potential to autonomously provide data to an artificial intelligence from background web-services can enable these characters to behave more realistically and dynamically; conversational agents such as those of Daden (<http://www.daden.co.uk/chatbots.html>) have demonstrated this capability in second-life, and extending this technique to large-scale virtual recreations of real-world locales using the geocoding approach is an interesting avenue for future work.

4 Case Study

In this section, we present our working proof-of-concept which integrates the key concepts introduced in Section 3 into a working software platform. We present this firstly in terms of the high-level architecture, applicable to any real-world model, and then secondly with respect to a case study using a large-scale model of Ancient Rome developed as the principal output of the Rome Reborn project (Guidi et al., 2005).

4.1 System Architecture

Our prototype solution to automated annotation, using the principles described in Section 3, is described in this section. The key processes - coordinate selection, return filtering and conveying information to the user – are achieved through integration of the JME and LOBO APIs within a proprietary core engine. Input and output are handled again via an external API, in this case Java's Swing. On a more general level, this can be seen as a discretisation of the central tasks of rendering the model in real time, providing a web-interface, and providing a common user interface which integrates both the textual data retrieved from searches with the rendered three-dimensional scene. Therefore, these components may be interchangeable to suit other hardware platforms (such as mobile devices) or languages as required. We therefore focus our discussion on the core engine, shown in Figure 1.

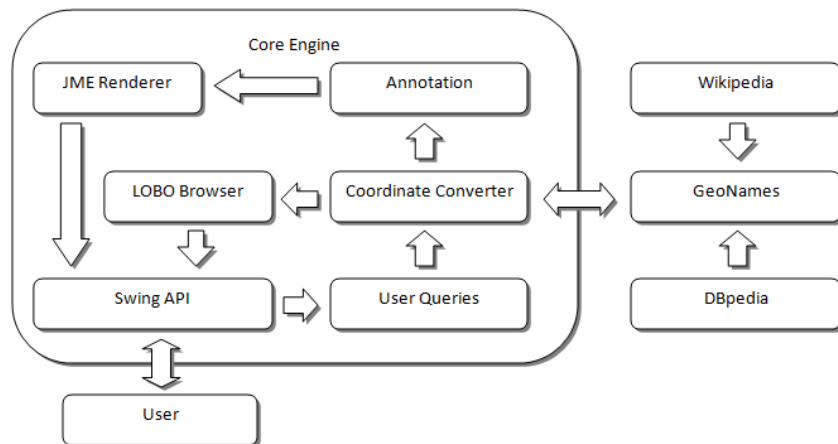


Fig. 1. Architecture using semantic web services to provide educational content and annotation for virtual representations of real-world locations both automatically and through user request

Automated annotation is achieved prior to run-time by automatically generating queries based on the specified real-world coordinates of the model. These allow a large volume of data to be captured and represented as information points within the virtual world, as discussed later in this section (see Figure 3). These points are passed to the rendering engine and hence used to create content, derived from processing of the raw XML data returned from web-queries. In this case, web-queries are directed

to the GeoNames integration with Wikipedia (<http://www.geonames.org/>) to obtain semantically-annotated data in the form of XML in response to input latitudes and longitudes. This is then processed and used as the basis for the construction of Aand minimises unnecessary web traffic generated by multiple queries with identical returns.



Fig. 2. Real-time rendering of the Rome Reborn model

Within the architecture illustrated in Figure 1, GeoNames is used as a bridge into both the non-semantic Wikipedia, and the semantic DBpedia service. Through the conversion of coordinates from virtual to real space (and vice versa), annotation is generated and passed to the rendering engine, and user queries are also handled using an integrated browser supported via the LOBO API. These allow the user to request data on their current location, which is spatially filtered via the viewport orientation, and temporally filtered by keyword and date comparison to the semantic data returned. This allows the returned XML to be filtered to only provide articles relevant to Ancient Rome since the more general data provided by the GeoNames service is not temporally restricted. *Query generation* both with regards to automated and user requested information is performed by constructing an HTTP request and passing the filters to the GeoNames service within the URL via CGI scripting.

The efficacy of the keyword-based filtering system is dependent on the richness of the semantic annotation of the returned data. In our proof-of-concept, this is restricted

to that data returned by the GeoNames service: the Wikipedia article title, synopsis, and date if provided are queried, although the date is often held non-semantically within the synopsis. The increasing drive towards creating semantic Wiki technology will have long-term benefits for this approach, enabling more accurate filtration as well as allowing information to be returned in more versatile forms. The prototype currently returns information to the user by filtering supplied XML via an XSL stylesheet: future work described in Section 6 will explore the use of this information to drive dialogic interactions with virtual characters. In the next section, we describe the application of this architecture to the model of Ancient Rome shown in Figure 2.

4.2 Rome Reborn

The proof-of-concept developed as part of this research integrates the GeoNames service, Wikipedia, and the Rome Reborn model (Guidi et al., 2005) using the architecture described in Section 4.1 to provide instant and automated semantic annotation of the 3D model with over 250 articles. A Java application was developed which allows the user to navigate through the model in real-time. Figure 2 shows the real-time Java/OpenGL render of the whole model, which includes prominent features such as the Colosseum, Basilica of Maxentius, Tiber River and Ludus Magna. The user is able to navigate through the model using keyboard and mouse in a standard first-person interaction paradigm, with their input affecting the position of the view point in virtual space. A selection of sliders allows cosmetic effects such as lighting and fog to be changed dynamically.

Performance is achieved through a discrete level of detail (LOD) approach, which segments the city into areas with three levels of detail, selected dependent on distance from the viewport. This allows for flexible management of performance by controlling the distance at which various levels of detail are selected. Further performance can be achieved by manipulating the far clip plane, ensuring ~30fps can be maintained. With respect to information retrieval and processing, the use of a local cache, coupled with the fact retrieval is either done prior to user interaction (in the case of annotation) or limited to the rate at which the user explicitly requests information (rarely more than once per second), results in the ability to respond to a user request immediately.

Virtual to real coordinate conversion was achieved using the technique described in Section 3, with three reference points taken at the most prominent structures and joined with the real-world equivalents as defined on Google Earth to form point sets. As the user moves the view point, requests to the Geocoding service are automatically generated by performing this translation on the view point location. Hence, this prototype uses purely distance-based measures of content relevance - the nearer the view point is to a location, the more likely it is to be returned by the Geocoding search as the closest point within the database, we refer to this as 'proximity searching'. To create the information service, a second pane is added (Figure 3) which uses the LOBO API to add a pure Java web browser within the application. Information points are loaded into the scene as simple geometric objects by querying

the Geocoding service for the 250 points nearest the centre, converting these coordinates to the virtual coordinate system, and adding them to the world. When a user clicks on one of these points in the 3D space, the Geocoding service (or cache) returns XML data centred on that point, which includes a title, summary, and link to the Wikipedia article nearest the queried latitude and longitude. This XML is, in this case, filtered through direct parsing hard-coded into the application, as well as a generic XSL style-sheet which formats the data to present it to the user. The link to the style-sheet is added by directly inserting a line to the XML during processing. If the user wants further information on a point, a link is provided to the Wikipedia article. It thus provides a simple example of the return filtering process. Local caching is used as shown in Figure 1 to minimise unnecessary web traffic. The solution demonstrates a simple proof-of-concept, showing all three components of the model described in Section 3 working to provide autonomous and dynamic information to the user as they explore the model.

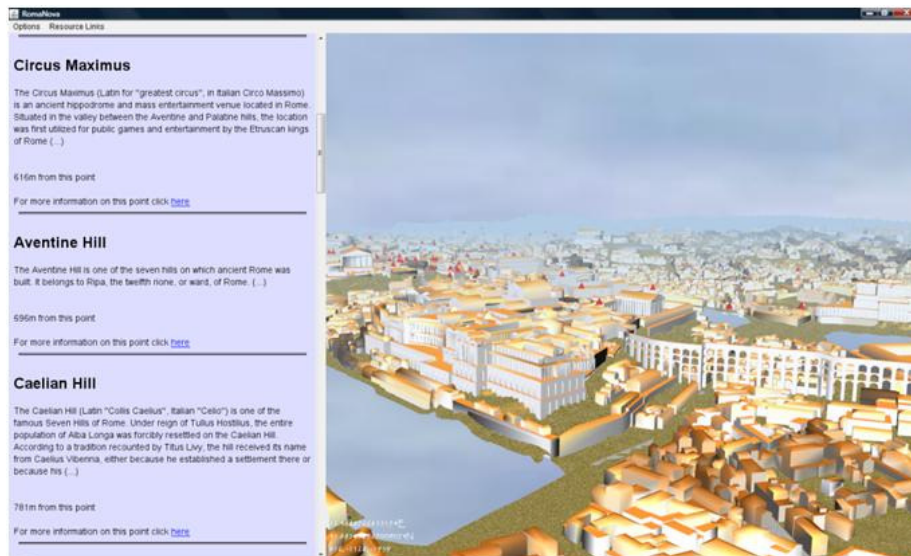


Fig. 3. The interactive application with information queried using GeoNames (left) rendered using XML/XSL

5. Discussion

The collaborative and dynamic nature of Wikis makes them an interesting area for pedagogic design. A central concept to Wikis is the notion of users as producers and evaluators, as well as consumers of content, and exploring this potential in a 3D virtual space within the proof-of-concept suggests several issues that need to be tackled. Firstly, the abstraction of the Wiki paradigm from the familiar interface may result in users failing to recognize it as such, and hence behave only as content consumers. Supporting the transition of the Wiki concept to different representational media requires that users continue to interact as content producers as well as

consumers, and this has repercussions for how user interaction is modeled and how interfaces are designed. A potential pedagogic advantage of an environment such as that developed is the facilitation of experiential (Kolb, 1984) or exploratory (de Freitas and Neumann, 2009) models of learning. The ability to immerse a learner within a detailed 3D environment, and utilize semantic services to provide detailed content and the ability to autonomously handle information requests and provide increasingly dynamic environments may have direct benefits to learning transfer. In the context of this paper, whilst preliminary qualitative work reinforces this hypothesis, significant challenges exist in defining how virtual learning environments can be accurately assessed. This is particularly the case where principal desired outcomes lie beyond the simple recollection of facts – a control study of virtual versus real scenarios may offer some insight in this respect, but would fail to reflect to typical role of virtual worlds as augmenting, rather than replacing, existing instructional techniques.

Furthermore, from an educator's perspective, one of the most prominent issues arising from the application of techniques such as those described within this paper is the transition of subject matter experts from content creators to content evaluators. As collaborative web-based knowledge bases expand, existing subject matter expertise is becoming increasingly available and accessible across disciplines. Furthermore, advances in how this information is represented (e.g. metadata formats) allow for versatility in how it is presented. Therefore, the role of educators and subject matter experts when designing learning environments increasingly becomes centred on the definition of information filters and presentation formats, so as to ensure that information is conveyed to learners in a valid and appropriate manner. Similarly, as virtual learning environments move towards experiential and situative pedagogies (Egenfeldt-Nielsen, 2007), and feature increasingly sophisticated intelligent tutors and characters, pedagogic design must support both learning within the environment itself as well as the integration of such environments across the curriculum as a whole. Whilst a key advantage of the technique described is that is capable of supporting exploratory learning, this infers that the usual cautions that should exist when creating exploratory learning experiences need to be considered. Foremost amongst these is the potential for the learner to deviate towards activities that fail to align with the desired learning outcomes. To overcome this, guiding the learner within the environment can be done in subtle ways using perceptual cues (Dixit and Youngblood, 2008), and integrating such models more fully with the methods described in this paper may be one avenue for creating experiences which guide the learner without constraining them in a perceivable way. The introduction of 'game elements' such as objectives, missions, or timed activities also has potential for increasing learning transfer when compared to open simulations (Mautone, 2008), and can also support more structured learning experiences within open, exploratory, environments.

Development of the working proof-of-concept identified a number of technical challenges. Firstly, although the visualization of the model itself is somewhat beyond the scope of this paper, rendering a large environment in real time is computationally intensive, and the overheads incurred by attempting work on such a scale can often

prove restrictive: for example, annotating the environment incurs additional performance overheads, and the level of annotation must be carefully balanced so as not to overload the user with information. Secondly, whilst GeoNames provides one potential link to semantic services, it is by no means the only such link which could be utilized. In our case study, we have demonstrated the case of using GeoNames to bridge into the non-semantic Wikipedia – however, careful consideration and selection of appropriate services is essential. In our case, bridging into Wikipedia was beneficial due to it containing the fullest collection of relevant information, though this is likely to change rapidly as services such as DBpedia offer increasing volumes of pure semantic content. In turn, this can be more fully utilized in different forms to add more depth to and variety to how information is represented and conveyed to the user.

One of the main issues regarding the use of automatically-generated educational content derived from the semantic web is the difficulty in ascertaining its accuracy and validity *in lieu* of a human expert. Doing so autonomously in a way which guarantees validity remains a substantial challenge, compounded by the need to also filter this data according to user needs. Furthermore, the dynamicism of web-based services and information, and the subsequent implications this has for instructional and educational programmes which are typically designed as repeatable courses (information will change over time as its web-based sources are edited, expanded, or removed), are an important consideration. Despite these drawbacks, the long-term advantages of evolving the techniques described within this paper are numerous: the large volume of freely available content allows for large volumes of relevant information to be rapidly integrated into the model, and at negligible cost compared to proprietary content development. For large, expansive areas, such as the city-scale model used in the case-study, these methods allow for more comprehensive and rapid annotation of content.

The more general challenges faced in developing and applying systems that integrate virtual worlds, web-based information, and intelligent tutoring for learning purposes must be addressed on both technological and pedagogic levels. This paper has presented several key technical issues, although the underlying pedagogy and, more fundamentally, purpose, of learning environments must always be a consideration in their long-term development and implementation. In the next section, we discuss some avenues for future work.

6. Conclusions and Future Work

This paper has further demonstrated the potential for the integration of information obtained from web-services into virtual learning environments. The solution is generically applicable: any real-world locale could be implemented using the approach described in the case study by simply changing the point set used when solving the equations presented in Section 3. The approach therefore has potential applicability to a wide range of learning environments, allowing developers to rapidly annotate content with information from a wide range of sources automatically and

dynamically. As mentioned in the previous section, developing pedagogies that realise the potential of this technology is a key area for future development. Open and exploratory environments may be capable of immersing and engaging learners, but if learning requirements are not met, they have limited use. Comparative evaluation of the various approaches that can be used to address this issue is a particularly relevant area for future study. The notion may also be introduced of using the results of queries to generate new queries autonomously, for example, Koolen (2009) demonstrate the potential use of Wikipedia pages, obtained as demonstrated in this paper through GIS coordinates, for book searches. Additionally, domain expertise can be modelled (White et al., 2009) to generate improved results.

On a more technical level, future work will focus around the latter two stages of the model, improving how information is sourced, filtered and conveyed to the user. This is a significant research challenge in many areas; in particular, using the information as a knowledge-base for intelligent tutoring systems driving virtual characters that behave and interact naturalistically requires advances in natural language processing, dialogue construction and pedagogy. Attempts to provide characters with a full, detailed knowledge-base must also consider the development of web-services as well as methods of content annotation to facilitate simpler integration. The integration and presentation of the information within the world in innovative ways is also an interesting area for future work, for example, weather patterns and air quality may be visualised in virtual spaces which provide information on real-world environments and systems.

A more sophisticated approach may be to embed semantic information into a virtual character as a knowledge base that the character can use to drive their own behaviour, and further enhance interactions with human users. This has the potential to be compatible with the hybrid architectures often used to control virtual humans (Conde and Thalmann, 2004, Donikian and Rutten, 1995, Sanchez et al., 2004), which are typically responsible for both low-level control of the agent such as navigation and obstacle avoidance, but also more complex interactions with the environment and other characters. Many challenges exist in realizing such techniques effectively: not only does it require a substantial amount of supplemental work, for example to animate and visualize the character, but increasing levels of realism and believability also imply increased challenges in creating characters able to adapt and behave dynamically. Consequently, the knowledge base may be limited to a specific context. Such techniques offer long-term potential for the application of semantic web-technology within virtual learning environments in a host of novel and interesting ways. However, the current state-of-the-art is often constrained by the large number of interrelated technical advances in many disciplines that are required to achieve these long term visions. In the next section, we describe a model which provides both a working, applicable approach for annotating worlds using existing technologies, whilst also accommodating and contributing towards these longer-term visions.

Finally, in this paper, we have focused on the user as a consumer rather than generator of content. Future potential exists for the use of virtual worlds to also allow users to create semantic Wiki content in innovative ways, by interlinking and

interacting with objects in virtual space. There is also an increasing demand for 3D content that is itself semantically-annotated (Spagnuolo and Falciendo, 2009). The methods described in this paper could provide a basis for allowing semantic annotation to be created for models such as Rome Reborn automatically by inverse geocoding. More significantly, as semantically represented 3D content becomes increasingly available, the ability to compose worlds autonomously using an integrated approach that adds content to the world based on its meaning and relevance to the learner, promises the potential to create sophisticated, adaptive learning environments.

References

1. Brusilovsky, P., Peylo, C.: Adaptive and intelligent web-based educational systems. *International Journal of Artificial Intelligence*. 13, 159--172 (2003)
2. Conde T., Thalmann, D.: An artificial life environment for autonomous virtual agents with multi-sensorial and multi-perceptive features. *Computer Animation and Virtual Worlds* 15, 311—318 (2004)
3. de Freitas, S., and Neumann, T.: The use of 'exploratory learning' for supporting immersive learning in virtual environments. *Computers and Education*, 52, 343—352 (2009)
4. Dietze S., Gugliotta, A., Domingue J.: Towards context-aware semantic web service discovery through conceptual situation spaces. In: *CSSSIA '08: Proceedings of the 2008 international workshop on Context enabled source and service selection, integration and adaptation*. 28, 1—8 New York, NY, USA. (2008)
5. Donikian, S., Rutten, E.: Reactivity, concurrency, data-flow and hierarchical preemption for behavioral animation. In: *Fifth Eurographics Workshop on Programming Paradigms in Graphics*. 137—153 Springer-Verlag (1995)
6. Dunwell I., Whelan J. C.: Spotlight interest management for distributed virtual environments. In: *14th Eurographics Symposium on Virtual Environments (EGVE08)* 1, 56--64 (2008)
7. Dixit P. N., Youngblood, G. M.: Understanding information observation in interactive 3d environments. In: *Sandbox '08: Proceedings of the 2008 ACM SIGGRAPH symposium on Video games*. 163--170. New York, NY, USA (2008)
8. Egenfeldt-Nielsen S.: *Beyond Edutainment: The Educational Potential of Computer Games*. Continuum Press (2007)
9. Guidi G., Frischer B., De Simone M., Cioci, A., Spinetti, A., Carosso, L., Micoli, L., Russo, M., Grasso, T.: Virtualizing ancient rome: 3d acquisition and modeling of a large plaster-of-paris model of imperial rome. In: *Proceedings SPIE International Society for Optical Engineering*, 5665, 119--133 (2005)
10. Goldberg D. W., Wilson J. P., Knoblock C. A.: From text to geographic coordinates: The current state of geocoding. *URISA Journal*. 19, 33—46 (2007)
11. Kolb, D. A.: *Experiential learning : experience as the source of learning and development*. Englewood Cliffs, N.J: Prentice-Hall. (1984)
12. Koolen M., Kazai G., Craswell N.: Wikipedia pages as entry points for book search. In *WSDM '09: Proceedings of the Second ACM International Conference on Web Search and Data Mining* 44—53. New York, NY, USA (2009)
13. Kleinermann, F., Mansouri, H., Troyer, O. D., Pellens B., Ibanez-Martinez, J.: Designing and using semantic virtual environment over the web. 53--58.

14. Krieger N.: Place, space, and health: GIS and epidemiology. *Epidemiology* 14, 4 380—385 (2003)
15. Kallmann M., Thalmann D.: Modeling behaviors of interactive objects for real-time virtual environments. *Journal of Visual Languages and Computing* 13, 2, 177—195 (2002),
16. Mrissa M., Dietze S., Thiran P., Ghedira C., Benslimane D., Maamar Z.: Context-based Semantic Mediation in Web Service Communities. Springer, Berlin, (2009)
17. Meli M.: Knowledge management: a new challenge for science museums. In: *Proceedings of Cultivate Interactive*, 9 (2003)
18. Marcos G., Eskudero H., Lamsfus C., Linaza M. T.: Data retrieval from a cultural knowledge database. In: *Workshop on Image Analysis for Multimedia Interactive Services Montreux, Switzerland*, (2005)
19. Mautone T., Spiker A., Karp M.: Using serious game technology to improve aircrew training. In: *The Interservice/Industry Training, Simulation and Education Conference (I/ITSEC)* (2008)
20. Pierce J. S., Pausch R.: Colorplate: Navigation with place representations and visible landmarks. *Virtual Reality Conference, IEEE*, 288 (2004)
21. Rushton, G., Armstrong M., Gittler J., Greene B., Pavlik C., West M., Zimmerman D.: Geocoding in cancer research - a review. *American Journal of Preventive Medicine* 30, 2 516—524 (2006)
22. Rebolledo-Mendez, G., Dunwell, I., Martinez-Miron, E. A., Vargas-Cerdan, M. D., de Freitas, S., Liarokapis, F., Garcia-Gaona, A. R.: Assessing neurosky's usability to detect attention levels in an assessment exercise. In: *Proceedings of the 13th International Conference on Human-Computer Interaction. Part I. Berlin, Heidelberg*, 149—158 (2009)
23. Sundstedt, V., Debattista K., Longhurst P., Chalmers A., Troscianko T.: Visual attention for efficient high-fidelity graphics. In: *SCCG '05: Proceedings of the 21st spring conference on Computer graphics*, ACM, New York, NY, USA, 169—175 (2005)
24. Spagnuolo, M., Falcidieno B.: 3d media and the semantic web. *IEEE Intelligent Systems* 24, 2 90—96 (2009)
25. Sanchez, S., Luga H., Duthen Y., Balet O.: Bringing autonomy to virtual characters. In *Fourth IEEE International Symposium and School on Advance Distributed Systems. Published in Lecture Notes in Computer Science*, 3061, Springer (2004)
26. Troyer, O. D., Kleinermann F., Mansouri H., Pellens B., Bille W., Fomenko V.: *Developing Semantic VR-Shops for E-Commerce*. Springer, London, (2006)
27. Van Dijk E., Op Den Akker H. J. A., Nijholt A., Zwiers J.: Navigation assistance in virtual worlds. *Informing Science, Special Series on Community Informatics*, 6, 115—125 (2003)
28. White, R. W., Dumais S. T., Teevan J.: Characterizing the influence of domain expertise on web search behavior. In: *WSDM '09: Proceedings of the Second ACM International Conference on Web Search and Data Mining*. New York, NY, USA, 132--141 (2009)
29. White, M., Mourkoussis N., Darcy J., Petridis P., Liarokapis, F., Lister, P., Walczak, K., Wojciechowsky, R., Cellary, W., Chmielewski J., Stawniak M., Wiza W., Patel M., Stevenson J., Manley J., Giorgini F., Sayd P., Gaspard F.: Arco: an architecture for digitization, management and presentation of virtual exhibitions. In: *CGI '04: Proceedings of the Computer Graphics International*. IEEE Computer Society, Washington, DC, USA 622—625 (2004)
30. Yee, H., Pattanaik, S., Greenberg D. P.: Spatiotemporal sensitivity and visual attention for efficient rendering of dynamic environments. *ACM Trans. Graph.* 20, 1, 39—65 (2001)