A Web-Based Medical Multimedia Visualisation Interface for Personal Health Records

Michael de Ridder, Liviu Constantinescu,
Lei Bi, Youn Hyun Jung, Ashnil Kumar,
Jinman Kim, David Dagan Feng
School of Information Technologies
The University of Sydney
NSW, Australia
mder9143@uni.sydney.edu.au
{liviu, lei, youn, ashnil}@it.usyd.edu.au
{jinman.kim, dagen.feng}@sydney.edu.au

Michael Fulham
Department of Molecular Imaging
Royal Prince Alfred Hospital
Sydney, Australia
Sydney Medical School
The University of Sydney,
NSW, Australia
michael.fulham@sydney.edu.au

Abstract

The healthcare industry has begun to utilise web-based systems and cloud computing infrastructure to develop an increasing array of online personal health record (PHR) systems. Although these systems provide the technical capacity to store and retrieve medical data in various multimedia formats, including images, videos, voice, and text, individual patient use remains limited by the lack of intuitive data representation and visualisation techniques. As such, further research is necessary to better visualise and present these records, in ways that make the complex medical data more intuitive. In this study, we present a web-based PHR visualisation system, called the 3D medical graphical avatar (MGA), which was designed to explore web-based delivery of a wide array of medical data types including multi-dimensional medical images; medical videos; text-based data; and spatial annotations. Mapping information was extracted from each of the data types and was used to embed spatial and textual annotations, such as regions of interest (ROIs) and time-based video annotations. Our MGA itself is built from clinical patient imaging studies, when available. We have taken advantage of the emerging web technologies of HTML5 and WebGL to make our application available to a wider base of users and devices. We analysed the performance of our proof-of-concept prototype system on mobile and desktop consumer devices. Our initial experiments indicate that our system can render the medical data in a fashion that enables interactive navigation of the MGA.

1. Introduction

Modern healthcare utilises numerous forms of multimedia data, including multi-dimensional medical images, patient and specimen photos, videos, sensor and text-based data, in a wide range of clinical applications [1-2]. When patients are provided with such data (typically in print, email or on an optical disc), most simply file it away without understanding or using the information [3]. To facilitate patient involvement and maintenance of their own medical data, as well as to improve understanding, many healthcare providers (with the support of government initiatives) now provide centralised electronic personal health record (PHR) systems [4-5]. Recent PHRs have been designed to serve as user-friendly, patient-facing digital repositories that consolidate an individual's medical history and provide tools for communication [6-7]. Further, PHRs are being designed to be accessible on tablet and smart-phone devices in the growing trend towards mobile healthcare.

To increase availability, compatibility and accessibility while potentially reducing costs associated with hospital infrastructure, there is an interest in moving medical data to web-based cloud computing platforms for storage and maintenance [8]. This interest has been extended to practitioner-facing electronic health record (EHR) systems and the picture archiving and communication system (PACS) [9]. Web-based applications for accessing these data have become the most effective way of building upon current cloud computing system capabilities. Such web-based solutions should support desktop and mobile clients, and must be able to supplement existing

repositories by offering improved portability and interoperability for the health data.

In this study, we present a novel web-based PHR visualisation system, called the medical graphical avatar (MGA), which was designed to explore web-based delivery of a wide array of medical data types including multi-dimensional medical images, e.g. positron emission tomography - computed tomography (PET-CT); magnetic resonance imaging (MRI) and X-ray; medical videos, e.g. endoscopic; text-based data and metadata, e.g. test results; and supplementary notes, e.g. textual / spatial annotations. Using the data processing techniques for medical data in PHR described in our previous study [10], we extracted 'mapping' information, such as focus body parts, and regions of interest (ROIs), for each medical event in the patient's history. Employing this information, we presented the medical data 'in-situ' using our avatar in the WebGL context.

3. Method

3.1. System Overview

Figure 1 shows the core interface for our MGA, which consists of three parts. Figure 1(a) is the Web-GL context which is used to present all spatial data, including highlighting, in red, the focus body parts in the avatar's anatomy for medical events; The top of Figure 1(b) is split into three tabs, the first displays textual information, the date and title, for events in the current year; the second presents detail of the selected event; and in the third, a hierarchy of the avatar's anatomy (detailed in Figure 2), which can be used to hide objects, e.g. skeleton and liver, in the Web-GL context. The bottom of Figure 1(b) is a navigation panel with the core rotation, pan and zoom for the Web-GL camera, with space for additional controls, for video playback and image stack traversal, which are added when focusing on timeline events containing image stacks or video. Figure 1(c) is the timeline, which is navigated through scrolling and dragging. It allows for an overview of the medical history at a glance, while also updating the textual data and spatial information presented in Figure 1(a, b) during navigation.

The architecture of the system is designed to support data flow between these three components. Medical data are divided into discrete events, and processed via the following steps: (i) textual details of the date, a description and the focus body part in the anatomy are parsed for the timeline and 3D location on the avatar; (ii) the raw data are filtered into a contextually appropriate viewer widget: an image

viewer, a video viewer, or the more general "miscellaneous" item viewer, which serves as a baseline in the interest of allowing our system to handle all data types; (iii) spatial regions of interest (ROIs), e.g. tumours, are extracted to be presented in the 3D scene.
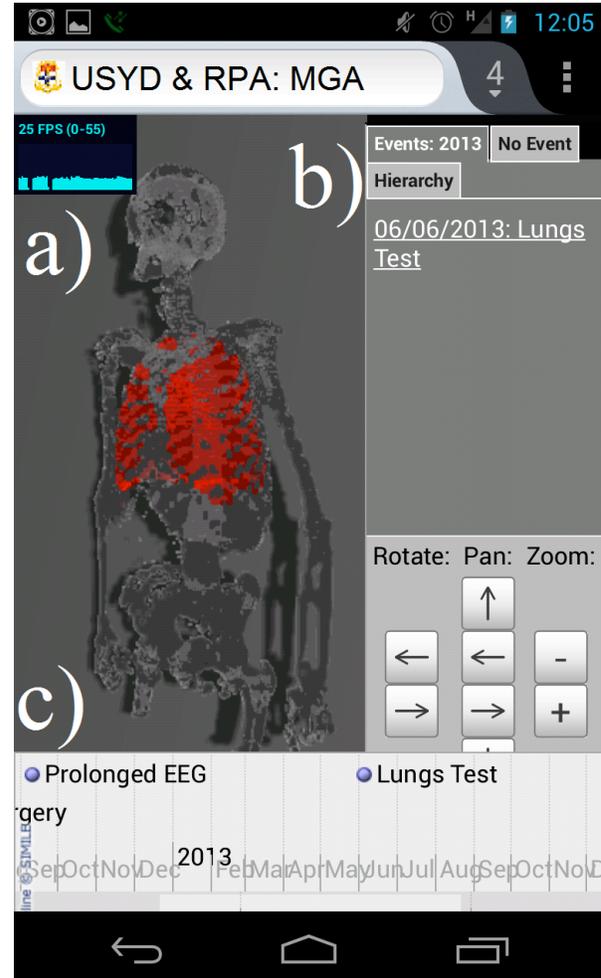

Figure 1. The main graphical interface of the MGA running in Firefox on a Google Nexus 4 smartphone.

3.2. Segmentation and Registration

Whole-body PET-CT image volumes were used to generate the 3D geometry of our sample MGA. These images were obtained from anonymised PET-CT scans from patient studies. Major organs and structures were segmented from both the PET and CT data. From the CT, the following segmentation steps were applied (automatically, unless noted otherwise): (i) multi-atlas registration with a weighted decision function was used to segment the liver and the descending aorta [11]; (ii)

lung structures were segmented by adaptive thresholding followed by estimating the two largest tissues [12]; and (iii) the skeleton was segmented by intensity thresholding based on Hounsfield units (HU) of the bones [13]. From PET, manually adjusted intensity thresholding was used to define a seed point from the brain, heart, bladder and tumours; these seeds were then used with grow cut algorithm [14-15] to segment the structures.

For patient data comprising of multiple image scans (e.g., to assess response to treatment), we used image registration [16] to morph all the images in the series to occupy the same spatial coordinates (morphed to the initial image). This allows our system to visualise changes in the patient data across multiple scans.
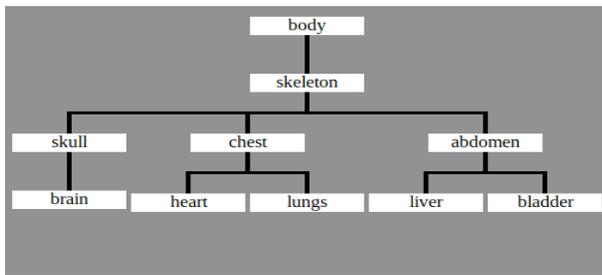

Figure 2. Anatomical hierarchy used by our MGA system

3.3. 3D Mesh Construction

In our avatar design, we partitioned the human body into a ten-node hierarchy, as shown in Figure 2. Here, each node represented a key segmented structure from our image data, patient PET-CT, or a region to group by, and was used in loading meshes as well as mapping medical events. For patient studies without an image data set, we used the Zubal Phantom [17], a simulation of the major anatomical regions of the human body, to provide these structures. Starting from the segmented volumetric images, we first translated these objects into polygonal surface meshes using the IntervalThreshold, WEMIsoSurface and WEMSave modules in MeVisLab [18]. For desktop use, high-quality polygon models were constructed, by passing the full resolution of the images through the WEMIsoSurface module. To improve the performance of our MGA on mobile devices, these 3D meshes were down-sampled by the WEMIsoSurface module to one third of the full voxel density. Next, meshes were exported as Wavefront Obj format and then converted to JSON using the WebGL-based Three.js Javascript graphics library [19].

The display method for the MGA traverses the hierarchy in a top-down manner, searching for a correctly named file at each node. This method uses the year of the event as a base in order to update the scene with relevant anatomical structures. Thus, our series of temporally sorted body meshes and ROIs ensures that the patient is always seeing the most contextually relevant and temporally-accurate results possible, allowing for the MGA to accurately model growth and change in the patient and their medical condition(s).

3.4. Data Representation

To present the data contained within medical events as a 3D scene, our MGA system makes intense use of the capabilities of WebGL, a JavaScript API that enables implementing interactive 3D graphics directly in the web-browser. Further, we utilised the Three.js JavaScript library to facilitate our use of WebGL. For the purpose of delineating anatomical structures, each of our segmented meshes is stored as a separate JSON object, which our system then renders in real-time. This reduces the need for any sort of segmentation in the browser, allowing us to simply align the pre-segmented pieces in the scene. In order to increase the user's ability to differentiate these meshes, and thus become aware of the important ROIs, our system also applies texturing, colouration and translucency to each object. This information is then fed to the GPU for processing via Three.js. Finally, our MGA is rendered directly into the web-browser, which can then be overlaid with further WebGL and HTML5-based visual elements. Examples of this rendering and the elements are shown in the four panels of Figure 3, wherein each frame shows the MGA scene at a different time point, and with different medical events selected.

Taking advantage of these capabilities, we developed a specific viewer component to support image volumes in the standard DICOM format. The image viewer is able to present a cross-sectional image stack 'in-situ' over the avatar, allowing correlations between the two to be easily discerned. The method used to align these visualisations involves locating the focus body part in the 3D anatomy using hierarchical traversal, and determining its bounding box. A 2D plane is rendered at that position and then textured with the image data. As a user navigates the slices of the image data, the plane is moved forward and backward in the bounding box and retextured, giving the appearance of a live cross-section. The appearance of such a cross-section is shown in Figure 3(a).
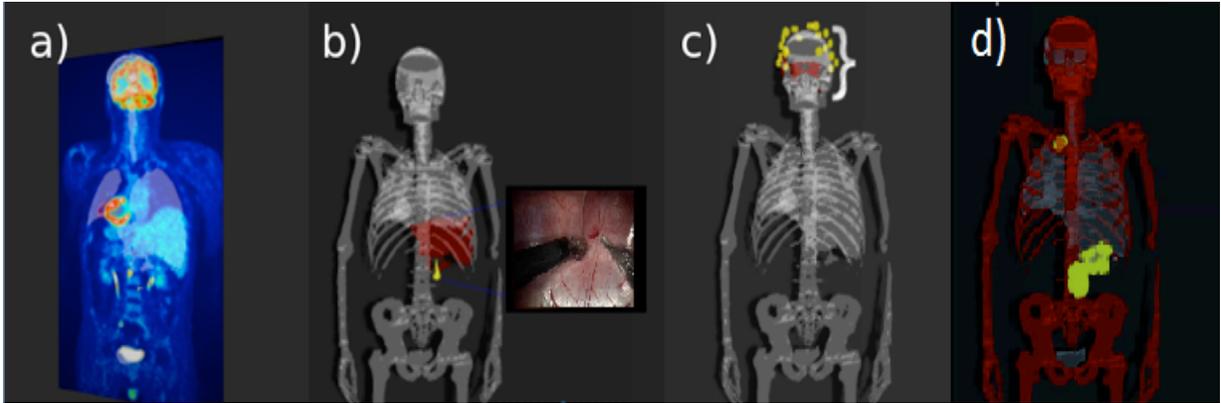
Figure 3. MGA Avatar with medical multimedia viewers. a) 3D image viewer with hidden skeleton; b) Video viewer with a yellow marker representing the anatomy related to an annotation; c) Miscellaneous data viewer, shows EEG electrodes (yellow); and d) no data viewer with ROIs (of tumours) highlighted in yellow and focus body parts rendered in red.

We further implemented a component to act as a video viewer, as shown in Figure 3(b). This similarly constructs the 2D plane, using hierarchy traversal to locate the focus, but instead of texturing with a single image, a HTML5 video element is embedded in the plane. In the 3D scene, the resulting video is situated to one side of the avatar, such that no part of the body is obscured by its content. This allows us to use the avatar itself to convey meaningful contextual information that could potentially aid user understanding. Yellow cones are rendered using the bounding box of the focus to create spatial annotations in 3D space on the avatar, corresponding to textual annotation, allowing the patient to contextualise the textual annotation information. To do this, each annotation defines an x and y coordinate ranging from (0,0) to (1,1). The cone is then rendered at the given coordinates on the front face of the bounding box.

With such a wide range of medical data used in clinical practice, it was not feasible to develop specific viewers for every possible medium. As such, a miscellaneous viewer component, Figure 3(c), was integrated to handle unfamiliar data types. It works by spatially indicating the anatomical focus of the event it is displaying, in addition to the relevant metadata. The notation used is a right-brace symbol, encapsulating the area in question. This allows for the visualisation to add additional spatial cues to aid understanding, regardless of the type of medical data fed into the system.

To display medical events in a way that allows for intuitive navigation and browsing, our MGA uses a timeline. We automated insertion of events into our timeline via textual parsing, using the SIMILE timeline toolkit [20]. This toolkit ensures that our timeline function is compatible with every major web-browser, and facilitates navigation via scrolling and dragging. Our system adapts the functionality of the SIMILE timeline for our purposes, however, integrating it with our hierarchy view and avatar visualisation such that when clicked, events on the timeline are represented using the appropriate data viewers and on the appropriate avatar structure.

4. Results

4.1. Performance

We tested the computational performance of our MGA by running the system in its default state and with each of the viewers present, on a wide range of standard consumer-level devices and in a number of the most common and popular HTML5 and WebGL compliant browsers. To normalise these results, tests were performed as the only active task on each system; access to the MGA web page on desktops was through localhost, while mobile devices were tested from within the same wireless network. Each test was performed 10 times and the average was taken. Table 1 summarises the observed results in frames per second (FPS). We compared these results with metrics specified by Nielsen [21], whereby a 'seamless' response occurs at greater than 100 FPS, while high performance ranges from 10 to 100 FPS. The two desktop/laptop machines with dedicated graphics cards (Asus and Macbook Air) perform well with $\geq$ 30 FPS; whereas the Dell desktop without the graphics card faired poorly in comparison. Among the smartphone / tablet evaluations, the Nexus 4 smartphone clearly performed best with a consistent 25 FPS, while other

Table 1. Performance of MGA system on a range of consumer-level desktop and mobile devices. Two popular web browsers supporting WebGL was used (Google Chrome and Mozilla Firefox (FF))

Scene Content	Asus M515N-c1		Macbook Air A1369		Dell Optiplex 760		Google Nexus 4	Google Nexus 7		Samsung Galaxy Tab 7.7	
	Chrome	FF	Chrome	FF	Chrome	FF	FF	Chrome	FF	Chrome	FF
Avatar Only	30.0	30.0	35.5	32.1	7.8	7.7	25.0	9.2	9.0	2.0	2.1
+ Misc Viewer	30.0	30.0	35.2	32.4	7.8	7.8	25.0	8.9	8.8	2.1	2.0
+ Image Viewer	30.0	30.0	35.8	32.0	7.7	7.3	25.0	8.9	8.8	1.6	1.7
+ Video Viewer	30.0	30.0	30.6	29.9	7.6	7.3	25.0	6.5	6.8	<1	<1

devices (Nexus 7 and Galaxy Tab) didn't reach 10 FPS. Note that the 3D meshes used in the rendering were down-sampled for the less powerful mobile devices. Also note that other popular iOS devices were not tested, as no WebGL compatible browsers were available on them at the time of testing.

4.2. Data Mapping and Navigation

To determine if our MGA correctly updates the 3D scene through timeline navigation, a representative set of clinical data was constructed from multiple anonymised clinical datasets, including PET-CT images (acquired with a Siemens Biograph TruePoint mCT scanner), surgical video (public data set available from [22]), textual (our own simulated data) and spatial data, derived from our images as explained in Section 3.2. For each of the events, we noted a set of rendering objects that should appear with its associated detailed viewer object. With these criteria defined, we manually selected each event and measured the MGA's ability to correctly render the scene. Through this we found that ROIs, such as tumours, appeared and disappeared correctly (within one second, on the desktop devices), the viewers were correctly placed in the scene and the avatar itself updated with navigation of the timeline.

4.3. Usability

To measure the usability of our system, we conducted a case scenario evaluation of our interface design against standard usability metrics outlined by Nielsen and Tory and Möller [21,23]. These contain a set of standards that should be followed when designing an intuitive and inclusive interface. Using the terminology from these investigations, we found that while our MGA did not provide navigation history, the interface of the system was domain specific, reused common tools, appropriately highlighted elements and made core navigation elements always available.

Therefore, our interface design met most of the main criteria, such that it may be able to facilitate understanding of the medical data.

5. Discussion and Future Work

In this study, we demonstrated the capability of our MGA to map many different types of medical data to a personalised avatar by extracting information about their temporal and spatial context. Our preliminary results demonstrate that the visualisation approach we have adopted can potentially be used to improve patient understanding of the medical data by extending the display, filtering and navigation capabilities of current PHR systems.

As of the time of our evaluation, there were several web browsers that did not support WebGL, e.g. Internet Explorer and Safari. However, since our evaluation, partial support has been introduced into the latter and we anticipate that all browsers will soon adopt the library. Further, while the rendering performance on two of the mobile devices was below the recommended level, improvement in results between Nexus4 (smartphone) and Nexus7 (tablet) that were released less than one year apart, suggests that future mobile technology will rapidly improve performance and hence be able to reach higher FPS performance

Moreover, due to the nature of web applications, online repositories, such as Microsoft Healthvault and ontologies, e.g. SNOMED-CT [24], and clinical practice guidelines, may be integrated in future. This would assist in providing a more thorough representation of a patient's medical history by presenting a greater depth of information when requested by the patient.

Security is another important issue, which our current prototype did not address. Nevertheless, as a web application, our system can easily adopt security standards such as secure sockets layer (SSL)

encryption. Thus, a major component of future research will be in ensuring data security and confidentiality through adopting these security standards.

Due to the patient-facing nature of our MGA system, a priority of our future work is to perform a patient-focused usability study. This can occur with further development of our system, as outlined in the previous points, so that we can explore whether the current interface aids understanding and also garner suggestions to improve the benefits of the system for patients and other end-users.

6. Conclusion

Our MGA offers a new method of preprocessing and displaying a large array of medical data on a personalised, image-driven avatar. We have shown that an advanced level of medical data visualisation can be achieved directly in a web-browser through the use of HTML5 and WebGL technologies. Using these powerful standards, our MGA can be accessed on a wide variety of devices and in any of the major operating systems. In addition, as WebGL and HTML5 represent the future of web applications, and already enjoy substantial adoption among mobile device vendors, it is likely that a substantial proportion of new and emerging mobile devices will support our MGA.

Our data indicate that our MGA is able to display spatial and temporal contextual information, available in all forms of medical data. Further, our system plots these data on a timeline to allow intuitive browsing and exploration of an individual's medical multimedia data. Our prototype system exhibited good performance on a range of consumer devices in visualising and navigating a wide range of clinical data types.

7. Acknowledgements

8. References

[1]   R. Brooks and C. Grotz, "Records: How Healthcare Providers Are," *J. Bus. & Econ. Research*, vol. 8, no. 6, pp. 73–84, 2010.

[2]   N. Archer, U. Fevrier-Thomas, C. Lokker, K. a McKibbon, and S. E. Straus, "Personal health records: a scoping review.," *JAMIA*, 18, no. 4, pp. 515–22, 2011.

[3]   J. S. Kahn, V. Aulakh, and A. Bosworth, "What it takes: characteristics of the ideal personal health record.," *Health affairs (Project Hope)*, vol. 28, no. 2, pp. 369–76, 2009.

[4]   Microsoft, "Microsoft Health Vault," 2012. [Online]. http://www.microsoft.com/en-us/healthvault/.

[5]   Australian Government: Department of Health and Ageing, "Personally Controlled Electronic Health Records," 2012. [Online]. http://www.health.gov.au.

[6]   M. John D. Halamka, MD, Kenneth D. Mandl, MD, MPH, Paul C. Tang, "Early Experiences with Personal Health Records," *JAMIA*, vol. 15, no. 1, pp. 1–7, 2008.

[7]   L. Constantinescu, J. Kim, and D. Feng, "Integration of interactive biomedical image data visualisation to Internet-based Personal Health Records for handheld devices," *Healthcom*, pp. 79–83, Dec. 2009.

[8]   M. L. Braunstein, *Health Informatics in the Cloud.* Springer, 2013.

[9]   O. Ratib, "From PACS to Clouds," *Eur. J. Radiol.*, vol. 78, no. 2, pp. 161–162, 2011.

[10]   M. de Ridder, L. Bi, L. Constantinescu, J. Kim, and D. Feng, "Data Processing and Presentation for a Personalised, Image-driven Medical Graphical Avatar," *EMBC*, 2013. (accepted)

[11]   L. Bi, J. Kim, L. Wen, and D. D. Feng, "Automated and Robust PERCIST-based Thresholding framework for whole body PET-CT studies."

[12]   S. Hu, "Automatic lung segmentation for accurate quantitation of volumetric X-ray CT images.," *IEEE Trans. Med. Imag.*, vol. 20, pp. 490–498, 2001.

[13]   R. Manniesing, B. K. Velthuis, M. S. van Leeuwen, I. C. van der Schaaf, P. J. van Laar, and W. J. Niessen, "Level set based cerebral vasculature segmentation and diameter quantification in CT angiography.," *Med. Imag. Anal.*, vol. 10, no. 2, pp. 200–14, Apr. 2006.

[14]   J. Kim, L. Bi, L. Wen, and D. Feng, "Cellular Automata Diffusion Filter based Interactive Tumour Segmentation for Positron Emission Tomography," in *EMBC*, 2013. (in press)

[15]   V. Vehznevets, et al., "GrowCut:Interactive multi-label ND image segmentation by cellular automata," *Proc. of Graphicon*, 2005.

[16]   G. Nyirenda, J. Kim, L. Wen, and D. Feng, "Automated segmentation of tumour changes in temporal PET-CT data," *ISBI*, no. 9, pp. 1699–1702, 2012.

[17]   G. Zubal, "The Zubal Phantom." [Online]. http://noodle.med.yale.edu/zubal/.

[18]   MeVis Medical Solutions AG, "MeVisLab." 2012.

[19]   J. Etienne, "Three.js," 2012. [Online]. https://github.com/mrdoob/three.js.

[20]   SIMILE Widgets, "SIMILE Timeline," 2009. [Online]. http://www.simile-widgets.org/timeline/.

[21]   J. Nielsen, *Usability Engineering*, 2nd ed. Morgan Kaufmann, 1994.

[22]   "Hamlyn Centre Laparoscopic/Endoscopic Video Datasets." [Online]. http://hamlyn.doc.ic.ac.uk/vision.

[23]   M. Tory and T. Möller, "Human factors in visualization research.," *IEEE Trans. Vis. Comput. Graphics*, vol. 10, no. 1, pp. 72–84, 2004.

[24]   International Health Terminology Standards Development Organisation, "SNOMED CT." 2012.