

Semantic Annotation of Ubiquitous Learning Environments

Mark J. Weal, Danius T. Michaelides, Kevin Page,
David C. De Roure, *Fellow, IEEE*, Eloise Monger, and Mary Gobbi

Abstract—Skills-based learning environments are used to promote the acquisition of practical skills as well as decision making, communication, and problem solving. It is important to provide feedback to the students from these sessions and observations of their actions may inform the assessment process and help researchers to better understand the learning process. Through a series of prototype demonstrators, we have investigated the use of semantic annotation in the recording and subsequent understanding of such simulation environments. Our Semantic Web approach is outlined and conclusions drawn as to the suitability of different annotation methods and their combination with ubiquitous computing techniques to provide novel mechanisms for both student feedback and increased understanding of the learning environment.

Index Terms—Semantic annotation, Semantic Web, ubiquitous computing, case study.

1 INTRODUCTION

THIS paper presents a series of prototype demonstrators that have looked to evaluate the use of semantic annotation as part of a skills-based learning environment to better understand *how* students learn.

Simulations are used to promote the acquisition of practical skills as well as decision making, team working, communication, and problem solving [1]. They can be incorporated into assessment of student performance [2], which brings a requirement that the approaches for assessment and feedback need to be sound, valid, reliable, feasible, educational, and of course acceptable to practitioners [3]. Through the simulation, the student experiences are designed to be exactly as they would experience in the workplace in real time. The University of Southampton has such a clinical skills laboratory.

The laboratory mimics the reality of ward life in both its behaviors and resources, equipment, clinical charts, wall displays, and phones (see Fig. 1). The ward is equipped with computerized and interactive SimMan¹ mannequins, noncomputerized mannequins, and a range of equipment that are purposively arranged to provide clinical activities for the students. SimMan can be programmed to mimic varying medical conditions as the scenario progresses [4]; an inbuilt loudspeaker allowing a remote operator to provide the patient with a voice.

The students receive a report that typifies genuine practice and engage in the scenarios designed for the occasion. The students are given a plethora of tasks and the computerized mannequins can be programmed to alter their parameters to a point of significant deterioration when emergency responses would be required. These activities provoke the students to move themselves and equipment around the ward, to interact with each other and the supervising staff members, and to use the telephone. This also means that not only is concurrent activity taking place in different parts of the ward but there is plenty of background noise and movement.

The ward is viewable from a central control room via six ceiling mounted cameras, each controllable with a 360 degree viewing angle and microphones suspended above the beds to record audio. The cameras are remotely controlled from an adjacent room, where teaching staff can monitor the students through the audio/video streams and direct proceedings without interrupting the ongoing simulation. When the students and mentors are “immersed” in the simulation and behaving “as in real practice,” the use of captured video data can provide important information about their performance. This type of activity is an integral part of the curriculum; skills-based learning is being developed as part of a national agenda to help ensure that practitioners are “fit for practice” [5].

The School of Health Sciences team has an ethical and governance framework to address the ethical, legal, and governance issues that arise through the collection of data concerning students, staff, patient related data, and the role of the researchers. In this case, the key issues were 1) the access of patient related material to the research team; 2) the collection, storage, and dissemination of staff and student data that could not be anonymized. The students and staff members willingly participated and contributed to the debriefing session of the trials with suggestions for improvement and an outline of their experiences. Previous findings have shown that being filmed as part of an assessment activity of this type does not significantly

1. <http://www.laerdal.com/SimMan>.

- M.J. Weal, D.T. Michaelides, and K. Page are with the Electronics and Computer Science Department, University of Southampton, Southampton, Highfield SO17 1BJ. E-mail: {mjw, dtm, krp}@ecs.soton.ac.uk.
- D.C. De Roure is with the Oxford e-Research Centre, University of Oxford, 7 Keble Road, Oxford OX1 3QG. E-mail: david.deroure@oerc.ox.ac.uk.
- E. Monger and M. Gobbi are with the Faculty of Health Sciences, Building 67, Highfield, Southampton SO17 1BJ. E-mail: {e.monger, m.o.gobbi}@soton.ac.uk.

Manuscript received 1 Apr. 2011; revised 24 Nov. 2011; accepted 6 Dec. 2011; published online 12 Dec. 2011.

For information on obtaining reprints of this article, please send e-mail to: lt@computer.org, and reference IEEECS Log Number TLTSI-2011-04-0038. Digital Object Identifier no. 10.1109/TLT.2011.37.



Fig. 1. The clinical skills laboratory.

modify student behavior any more than having assessors physically present in the room [6].

Audio and video together present a highly detailed capture of an activity; perhaps too detailed, because reviewing a recording can be as time consuming as the original activity. Similarly, annotating a video by hand can be an intensive and laborious process and often involves reviewing the entire digital record. One approach is to make annotations “live,” during the teaching session, although this is unlikely to be a comprehensive record of events and precludes the full engagement in the activity itself. Ubiquitous computing technologies and techniques provide us with an additional mechanism to capture annotations on events that take place in the clinical skills laboratory and from sources that have a low impact and overhead on the participants.

Annotations are at their simplest, just metadata, but by harvesting annotations with meaning, defined by ontologies—semantic annotations—we aim to fuse metadata sources together. These various linked data sets could include: annotations made by educators observing the unfolding scenarios; linked data automatically produced from the SimMan system logs as it records its progress through its programmed sequence and any sensed interactions with it; annotations made by observing students; and location information recorded from a tracking system. By combining both automatic annotation gathering with manual annotation techniques, we aim to provide a much richer data sets to help shed new light on how and why students are learning. Through a series of demonstrators, we will show how the combination of location-based semantic information with manually authored semantic annotations can begin to provide answers to questions such as *What did the student do?*, allowing the explicit connection of person and activity into a machine readable form. This could lead to improved assessment of student learning, facilities for student self-reflection, and further research into understanding student learning in skills-based environments.

2 RELATED WORK

There have been a number of projects that have looked at the application of semantic annotation of pervasive spaces. Some have been within the educational context but others

share common goals of connecting activity/task to time, place, and person.

The museum experience described in Hatala et al. [7] is a good example of the use of semantic descriptions in a real (and real-time) application. It uses inference rules alongside user models and content descriptions, and involves several ontologies. The “Semantic Smart Laboratory” work [8] uses RDF from the very first stage of capturing the activities of chemists working in a laboratory, as well as a sensor network to capture laboratory environmental conditions. This is used to establish a complete provenance trail through to scholarly output, enabling researchers to chase back to the original data.

The Task Computing project at Fujitsu Labs America [9] applies Semantic Web technologies (RDF, OWL, DAML-S) and Web Services (SOAP, WSDL) to pervasive computing, aiming to “fill the gaps between tasks and services.” Users see the tasks that are possible in their current context and are assisted in creating complex tasks from simpler tasks, which can then be reused.

The agents research community has also applied ontologies in the pervasive area, such as the FIPA Device Ontology specification (see <http://www.fipa.org>) which enables agents to pass profiles of devices. The Standard Ontology for Ubiquitous and Pervasive Applications (SOU-PA) [10] is a comprehensive example of an ontology. It includes vocabularies to represent intelligent agents with associated beliefs, desires, and intentions, time, space, actions and events, user profiles, actions, and policies for security and privacy. In our context, one can envisage agents which work with the accumulating knowledge—perhaps automating elements of assessment or flagging errors as they occur.

Many spatial annotation efforts are emerging. For example, accumulation of annotations in a spatial region is the basis of the OpenGuides “WIKI” city guides (<http://openguides.org>). The Basic Geo vocabulary is used in Locative Packets for spatial annotation (locative.net). The Open Geospatial Consortium (OGC) pursues standards for geospatial and location-based services (<http://www.opengeospatial.org>). Although location is at a much coarser granularity than in our work, some of the underlying principles are transferable.

Hypermedia links are another way of expressing associations between things, and the hypermedia research community has a long history of working with these associations as “first class citizens,” as they will be in our ontology. This was originally achieved using, for example, XML and XLink technology, but now increasingly uses RDF. Recent work on digital-physical linking illustrates the extension of these ideas into the physical world [11].

Work on authoring and design for ubiquitous systems tends to have concentrated on the system designer, with the assumption that they would also be deploying and maintaining the system. Work, such as the iStuff framework, provide an interface for connecting and orchestrating devices in a ubiquitous system [12]. The Urban Tapestries project looked at the idea of public authoring [13], where members of the public could create locations in an ubiquitous system by uploading GPS co-ordinates,

and then attach media items such as notes or photos, either in-situ using a PDA or at a later point on a web site. In this way, Urban Tapestries hands some of the design to the users, and allowing nontechnical people to create a ubiquitous experience. Similarly, M-studio allows users to create an experience by authoring content [14], delivering video content to PDAs according to their location. A graphical authoring tool allows authors to place video content at locations, but also supports storyboarding and simulated located playback, so authors can check the effect of movement on their narratives.

Topiary is a rapid prototyping system that uses a high level of abstraction (people and places, rather than sensors and devices). Topiary allows authors to storyboard, situate, and simulate information placed into a geographic environment on a map [15]. Topiary also supports automatic pathfinding and more advanced trigger conditions based on user and place (such as user1 and/or user2 are near, etc.). Such spatial triggers are echoed in the spatial inference we will discuss later. The eDiary [16] allowed architecture students to record their path during a site visit using a handheld device which would map photos and notes to a map of the site. Later on this could be edited on a PC and the path calibrated to the map. Nodes of the path, representing locations where notes had been taken, could be moved or expanded. The annotated map then was used in multimedia presentations on the site visit. This two-phase approach of bookmark annotations with subsequent refinement is something that has emerged from our findings.

There are numerous video annotation systems, two of which we would note. The DIVER system [17] allows users to attach textual annotations segments of video with a view to fostering collaboration. The LORAMS framework [18] notes the time consuming nature of such manual textual annotation and seeks to mitigate against this through the use of simple RFID markers with which users can perform searches on an annotated video set. Other work has focussed on automatically identifying events within video streams [19].

Initial uses of skills-based learning environments in nursing education have traditionally been very task specific, for example managing a cardiac arrest [20], or performing a specific intervention, for example giving an injection [21]. These simulations are generally short and easily objectively marked according to a defined set of criteria, "You do this, in this order." It is also of note, however, that many educational institutions have invested in facilities in simulated environments and use video for a variety of educational purposes. For example, the analysis and assessment of student performance and or competence, the analysis of events [22] or processes [23], and Objective Structured Video Examination [24].

3 SEMANTIC ANNOTATION

The Semantic Web is designed to express meaning. Originally designed to "bring structure to the meaningful content of Web pages. Its unifying logical language will enable these concepts to be progressively linked into a universal Web" [25].

Two key technologies underpinning the Semantic Web are: Extensible Markup Language (XML) and the Resource Description Framework (RDF). Fundamentally, RDF data describe "things," even if they cannot be directly retrieved on the Web (they just need to be identified.) RDF was created as a framework for metadata to provide interoperability across applications that exchange machine-understandable information on the Web. It has a very simple relational model which accommodates structured and semistructured data, and in fact can be seen as a universal format for data on the Web, providing greater interoperability and reuse than XML alone. Although designed to express the meaning of Web pages, the technologies are well suited to describing data of other forms and the notion of what constitutes the Semantic Web has evolved [26].

Part of the added value of the Semantic Web approach is the "network effect" that can be achieved by having metadata accumulate about the same things—those things then effectively interlink different pieces of knowledge, forming rich structures. For example, information about relationships between people (friend-of-a-friend, coauthorship, etc.) accumulates on those people to describe communities of practice. Similar effects are achieved when the metadata describes regions in time or in space, and there are RDF vocabularies (such as the Basic geo vocabulary, <http://www.w3.org/2003/01/geo>) for spatially located things.

One of the important roles for RDF in pervasive computing, together with the associated Web Ontology Language (OWL) which is used to describe shared vocabularies, is in describing context [27]. A variety of notions of context may be expressed, including location and user tasks [28]. Ontologies can also be used to describe device capabilities, for example to facilitate content delivery to devices with diverse characteristics [29].

So what do we mean by semantic annotation? This work builds on earlier work which looked at adding hyperstructure to video collaborations [30], [31]. Mechanisms for capturing annotations from the skills-based sessions have been developed by combining Semantic Web technologies and techniques previously applied to enhanced field trips for children [32]. These annotations can be attached to people, as in FOAF (friend of a friend networks) [33], physical objects [11], or tasks [9].

Annotations describing an activity space can then be used to generate an index into the video structure through which the detail-rich record can be more effectively used.

The uptake of Semantic Web technologies in education has been slow, with the main uses being in the creation of well-formed metadata for repositories [34]. Web 2.0 systems have also enabled lightweight knowledge modeling approaches (typically folksonomies) based around techniques such as community tagging, clustering, and community authoring [35]. The coming together of Web 2.0 technologies and semantic technologies are proposed as an inevitable development of existing technologies [36].

More simplistic keyword tagging approaches could be used such as those employed by Flickr (<http://www.flickr.com>) and del.icio.us (<http://del.icio.us>), however, these

benefit from scale and we felt that lightweight ontologies may give us the following advantages:

- Well-formed metadata providing consistency in the data.
- This makes for easier comparison across data sets as equivalence can be established more easily.
- The formal description allows for the relationships between concepts to be mapped out more easily than with a looser, keyword-based tagging system.
- Lightweight annotations made real time can then become more complex afterward as more detail is added or they are combined.
- Data can be exported and shared with a guaranteed shared vocabulary. Interoperability is one of the cornerstones of the Semantic Web and allows researchers to more easily share data and provide machine readable versions for software agents.
- The Semantic Web-based annotations provide an underlying data sets that allow for rule-based analysis or complex inferences (such as those supported by the JENA framework [37]).

4 SKILLS-BASED LEARNING ENVIRONMENTS

The ability to track and annotate people, equipment, and actions through simulated hospital ward activity can have many benefits. Simulated ward environments enable the safe and ethical development of tools that could subsequently be deployed for use in the health care practice, for example, designing location trackers that are health and safety compliant within a clinical or perhaps home setting. In their discussion of the suitability of industrial methods to improve the quality and efficiency of health services, Young et al. [38] recommended the use of simulations to identify how these methods would translate before the benefits could be realized. The problems of evaluating simulations themselves have been highlighted by Brailsford [39] and the digital record collected using our approach may provide new mechanisms to evaluate the efficacy of the many simulation models currently in use.

One important component of simulation activities is the feedback to students. In the emerging literature about simulation in healthcare, and other fields, there are mentions of feedback, and attempts to address the logistical issues. For example, Roberts et al. [40] describe using Discourse Analysis to elicit elements of good and poor communication in medical students.

Such audio-visual equipped simulation environments are used increasingly for the education and training of a range of students and also staff. The data generated can provide a rich resource for educational and research analysis of student performance as well as their interactions with each other and objects/equipment located in the environment (as shown in Fig. 2). Noting how people and equipment are physically located, move and interact with each other in response to events provides educational and management insights into the efficiency of these movements, logistics, ergonomics, environment design, team working, and leadership styles. Within the cost constraints of modern health services, strategies to improve design,



Fig. 2. Video capture of a training session.

process, and human performance are ever present. Other approaches have sought to make similar illuminations on how coordination is achieved with nondigital artifacts in clinical settings [41].

The ability to assess student performance and provide timely feedback is a huge challenge when the ward facilities have several cameras in simultaneous use and there are large numbers of students requiring such feedback. Indeed “finding” the student in the ward environment is a crucial step when the student may be visually “off camera,” yet captured by a different camera/microphone and the location device. Our approach to alleviating the problems associated with just using video data is to augment this with manual and automatically authored annotations. These annotations enable the marking of the events. Favela et al., [42] have shown how work activities can be derived from contextual information when it is available. The clinical incidents that form the basis of the simulated learning activity can generate time markers that should stimulate student responses. For example, the computerized patient (SimMan) can, at a predetermined moment, exhibit altered sensory data (e.g., pulse, blood pressure, oxygen levels). In this instance, we are arguing that the development of robust tools to monitor, track, annotate, and analyze data from people and equipment, provides the test bed for simulation scenarios that are realistic rather than speculative or hypothetical and have the potential to provide new methods to trial innovation and improvement processes. Our desire to match or triangulate data from the different data streams tests our ability to handle and translate these data into meaningful evidence for the educator, manager, or researcher until we can create technical, ethical, and practice solutions to these challenges.

Examining the process of these simulated scenarios, the following activities can be identified:

- Active participation by the students in a session. This will include the live observation of the activity by mentors performing roles within the scenario as well as lecturers observing the session from the control room.

- Peer observation of other students. While one group of students is engaged in a scenario, a second group is observing.
- Debriefing sessions. Immediately after the activity a debrief session is held, facilitated by the lecturers who participated in the scenario.
- Self-reflection after the event. The students will reflect on the activity and their own performance at a later time.
- Educator reflection on activities. The mentors and control room observers may wish to reflect on the activity in order to assess individual students performance or refine the scenarios.
- Educator reflection across activities. The educators may wish to consider a series of scenarios in terms of their efficacy as a learning tool, or, to address more specific research questions, for example examining how hygiene and infection control approaches are being used by the student cohorts as a whole.

For each of the activities described above, we can see how the use of captured video may be utilized where appropriate as well as examining what questions the students or educators might wish to be asking.

4.1 Active Participation in Sessions

What are the students doing? Where in the scenario are we? The activities of the students will be dependent on where they are in the scenario, the other related contextual factors like the current state of the patient and equipment functioning; and the parallel activities of other participants in the scenario. Being able to identify the precise position within the sequence of the scenario and how this relates to concurrent parallel actions may be an important part of the monitoring process. Often, cues to the current scenario position may not be inferred from the video alone (the patient has stopped breathing), and with multiple participants performing actions simultaneously, not necessarily visible from a single camera, being able to indicate current actions of all the students is likely to be important.

4.2 Peer Observation of Other Students

What are they doing? What should they be doing? The act of observing may in part involve the observers identifying what individual students within the scenario are doing at any given moment in time. Identifying appropriate and inappropriate behaviors may well offer indicators as to their understanding of actions required. Sometimes, however, it may be useful to indicate what should be happening during the current phase in the scenario (to use the previous example, the patient has stopped breathing and action is required.)

4.3 Debriefing Sessions

This is what you did? What were the interesting moments to discuss? Debriefing is an important component of these skills-based learning sessions providing the students with timely feedback on their performance and promoting reflection in their learning. During the debrief session, the mentor will wish to highlight key moments in the scenario in order to indicate both good practice and potential areas for improvement. Although the video provides a mechanism to

replay these key moments, identifying them is often problematic as the mentor is required to either estimate when the key moment took place either directly, or by narrowing in by moving forward or backward from remembered positions. Structuring of the activities in the timed settings of the scenario programmed into the mannequins is lost during the simple video capture process. This is more problematic when the mannequin responses are altered remotely in response to student behaviors.

4.4 Self-Reflection After the Event

What did I do? What should I have done? What did the other students do? How does my performance compare to others? How does my performance compare to previous sessions I did? During the activity the students are focussed on the task and these are designed to be both engaging and often stressful. The debrief sessions can provide feedback to the students as to how they did but will inevitably not cover the totality of any individual student's contribution. Access to the video allows the students to review the activity at a later date but may suffer from a number of problems. The multiple camera setup means that a student may move out of shot of one camera and into shot on a second while moving around the ward environment. Although all the camera feeds can be recorded and made available to the students, knowing when to switch cameras and to which other camera is not straightforward. Although able to visually observe what they did, students will not always be able to identify errors or inappropriate actions when they occur that may be caught by other observers. There will also not be an indication of missed alternative courses of action that might have been more appropriate. Although there may be a large corpus of videos it will not be possible for a student to compare how they dealt with a situation compared to their peers as identifying similar situations across the video data set would be a nontrivial manual task. Similarly, a student may wish to compare their performance with one of their own from a session carried out the previous year. Again, this type of benchmarking is difficult through the use of video alone.

4.5 Educator Reflection on Activities

What did the student do? What feedback should they have? What assessment can I make from their performance? What is the educator noticing? In addition to the students understanding what they did it will also be necessary for the educators to identify the actions of individual (and groups) of students. This will be required for immediate debriefing as well as identifying appropriate feedback to give to students. This could take the form of assessment "have they followed appropriate infection control procedures," or more long-term assessment of abilities, "have they demonstrated improvement from the similar activity recorded in the previous year." Assessment directly from the video may be straightforward with individual events, "did they put the oxygen mask on the patient correctly," but other types of activity occurring over a period of time may be harder to identify, "Now they are touching this patient, have they washed their hands since they touched the last patient?"

4.6 Educator Reflection across Activities

How did the students operate as a group? What common patterns of error can I identify? What have the observers of the

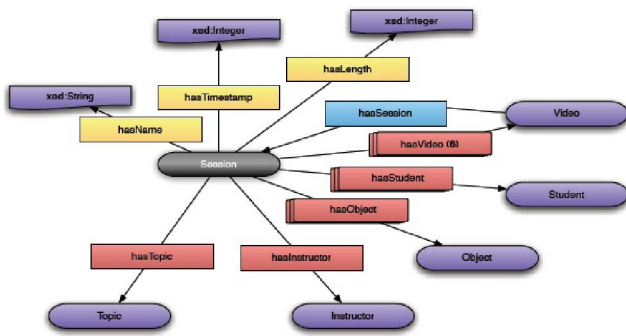


Fig. 3. The system ontology session entity.

session noticed or not? As well as identifying the actions of individual students, educators may wish to examine the actions of groups of students across a number of sessions. This could be to identify certain types of behavior, group formation, learning styles, or it could be to try and identify areas that required reinforcing, “are there large numbers of students that aren’t remembering to adjust the bed heights to an appropriate level before treating patients?” Some researchers may just be interested in certain types of activities, infection control, ergonomics, and being able to identify events and actions associated with these for research purposes is nontrivial using just the video data. Finally, it may be that the subject of the research is the educators themselves and identifying what it is that they are noticing when they observe students carrying out these activities.

Our approach to attempt to provide answers to some of these questions is through the use of manually authored and automatically generated semantic annotations.

5 THE DEVELOPMENT OF AN ONTOLOGY

A number of ontologies underpin these demonstrators. A system ontology was constructed that contains all the entities describing the videos, sessions, and participants. Fig. 3 shows the session entity and how this links the various videos of the session, the students, instructors, and objects of interest.

Fig. 4 shows the annotation entity in the ontology and how it connects annotations to the video(s) that it annotates, the author of the annotation and the session in which the annotation occurs.

For the creation of specific annotations, a domain ontology was developed to contain the domain specific annotation information. This allowed the underlying video annotation framework to be independent of the specific context of annotation. The nursing domain ontology was developed through a series of workshops, observational sessions, and discussion groups. Having identified types of annotation, we have constructed an ontology representing the range of annotations applicable in the scenarios. The ontology provides the basis for the annotation interfaces developed.

For annotation to be successful, it is important to design cues/prompts that are easily recognizable and familiar to the users. Two ways of achieving this are through naturalistic time sequenced observation or through the

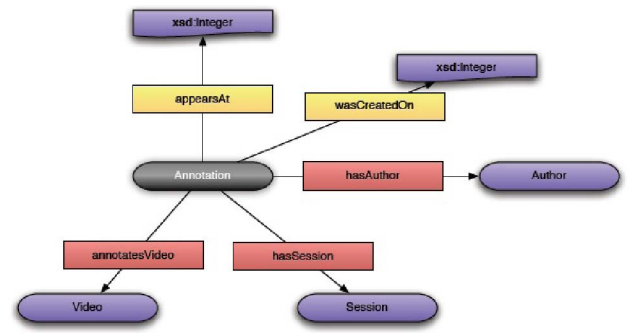


Fig. 4. The annotation entity.

use of established observational schedules. In our case, we have used naturalistic time sequenced observation. These have been clustered into themes according to discipline specific relationships. For example, “taking a pulse” appears under a heading of “taking and recording vital signs.” The individual activity of the pulse can then be broken down into further components such as “looking at watch,” “feeling pulse,” etc. The ontology was not intended to be in any way comprehensive nor to encompass all pervasive activities as has been attempted with other taxonomies [43]. These annotations, although possibly using medical terminology, are more naturalistic observations, and the ontologies developed are not intended as a mechanism for sharing clinical knowledge as is supported by other systems [44], [45].

The ontology was modeled using the Protégé ontology editor, with a base ontology describing the structure of these annotations coupled with domain specific instances of these annotations along with mechanisms for timestamping. The ontology allows for the construction of annotations about objects and events and the relationships between them. The notion of an EventWeb as opposed to a document Web has been proposed by Jain [46]. The ontology was designed to be easily extensible with the ability to add annotation describing specific research areas at a later date. Records of individuals are not kept within the annotations for issues of governance and security. An additional location ontology developed was used to describe the ward space. Fig. 5 shows an XML fragment of the domain ontology.

```
<Annotation name="hands/hygiene">
  <Annotation name="wash"
    type="Specific-Actions" />
  <Annotation name="touch hand to hand"
    type="Specific-Actions" />
  <Annotation name="touch soap"
    type="Specific-Actions"
    equipment_used="soap" />
  <Annotation name="touch self"
    type="Specific-Actions" />
  <Annotation name="touch something else"
    type="Specific-Actions"
    uses_equipment="true" />
  <Annotation name="gloves"
    type="Specific-Actions"
    equipment_used="gloves" />
</Annotation>
```

Fig. 5. A sample of the domain ontology (in XML).

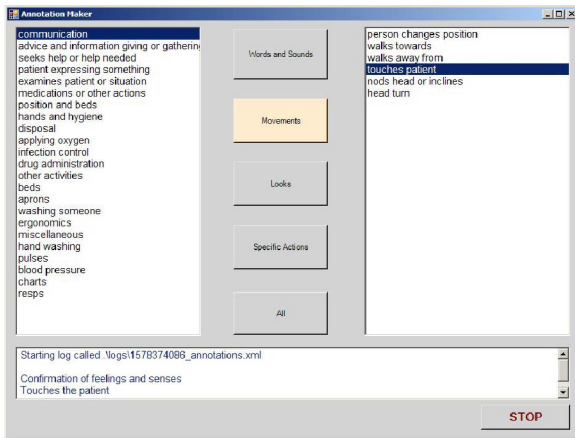


Fig. 6. Tool for adding textual annotations.

6 THREE CASE STUDIES

In order to explore the creation and use of annotations built upon the ontology, three small scale trials were carried out, the creation of manual annotations by observers real time, the use of audio annotations by observers of the videos of sessions, and the combination of annotations created automatically through location tracking with those recorded manually by observers. These trials were constructed around an existing learning session involving a group of students being tasked with monitoring the health of an admitted patient, whose condition will deteriorate over the course of the scenario ultimately resulting in pulmonary arrest. The scenario includes routine tasks (observations), communicating with members of a team and responding in an emergency situation. The feasibility of capturing on the fly annotations was evaluated along with the utility of annotations captured both through manual annotation and automatically through location sensing. The objectives of the case studies were to assess the usability of a manual annotation system in generating annotations in real time, better understand the ontologies required in the capturing of these annotations, and to investigate the utility of annotations captured in this way for various reuse possibilities. The methods employed were system logging, coupled with video recordings of the scenario. Participant observation was carried out of the annotation sessions, participants engaged in think aloud processes during the annotation and interviewing of participants was carried out postsession.

6.1 Manual Textual Annotation

The first approach taken was real-time manual annotation of the activities. By real time, we are referring to an observer monitoring the session via a video feed in the control room and recording annotations of what they observe through an annotation tool. These annotations can then be used for debriefing the students, providing feedback to the students during self-reflection at a later time, or for analysis of the activities by researchers interested in student learning.

Video annotation systems for other domains exist, for example, news production [47] where the focus is on more explicit description of content or for collaborative annotation of video [48]. In the case of our textual annotations, the authoring process can occur both in real time or postsession,

```
<annotationfile date="2007-06-20 12:46:09Z"
  timestamp="12707973">
  <videofile name="Sessions_20-6-07.avi"
    startindex="3804"/>
  <annotation id="annotation_12717707"
    type="look" timestamp="10">
    <name>at paperwork charts</name>
    <description>Looking at paperwork charts
    </description>
  </annotation>
  <annotation id="annotation_8723513"
    type="words" timestamp="407"
    communication="true"
    involves_others="true">
    <name>solicited information</name>
    <description>Solicited information
    </description>
  </annotation>
```

Fig. 7. A fragment of annotation file (in XML).

with the annotations potentially reused in a variety of ways. An example of this would be coarse annotations made during the exercise being used as an index for creating more detailed annotations about specific activities and events at a later point.

A simple interface was built so that an observer in the control room can, while monitoring a session, quickly capture events as they occur using the ontology; an event is time stamped and recorded when selected using a mouse in the tool (see Fig. 6).

Although the observational tool has been designed to be used simply and quickly, thus distracting the annotator from the video feed for a minimal period, the cognitive overhead of annotating in real time is still significant and volume and detail of annotation correspondingly limited. We will later turn to other, automated, pervasive sources of semantic annotations. By extending our ontology and mapping to others, we will be able to combine these annotations and produce structures for navigation and review.

Fig. 7 contains a snippet of a log file (in XML) generated as part of the real-time annotation session. This first prototype was using an XML version of the domain ontology as the system ontology was still in development.

The annotation tool was used by two lecturers observing four separate sessions of the same scenario. One hundred sixty-two annotations were made during the 30 minute sessions on average. The majority of the annotations were made at the top level (65 percent). In some cases, these were serving as placeholders for gaps in the ontology. As one participant commented "Some annotations commands were missing, although they were covered by the central command stems. I therefore used approximations." In some cases, more general annotations (top-level categories) were used as placeholders with the intention of adding more details through further annotation offline later on.

Generally, it was felt that the coverage of the ontology was adequate although only a subset was used for this specific scenario. It was felt that "the more annotations there are, the harder it is to remember where things are at the beginning," so the ability to scope the visible annotations for specific scenario may well improve this. Annotation was not limited to what was visible in the video, however, one

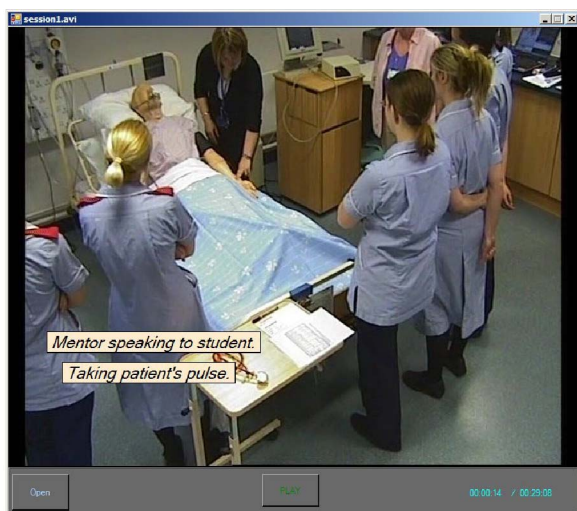


Fig. 8. Replaying the annotations.

participant articulating that “because there was more movement in and off screen, it was also possible to annotate sound that was off screen (for example, hand washing from the sinks).” Although as the participant went on to say, “the annotation of course can’t identify who was doing something, rather it annotates what is happening, unless the who has been entered into the annotation vocabulary.” All the participants felt they were able to broadly keep pace with the activities but were aware when annotation events that happened simultaneously or in quick succession that the annotations were often delayed. There were sufficient periods of inactivity during the scenario that this did not cause problems however. The participants were also working in the knowledge that they were not being tasked with providing a comprehensive annotation set for the scenario but just to annotate things as they observed them.

Problems encountered included that inability to delete or undo an annotation when errors occurred, as is quite possible if trying to annotate during a session. One participant also felt that it was important to indicate that nothing was happening, wanting to “note no change.”

The annotations captured can provide an index into the video for use in debriefing as the named annotations provide cues for the mentor that help them identify points of interest. If the mentor wishes to jump to the point in the video where the students attempted to resuscitate the patient then annotations such as “moved crash trolley,” “fitted oxygen mask,” etc., would implicitly identify the period of interest. Reflecting on our earlier posed questions, annotations in this form can help answer *where in the scenario are we* and *what are the students doing?*

Similarly, the textual annotations recorded could be played back to the students as overlay captions on the video. A prototype tool can be seen in Fig. 8 that illustrates how this would look. In this case, the students are able to supply the connection between the annotation and who is carrying out the activity visually from the video. The student about whom the annotation “Taking patient’s pulse,” refers to, can be visually identified from the video. Here again, the knowledge of the student can make the connections between the annotated actions and the person

performing those actions. More detailed descriptions of the textual replay tool and scenarios can be found in [49] and [50]. Analysis of the capture annotations from this demonstrator however suggests that the use of book-marking would make many of the annotations less useful to anyone but those having made the annotations. Before they could be presented to the student filtering, or a second pass of more detailed annotation, would be required. The suggestion from the study is that annotations created are unlikely to be directly reusable in multiple contexts. If the intention is for their immediate reuse by students for feedback then they will need to be authored with this specifically in mind. In our case, the annotators were reflecting on student actions with a view to assessment, so some of the annotations would not be easily interpreted directly by students.

Although providing a good quick index into the material, or detailed feedback to the students when accompanied by the video, the data are not complete, relying on the viewer to make the connections between person and action. In simple cases, this is a straightforward process. For the annotation “raises the bed” the action is likely performed by a single individual who can easily be identified from the video. When analyzing the observations of a number of different observers it will not always be possible to identify who is being referred to with less visual annotations such as “display of tacit knowledge” or “Routine scan of monitor.” When examining the process of learning itself, we may be interested in when observers are noticing the same thing. This type of enquiry can help begin to identify what it is that identifies a good student to an observer. The observers may be conscious of why they identify the student as being good, however, it could be an act of “coup d’oeil” or “the power of the glance,” the ability to see and immediately make sense of a situation even if the individual contributing factors might only register at a subconscious level. To investigate this, further a new approach was taken. Here, observers explicitly identify the participants on the video by clicking on them. To incorporate these more explicit indications of attention, textual annotation making was replaced by an audio commentary, with the commentary providing alternative answers to the questions *what did the student do?* and *what assessment can I make from their performance.*

6.2 Audio Annotation

For the audio annotation trials, five observers watched the same video of a session and performed a think aloud annotation of what they were seeing in the video accompanied by clicking on participants in the video as they discussed them. A stand alone annotation tool was created to facilitate this (based on the same ontology) with an accompanying playback tool that allowed the replay of the video alongside a number of possible audio annotation tracks (see Fig. 9.) The tool also allowed for the creation of manual annotations based on the previous ontology in a second phase. This provided a simple mechanism to transcribe the audio annotations into a more machine readable format.

When the annotations are replayed, the area clicked on is highlighted. Multiple annotation files can be loaded into the replay tool with multiple foci of interest highlighted

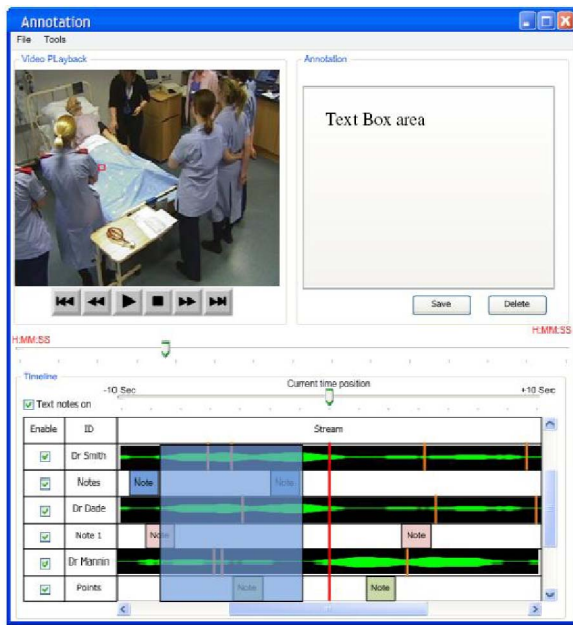


Fig. 9. Listening to the audio annotations.

on the video. This provides a simple mechanism to index into the video to help identify when a number of observers are noticing the same thing (*what are the educators noticing?*). The researcher can then play the audio recorded for the different observers at that point in the video to investigate what it is that is catching their attention and add textual annotations. Similarly, outliers, where an individual notices something different from the majority, may also provide potential areas for exploration. The visual markers are performing the function of narrowing the search field from the entirety of the audio commentaries to just those segments that may be of interest based on common focus of interest.

Through analysis of the resulting data, a number of events were noticed in the video where the multiple observers identified the same activity taking place. Although the sample is small, patterns do appear that seem to reflect the specialisms of the observers. Those more specialized in primary care were more likely to identify certain types of events taking place for instance. The approach was seen as useful in providing a research tool to better understand how educators observe students performing skills-based tasks, but the annotations again tended to be specific to this research question and were less likely to be useful for student feedback.

As with the textual annotations previously however, the replay and understanding of the activities relies on knowledge of the viewer. The underlying system can only record the region of the video identified at a given moment in time, it is up to the researcher to make the connection between that region and a particular individual displayed there, there is no explicit recording of individuals in this system. Because audio annotations are less machine processable than the textual annotations, like other video annotation systems such as DIVER [17] the system allows the annotator to attach textual annotations to segments of the video in a postprocessing mode. This could allow the annotator to attach information about an individual to an annotation but we might prefer to do this more automatically.

This would then allow us to interrogate the data to ask questions such as *when did student X wash their hands?* or *show me all the activities of student Y*. In order to begin to address the problem of identifying participants more explicitly and automatically, a third approach was adopted, that of location-based annotation.

6.3 Location-Based Annotation

In these trials, we combined the manually authored annotations with information gathered from a location tracking system deployed in the lab. Coupled with information from the ontology simple rule-based inferences allow the construction of more complex semantic information associating individuals taking part in the scenario with actions being recorded by the observers. Coyle et al. [51] have demonstrated the benefits of aggregating location information from multiple sources. Our approach also integrates information from multiple sources of which location is one. Other location-based systems have been deployed in Healthcare settings but with more specific aims such as indoor wayfinding for people with cognitive impairments [52] or tagging objects using RFID to aid information tracking [53].

The location tracking system used in this trial was a commercially available ultrawideband (UWB) radio frequency real-time location system called Ubisense.² It claims to provide location accuracy down to 15 cm in three dimensions (although we only used two) and real-time subsecond response. However, for the purposes of these trials we were less concerned with fine grained accuracy and it was hoped that the roughly rectangular ward environment was not expected to test its claims to operate in “challenging environments.” The Ubisense sensors were deployed at ceiling height. This was intended to minimize interference with the normal running of the lab and avoid distracting the students when they are immersed in the scenario. In a full-scale deployment, these commercial sensors can be easily installed permanently, but for our trials a temporary installation was put up and calibrated successfully in half a day. The calibration of the Ubisense sensors only fixes their position in relation to each other, and hence the tags in relation to the sensors. To provide a more detailed map of the room it was measured and the Ubisense tags were used to trace paths around the ward and identify the position of specific objects within it. This was carried out while the cameras were recording. This allowed the coupling of the measurements of the physical space with the abstract coordinate space of the sensor system helping construct a mapping between the physical room and the virtual coordinate system.

The sensors pick up battery powered Ubisense tags, which are small and light enough to be worn with an acceptably minor impact on the participants actions and behavior. Previous deployments of the technology by colleagues on other projects had found that tight grouping of people can cause degradation to the radio signal; the teaching scenarios inevitably involve the bunching of students around a patient, so we attached the tags to the epaulettes on the shoulders of the student uniform

2. <http://www.ubisense.net>.



Fig. 10. Replaying the scenario in the Digital Replay System.

(positioning of the sensors at ceiling height also helped minimize interference). Once a session had started it would be difficult to make technical adjustments, so we built in redundancy by instrumenting two of the students with a tag on both shoulders. This both allowed for possible error and expanded the data set available to explore of location consistency (two sensors moving together a fixed distance apart) and orientation of the student. Data gathered however suggests that orientation prediction would not be reliable using this technique in this configuration.

In addition to the student participants, other actors in the scenario, such as the mentors playing the roles of ward sister and doctor, also wore tags. Two pieces of mobile equipment in the lab, the dressing trolley and crash trolley, were also tagged as the scenario would involve the movement of these items by the students. The scenario was as before, with a monitored patient deteriorating suddenly and going into cardiac arrest.

The UbiSense software logs a time-stamped record of updated tag positions. A mapping was recorded manually between the tag IDs that appear in the logs and the participants and objects that were wearing the tags. This additional semantic data, form another source of information for the subsequent connection of the different information sources. As time is the key axis against which we wish to align our annotations and video, having a fixed point common to all media and annotations was necessary for synchronization purposes. To achieve this, we used the buttons on a UbiSense tag to register an event in the annotation log. The button press was synchronized to a verbal countdown on the video/audio recording (in the spirit of a clapperboard). A more permanent installation would synchronize all the capture machines to a common reliable clock source and have the different infrastructure systems more closely coupled.

A simulation was carried out with volunteer students and members of staff. Data were collected from the location tracking system as well as from two observers manually annotating in the control room using the previously described system. Post trial analysis was carried out on the data to investigate the efficacy of




Fig. 11. A participant at the sink hand washing.

combining the multiple annotation streams in order to identify new annotations.

To replay the data captured during the location tracking trials, the Digital Replay System (DRS) was used. DRS is a software tool to support the coordinated replay, annotation, and analysis of combinations of video, audio, transcripts, images, and system log files [54]. DRS enables time-based data—i.e., system recordings and audio/visual recordings—to be combined and replayed side-by-side and for annotations to be added to create new representations. DRS was extended to view and analyze the data from the location trial. Log file importers were written for the text-based observational annotations and the UbiSense location logs; a data viewer has also been written to visualize the UbiSense logs. These annotation sources are combined with the audio/video recordings from the ward and analyses can then be constructed that allow the user to navigate through the data from different standpoints (Fig. 10).

The analysis of captured data looked at the feasibility of generating additional “meta-annotations” through inference over the multiple annotations captured. The following worked example takes the actual data collected during the trial and examines how it could be connected together by following the RDF graph being constructed from the semantic annotations. Although the data seem to support the process it is acknowledged that this is a simple case but hopefully provides a clear and understandable example.

Fig. 11, shows one of the participants in the trial washing their hands at the sink. A manual annotation to this effect was created by the observer annotating from the video feed in the control room. However, as discussed previously, the manual annotation system alone means that the person performing the hand washing is unlikely to be identified.

The manual annotation is recorded as follows:

```
<annotation id="annotation_6659736"
timestamp="2008-07-31 10:40:13Z">
  <name>hand_washing</name>
  <description>Hand washing</description>
  <person></person>
</annotation>
```

or (person?, performs_activity, hand_washing).




Fig. 12. A plot of the captured location data around this time.

Tuples will be used throughout this worked example for convenience. We are also assuming that the information being combined is time synchronized, i.e., we are able to examine a snapshot of the activity space. The sink has a fixed location within the space and in our ontology we can identify the location at which “hand washing” takes place:

(hand_washing, has_location, sink).

By examining the location data around the time that the annotation was created we can infer who was present at the sink at that time (see Fig. 12). This location processing was not automated in our analysis and for a fully automated system we would imagine a polygon matching module that would provide a range of automatic matching techniques such as “within region,” “nearest,” “within x distance of,” etc. The location data taken from *tag136* are plotted in Fig. 12. The second plot in this case corresponds to the student who is visible in the foreground of the video (Fig. 11).

From this location information, we can establish that at the time of the manual annotation creation, *tag136* is the nearest tag to the sink:

(tag136, is_near, sink).

In turn, *tag136* can be identified as the participant “Eva.” From the RDF tuple, we can see how walking the RDF graph would allow us to fill in the initial gap in the first tuple and establish that:

(Eva, performs_activity, hand_washing),

demonstrating how we can construct more detailed semantic descriptions of the activities from the disparate sources of semantic information in the system. We might imagine other examples such as identifying who is fitting an oxygen mask through having information about which participants are stood at the head of the bed, or who is answering the phone at the nurses station by connecting participant location and an annotation sent from a sensor in the phone handset.

In addition, the measurements taken of the lab and calibration with the tracking system might allow us to identify regions of the space that are viewable from the cameras although it should be noted that the cameras are movable during recording so identifying which cameras a participant is viewable on from location information in this

way is likely to be problematic. Other approaches such as AR glyphs in the space might provide a solution to these problems. Such approaches have previously been adopted for indoor tracking systems [55].

Using the fusion of the Semantic Web information from the different annotation sources allows us to make inferences across the data to produce more detailed semantic understanding of the activity, which in turn can be used to provide additional functionality across the data. This worked example, although supported by the data, was not performed automatically and the consistency of the sensor data suggests that more work is needed on preprocessing and filtering of the raw location information. The mechanisms for performing the spatial queries described above would also need formalization for such a system to perform inferences of this type automatically. There are also known issues with using Semantic Web inferencing across triple stores in real-time, which would need to be addressed. Postprocessing of the annotations to create more detailed annotations would cater for the majority of expected uses. The examples chosen here have largely dealt with what could be termed spatial queries. Other interesting issues arise when dealing with temporal queries such as identifying repeated events. Further work is planned on investigating issues around performing inference on streams of RDF data.

The location tracking data have further potential as a primary data source. When synchronized with the activity data, we expect to extract and infer from the combined semantic annotations, it could provide valuable insight into streamlining ward layout. The same activity data could be fed back into the text-based annotation tool on a real-time basis, automatically informing the choice of annotations and further enhancing the process. If we know a participant is near the sink then washing related annotations could be offered as likely to be appropriate.

Our results indicate that the combination of location-based semantic information with manually authored annotations can begin to provide answers to questions such as *What did the student do?*, allowing the explicit connection of person and activity into a machine readable form.

7 CONCLUSION

Through the trialing of a series of small-scale prototypes, we have investigated the benefits of semantic annotation in understanding learning activities taking place within a simulated ward environment. The development of ontologies for both annotation and tracking offer interesting potentials for future modeling in complex environments (e.g., our own work and the social interaction ontologies outlined by Chen et al. [56] when they conducted audio-visual analysis of elderly people in a nursing home). Analysis of the logs has highlighted a number of issues with the nature of the annotations. The annotation act is purposeful to the participants and the annotations created are not necessarily generically useful. Those annotations created with the intention of student feedback were generally less useful for purposes of research analysis and vice versa. Interviews conducted showed the participants to

be comfortable with using the software and they felt the annotations they had created were useful.

Quality improvement initiatives require close attention to processes, interactions, and resources. The ability to simulate change and evaluate it before deployment could be crucial to effective implementation of new initiatives. We have demonstrated that analysis of real time activities offers huge potential once the appropriate techniques and tools have been more fully developed and tested. Preliminary usability evaluation of our systems suggests that practitioners find such annotation tools usable and can see benefits in the data that they produce. The outcomes of such work could offer insights into new and better ways of working; tools to train and educate staff to be more effective and self-reflective; strategies and tools to measure, collect, and analyze different data streams; and modeling of clinical environments to better reflect the activities within the environments.

The system could also be used by students in the longer term by allowing the students to make annotations of their activities during their placements. The reuse of the ontology would provide a link between their placement and the knowledge acquired in the university learning environment. Parallel work that considered the potential of video analysis in the assessment of student performance indicates that an annotation facility could help realize effective formative and assessment strategies [57].

We believe annotations derived from the location data to be a useful bridge between observational text annotations and the full video record of the session. We have shown how automatically gathered location information can be combined with manual authored annotations to provide more detailed descriptions of activities taking place within the learning space. By using extensible ontologies we expect to also integrate annotations from the SimMan mannequins, and extend capture to other pervasive sensors should they be installed (telemetry from other equipment, light switches, sensors on soap dispensers, etc.) It is the expressiveness, interoperability, and common vocabulary that can be constructed using RDF and Semantic Web technologies that makes it highly suitable for constructing these types of information systems. The representation of time in RDF triples is not trivial and issues of synchronization of annotation streams will be important to address.

The capturing of detailed annotations of student activities during skills-based sessions is also allowing researchers in nursing to look in more detail at the teaching process itself, and providing a record of what nursing educators see when they watch students carry out the scenarios. Analysis of this record may provide some insight into deeper research questions around the assessment and education of students in such sessions.

REFERENCES

- [1] M. Gobbi, E. Monger, G. Watkinson, A. Spencer, M. Weaver, J. Lathlean, and S. Bryant, "Virtual Interactive Practice: A Strategy to Enhance Learning and Competence in Health Care Students," *Proc. Medinfo*, pp. 874-878, 2004.
- [2] M. Seropian, K. Brown, J.S. Gavilanes, and B. Driggers, "Simulation: Not Just a Manikin," *J. Nursing Education*, vol. 43, no. 4, pp. 164-169, 2004.
- [3] R. Hays, H. Davies, J. Beard, L. Caldon, E. Famer, and P. Fincane, "Selecting Performance Assessment Methods for Experienced Physicians," *Medical Education*, vol. 36, pp. 910-917, 2002.
- [4] T. Blackburn and C. Sadler, "The Role of Human Simulators in Health-Care Training," *British J. Hospital Medicine*, vol. 64, no. 11, pp. 677-681, 2004.
- [5] T. D. of Health, "Delivery 21st Century IT Support for the NHS: National Strategic Programme," technical report, The Stationary Office, June 2002.
- [6] M. Gobbi and E. Monger, "Using Web Based Multimedia, Simulated and Virtual Practice to Assess Students Professional Practice Skills," technical report, HEA, 2008.
- [7] M. Hatala, R. Wakkary, and L. Kalantari, "Rules and Ontologies in Support of Real-Time Ubiquitous Application," *J. Web Semantics: Science, Services and Agents on the World Wide Web*, vol. 3, pp. 5-22, 2005.
- [8] G. Hughes, H. Mills, D. De Roure, J. Frey, L. Moreau, M.C. schraefel, G. Smith, and E. Zaluska, "The Semantic Smart Laboratory: A System for Supporting the Chemical Scientist," *Organic and Biomolecular Chemistry*, vol. 2, pp. 1-10, 2004.
- [9] R. Masuoka, B. Parsia, and Y. Labrou, "Task Computing—The Semantic Web Meets Pervasive Computing," *Proc. Second Int'l Semantic Web Conf.*, pp. 866-881, Oct. 2003.
- [10] H. Chen, F. Perich, T. Finin, and A. Joshi, "SOUPA: Standard Ontology for Ubiquitous and Pervasive Applications," *Proc. Int'l Conf. Mobile and Ubiquitous Systems*, pp. 258-267, 2004.
- [11] K. Grønbaek, P.P. Vestergaard, and P. Ørbæk, "Towards Geo-Spatial Hypermedia: Concepts and Prototype Implementation," *Proc. 13th ACM Conf. Hypertext and Hypermedia*, pp. 117-126, 2002.
- [12] R. Ballagas, M. Ringel, M. Stone, and J. Borchers, "Istuff: A Physical User Interface Toolkit for Ubiquitous Computing Environments," *Proc. SIGCHI Conf. Human Factors in Computing Systems (CHI '03)*, pp. 537-544, Apr. 2003.
- [13] L. G., "Urban Tapestries: Wireless Networking, Public Authoring and Social Knowledge," *Personal Ubiquitous Computing*, vol. 7, nos. 3/4, pp. 169-175, July 2003.
- [14] P. Pan, C. Kastner, D. Crow, and G. Davenport, "M-Studio: An Authoring Application for Context-Aware Multimedia," *Proc. 10th ACM Int'l Conf. Multimedia*, pp. 351-354, Dec. 2002.
- [15] Y. Li, J.I. Hong, and J.A. Landay, "Topiary: A Tool for Prototyping Location-Enhanced Applications," *Proc. Symp. User Interface Software and Technology (UIST '04)*, pp. 217-226, Oct. 2004.
- [16] G. Iacucci, J. Kela, and P. Pehkonen, "Computational Support to Record and Re-Experience Visits," *Personal and Ubiquitous Computing*, vol. 8, pp. 100-109, 2004.
- [17] R. Pea and R. Lindgren, "Video Collaboratories for Research and Education: An Analysis of Collaboration Design Patterns," *IEEE Trans. Learning Technologies*, vol. 1, no. 4, pp. 235-247, Oct.-Dec. 2008.
- [18] H. Ogata, T. Misumi, T. Matsuka, M.M. El-Bishouty, and Y. Yano, "A Framework for Capturing, Sharing and Comparing Learning Experiences in a Ubiquitous Learning Environment," *Research and Practice in Technology*, vol. 3, no. 3, pp. 297-312, 2008.
- [19] L. Ballan, M. Bertini, A. Del Bimbo, L. Seidenari, and G. Serra, "Event Detection and Recognition for Semantic Annotation of Video," *Multimedia Tools and Applications*, vol. 51, pp. 279-302, 2011.
- [20] D. Murray, J. Boulet, J. Kras, J. Woodhouse, M. Thomas, and J. McAllister, "Acute Care Skills in Anaesthesia Practice: A Simulation Based Student Performance Assessment," *Anaesthesiology*, vol. 101, no. 5, pp. 1084-1095, 2004.
- [21] R. Kneebone, D. Nestel, F. Yadollahi, R. Brown, C. Nolan, J.D.H. Brenton, C. Moulton, J. Archer, and A. Darzi, "Assessing Procedural Skills in Context: Exploring the Feasibility of an Integrated Procedural Performance Instrument (IPPI)," *Medical Education*, vol. 40, pp. 1105-1114, 2006.
- [22] J. DiGiacomo, W. Hoff, M. Rotondo, K. Martin, D. Kauder, H. Anderson, G. Phillips, and W. Schwab, "Barrier Precautions in Trauma Resuscitation: Real-Time Analysis Utilising Videotape Review," *Am. J. Emergency Medicine*, vol. 15, no. 1, pp. 34-39, 1997.
- [23] P. Ram, R. Grol, J. Rethans, B. Schouten, C. van der Vleuten, and A. Kester, "Assessment of General Practitioners by Video Observation of Communicative and Medical Performance in Daily Practice: Issues of Validity, Reliability and Feasibility," *Medical Education*, vol. 33, pp. 447-454, 1999.

- [24] P. Vivekananda-Schmidt, M. Lewis, D. Coady, C. Morley, L. Kay, D. Walker, and A. Hassell, "Exploring the Use of Videotaped Objective Structured Clinical Examination in the Assessment of Joint Examination Skills of Medical Students," *Arthritis and Rheumatism*, vol. 57, no. 5, pp. 869-876, 2007.
- [25] T. Berners-Lee, J. Hendler, and O. Lassila, "The Semantic Web: Scientific American," *Scientific Am.*, May 2001.
- [26] N. Shadbolt, T. Berners-Lee, and W. Hall, "The Semantic Web Revisited," *IEEE Intelligent Systems*, vol. 21, no. 3, pp. 96-101, July 2006.
- [27] T. Strang and C. Linnhoff-Popien, "A Context Modeling Survey," *Proc. First Int'l Workshop Advanced Context Modelling, Reasoning and Management*, Sept. 2004.
- [28] I. Millard, D.D. Roure, and N. Shadbolt, "Contextually Aware Information Delivery in Pervasive Computing Environments," *Proc. Location and Context Awareness*, pp. 189-197, Jan. 2005.
- [29] O. Lassila and M. Adler, "Semantic Gadgets: Ubiquitous Computing Meets the Semantic Web," *Spinning the Semantic Web*, D. Fensel, ed., MIT, pp. 363-376, 2003.
- [30] K.R. Page, D. Cruickshank, and D.D. Roure, "It's About Time: Link Streams as Continuous Metadata," *Proc. 12th ACM Conf. Hypertext and Hypermedia (Hypertext '01)*, pp. 93-102, Aug. 2001.
- [31] K.R. Page, D.T. Michaelides, S.J.B. Shum, Y.-H. Chen-Burger, J. Dalton, D.C. De Roure, M. Eisenstadt, S. Potter, N.R. Shadbolt, A. Tate, M. Bachler, and J. Komzak, "Collaboration in the Semantic Grid: A Basis for E-learning," *Applied Artificial Intelligence*, vol. 19, nos. 9/10, pp. 881-904, Nov. 2005.
- [32] M.J. Weal, D.T. Michaelides, D.E. Millard, D.C.D. Roure, K. Howland, and G. Fitzpatrick, "Card Based Metaphor for Organising Pervasive Educational Experiences," *Proc. IEEE Fifth Int'l Conf. Pervasive Computing and Comm. Workshops (PERCOMW '07)*, pp. 19-23, 2007.
- [33] D. Brickley and L. Miller, "FOAF Vocabulary Specification," technical report, WC3, 2005.
- [34] T. Tiropanis, H. Davis, D. Millard, and M. Weal, "Semantic Technologies for Learning and Teaching in the Web 2.0 Era—A Survey," *Proc. Society On-Line (WebSci '09)*, <http://eprints.ecs.soton.ac.uk/17106>, 2009.
- [35] P. Anderson, "What Is Web 2.0? Ideas, Technologies and Implications for Education," technical report, JISC Technology and Standards Watch, Feb. 2007.
- [36] J. Ohler, "The Semantic Web in Education—What Happens When the Read-Write Web Gets Smart Enough to Help Us Organize and Evaluate the Information it Provides?" *EDUCAUSE Quarterly Magazine*, no. 4, 2008.
- [37] J.J. Carroll, I. Dickinson, C. Dollin, D. Reynolds, A. Seaborne, and K. Wilkinson, "Jena: Implementing the Semantic Web Recommendations," *Proc. 13th Int'l World Wide Web Conf. Alternate Track Papers and Posters (WWW Alt. '04)*, pp. 74-83, 2004.
- [38] T. Young, S. Brailsford, C. Connell, R. Davies, P. Harper, and J.H. Klein, "Using Industrial Processes to Improve Patient Care," *BMJ*, vol. 328, no. 7432, pp. 162-164, 2004.
- [39] S.C. Brailsford, "Advances and Challenges in Healthcare Simulation Modeling: Tutorial," *Proc. 39th Conf. Winter Simulation (WSC '07)*, pp. 1436-1448, 2007.
- [40] C. Roberts, V. Wass, R. Jones, S. Sarangi, and A. Gillett, "A Discourse Analysis Study of Good and Poor Communication in an Osce: A Proposed New Framework for Teaching Students," *Medical Education*, vol. 37, no. 3, pp. 192-201, 2003.
- [41] J.E. Bardram and C. Bossen, "A Web of Coordinative Artifacts: Collaborative Work at a Hospital Ward," *Proc. Int'l ACM SIGGROUP Supporting Group Work (GROUP '05)*, pp. 168-176, 2005.
- [42] J. Favela, M. Tentori, L.A. Castro, V.M. Gonzalez, E.B. Moran, and A.I. Martinez-Garcia, "Estimating Hospital Work Activities in Context-Aware Healthcare Applications," *Proc. First Int'l Conf. Pervasive Computing Technologies for Healthcare*, pp. 1-10, 2006.
- [43] J.A. Muras, V. Cahill, and E.K. Stokes, "A Taxonomy of Pervasive Healthcare Systems," *Proc. First Int'l Conf. Pervasive Computing Technologies for Healthcare*, pp. 1-10, 2006.
- [44] M. Gustafsson, G. Falkman, F. Lindahl, and O. Torgersson, "Enabling an Online Community for Sharing Oral Medicine Cases Using Semantic Web Technologies," *Proc. Fifth Int'l Semantic Web Conf. (ISWC '06)*, pp. 820-832, 2006.
- [45] V. Kashyap, A. Morales, T. Hongsermeier, and Q. Li, "Definitions Management: A Semantics-Based Approach for Clinical Documentation in Healthcare Delivery," *Proc. Int'l Semantic Web Conf. (ISWC '05)*, pp. 887-901, 2005.
- [46] R. Jain, "Eventweb: Developing a Human-Centered Computing System," *Computer*, vol. 41, no. 2, pp. 42-50, Feb. 2008.
- [47] F. Nack and W. Putz, "Designing Annotation Before it's Needed," *Proc. Ninth ACM Int'l Conf. Multimedia (MULTIMEDIA '01)*, pp. 251-260, 2001.
- [48] M. Fraser, J. Hindmarsh, K. Best, C. Heath, G. Biegel, C. Greenhalgh, and S. Reeves, "Remote Collaboration Over Video Data: Towards Real-Time E-social Science," *Computer Supported Cooperative Work (CSCW)*, vol. 15, no. 4, pp. 257-279, Aug. 2006.
- [49] M.J. Weal, D.T. Michaelides, D.C. De Roure, M. Gobbi, E. Monger, and J.W. McDonald, "Semantic Annotation in Ubiquitous Healthcare Skills-Based Learning Environments," *Proc. Workshop Semantic Web in Ubiquitous Healthcare (ISWC)*, Nov. 2007.
- [50] J.W. McDonald, M.O. Gobbi, D. Michaelides, E. Monger, M.J. Weal, and D. De Roure, "Grid-Enabled Data Collection and Analysis: Semantic Annotation in Skills-Based Learning," *Proc. Fourth Int'l Conf. E-Social Science*, June 2008.
- [51] L. Coyle, S. Neely, P. Nixon, and A. Quigley, "Sensor Aggregation and Integration in Healthcare Location Based Services," *Proc. Pervasive Health Conf. and Workshops*, pp. 1-4, 2006.
- [52] A.L. Liu, H. Hile, G. Borriello, H. Kautz, B. Ferris, P.A. Brown, M. Harniss, and K. Johnson, "Implications for Location Systems in Indoor Wayfinding for Individuals with Cognitive Impairments," *Proc. Pervasive Health Conf. and Workshops*, pp. 1-5, 2006.
- [53] J. Bravo, R. Hervás, C. Fuentes, G. Chavira, and S.W. Nava, "Tagging for Nursing Care," *Proc. Second Int'l Conf. Pervasive Computing Technologies for Healthcare*, pp. 305-307, 2008.
- [54] C. Greenhalgh, A. French, P. Tennant, J. Humble, and A. Crabtree, "From Replaytool to Digital Replay System," *Proc. Third Int'l Conf. E-Social Science*, Oct. 2007.
- [55] G. Reitmayr and D. Schmalstieg, "Location Based Applications for Mobile Augmented Reality," *Proc. Fourth Australasian User Interface Conf. User Interfaces (AUIC '03)*, pp. 65-73, 2003.
- [56] D. Chen, J. Yang, R. Malkin, and H.D. Wactlar, "Detecting Social Interactions of the Elderly in a Nursing Home Environment," *ACM Trans. Multimedia Computing Comm., and Applications*, vol. 3, no. 1, p. 6, 2007.
- [57] M. Gobbi and E. Monger, "Using Web Based Multimedia, Simulated and Virtual Practice to Assess Students? Professional Practice Skills," technical report, HEA HSP Centre and Univ. of Southampton, 2008.



Mark J. Weal is a lecturer in the Web and Internet Science Group, School of Electronics and Computer Science, University of Southampton. His research interests include Web Science and the application of Semantic Web technologies in healthcare, e-learning, and pervasive systems.



Danius T. Michaelides is a senior research fellow in the Web and Internet Science Group, School of Electronics and Computer Science, University of Southampton. His current research interests include developing research environments for scientists.



Kevin Page is a research fellow in the Web and Internet Science Group, School of Electronics and Computer Science, University of Southampton. His current research interests include information systems for sensor nets.



Eloise Monger is a lecturer in Critical Care Nursing at the University of Southampton. Her research interests include the use of simulation and virtual interactive practice for educational development and research and research ethics in nursing.



David C. De Roure is professor of e-research in the Oxford e-Research Centre and the National Strategic Director for Digital Social Research. He focuses on the coevolution of digital technologies and research methods in and between multiple disciplines. He is a fellow of the IEEE.



Mary Gobbi is a senior lecturer in nursing and the Erasmus Coordinator for the Faculty of Health Sciences, University of Southampton. Her research interests include educational development and research with simulation and virtual interactive practice for patient safety, leadership, use of technologies in nursing, and accelerated student/practitioner competence.