

Location based semantic annotation for ward analysis

Mark J. Weal*, Danius T. Michaelides*, Kevin R. Page*, David C. De Roure*,
Mary Gobbi[†], Eloise Monger[†] and Fernando Martinez[‡]

*School of Electronics and Computer Science, University of Southampton, UK. Email: {mjw}{dtm}{krp}{dder}@ecs.soton.ac.uk

[†]School of Health Sciences, University of Southampton, UK.

Email: {m.o.gobbi}{em}@soton.ac.uk

[‡]Mixed Reality Lab, University of Nottingham, UK.

Email: fxm@cs.nott.ac.uk

Abstract—Semantic annotation has been used to combine varied information sources – gathered as unobtrusively as possible – and produce enhanced tools for working with digital resources. In this paper we describe trials carried out using a location tracking system and Semantic Web annotation technologies to analyse activities in a simulated ward environment. The motivation for semantic annotation of the space will be outlined along with the practicalities of the location based tracking system. The integration of location, annotations and video information will be discussed together with the technologies and approaches applicability to use in a real ward environment.

I. INTRODUCTION

This paper discusses the technical and practical issues encountered through the use of Semantic Web technology to capture information in a ward setting. The data collected included video streams, captured annotations, location tracking data of people (nursing students and staff) and ward equipment, and patient sensor data from a computerized mannequin (SimMan). The paper builds on previous work in the context of learning environments and video replay [1].

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. Audio and video together present a highly detailed capture of an activity; perhaps too detailed, because reviewing a recording can be almost as time-consuming as the original activity. Annotating a video ‘by hand’ is similarly intensive, as it often involves reviewing the entire record, or making annotations as the activity occurs - which is often onerous enough to preclude full engagement in the activity itself. We look, therefore, to augment our record with semantic annotations from as many sources as possible, and from sources that have a low impact and overhead on the participants.

Broadly speaking, there are pragmatic, policy and management reasons why it is important to develop the science and analysis of tracking and annotating people and equipment through simulated hospital ward activity.

First, the simulated ward environment enables the safe and ethical development of tools that could subsequently be deployed for use in the ‘real’ world of health care practice,

for example, designing location trackers that are health and safety compliant within a clinical or home setting.

Second, simulated environments, equipped with audio visual capture facilities, are used increasingly for the education and training of a range of students and staff. The data generated in these situations provides a rich resource for educational and research analysis of student performance, their interactions with each other and objects/equipment located in the environment. However 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 - the third challenge. Annotation enables the marking of the event, which can then be matched with the tracking data. Favela et al, [2] have shown how work activities can be derived from contextual information when it is available.

Fourth, the clinical incidents that form the basis of the simulated learning activity generate time markers that should stimulate student responses, for example the computerized patient (SimMan) who, at a predetermined moment exhibits altered sensory data (pulse, blood pressure, oxygen levels for example).

Fifth, 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, teamworking and leadership styles. Within the cost constraints of modern health services, strategies to improve design, process and human performance are ever present. Other approaches have sought to make similar illuminations on how co-ordination is achieved with non-digital artifacts in clinical settings [3].

Interestingly, in their discussion of the suitability of industrial methods to improve the quality and efficiency of health services, Young, Brailsford, Connell et al., [4] recommended the use of simulations to identify how these methods would




Fig. 1. The SimMan computerised mannequin.

translate before the benefits could be realized. In this instance, we are arguing that the development of robust tools to monitor, track, annotate, and analyse 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. The problems of evaluating simulations themselves have been highlighted by Brailsford [5] and the digital record collected using our approach may provide new mechanisms to evaluate the efficacy of the many simulation models currently in use.

However, 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.

II. BACKGROUND

This work builds on previous work in a number of areas.

A. Skills-based teaching environment

The test bed setting was a clinical skills laboratory that resembles a hospital ward. The ward is viewable from a central control room via 6 audio-equipped ceiling mounted cameras, each controllable with a 360 degree viewing angle. Similar to our original work [6], the setting mimicked the reality of ward life in both its behaviours and resources, equipment, clinical charts, wall displays, sounds and phones. Figure 2 shows the lab in which the simulations are carried out. For the purposes of this trial the ward was equipped with computerised and interactive SimMan®¹ mannequins, non computerised mannequins and a range of equipment that were purposively arranged to provide clinical activities for the students (see Figure. 1). SimMan can be programmed to mimic varying medical conditions as the scenario progresses [7]), while an inbuilt loudspeaker allows a remote operator to provide the patient with a voice. SimMan can also log

¹<http://www.laerdal.com/SimMan>



Fig. 3. Video capture of a training session.

the events it produces and these form another source of annotations captured during the sessions.

Three volunteer and consenting students and a supervising registered nurse wore their clinical uniforms and agreed to operate as a clinical team looking after the simulated patients in the ward. The students received a report that typified genuine practice and engaged in the scenarios designed for the occasion. The students were given a plethora of tasks to engage them and the computerised mannequins were programmed to alter their parameters to a point of significant deterioration when emergency responses were required. These activities provoked the students to move themselves and equipment around the ward, to interact with each other, the supervising staff member and to use the telephone. This also meant that not only was concurrent activity taking place in different parts of the ward but there was plenty of background noise and movement. The location tracking and video capture involved both fixed stationery items and sound sources (sinks, running water, telephones) but also mobile ones (e.g. students, resuscitation equipment and trolleys).

The six ceiling-mounted rotating cameras in each lab provide full video coverage and microphones are suspended above the beds to record audio. These 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. In this experiment, the camera angles were positioned to enable maximum coverage of the ward and to prevent the necessity to move cameras during the trial. Simultaneous recordings of the simulation were made from the six cameras, although only one was used for the purposes of annotation. This real time annotation was undertaken by a nurse member of the team viewing the camera that focused upon the computerised deteriorating simulated patient.

Previous work has developed an annotation system that allows academics in the control room to annotate the activities of the students in real-time according to a co-designed ontology [1]. This annotation system provides a basis for both



Fig. 2. The skills-based lab with beds, mannequins and equipment; also ceiling mounted cameras and Ubisense sensors.

student reflection, in the replaying of videos with overlaid annotations, and also for analysis of student learning. Additional work coupled to an on-line learning resource which overlaps with the learning objectives of the skills-based scenario has also been carried out [8].

The simulations are designed to promote the acquisition of practical skills as well as decision-making, team working, communication and problem-solving etc. [6]. They are incorporated into assessment of student performance [9] requiring that the approaches for assessment and feedback need to be sound, valid, reliable, feasible, educational, and importantly acceptable to practitioners [10]. Although simulated, the student experiences are exactly as they would experience in the workplace in real time. Hence, when the students and mentors are ‘immersed’ in the simulation and behave ‘as in real practice’ the video captured data provides important information about their performance. This 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” [11].

Although the driver for the creation of a semantic annotation for the skills-based lab activities has been its use in a learning context, the underlying infrastructure can be seen as annotating activities in a ward environment (albeit simulated) and in this paper we will describe how we consider it could be augmented with spatial annotations derived from an indoor location tracking system and might equally be repurposed in hospital or other environments to answer a different set of questions.

B. Ethical Framework

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 (a) the access of patient related material to the research team; (b) the collection, storage and dissemination of staff and student and data that could not be anonymised and (c) the ability to provide a sound educational experience for these volunteer students. The students and staff members willingly participated and contributed to the debriefing session of the trial with suggestions for improvement and an outline of their experiences.

C. Location Tracking

The Ubisense system, deployed in this project, is a commercially available ultra-wideband (UWB) radio frequency real-time location system. It claims to provide “15cm 3D accuracy, real-time sub-second response and proven robustness in challenging [...] environments”². The roughly rectangular ward environment was hoped not to be at the challenging end of the spectrum of environments despite the presence of a number of electrical devices within the space.

In line with the existing capture technology in the skills-based lab, the Ubisense sensors were also deployed at ceiling height. This was intended to minimise interference with the ‘normal’ running of the lab and avoid distraction from student immersion in the scenarios (Fig. 4).

²<http://www.ubisense.net/>



Fig. 4. Ubisense sensors installed in the ward.



Fig. 5. Ubisense tags.

The sensors were connected using Ethernet cabling back to a PC in the control room where the trial was monitored and the location data logged. In a more permanent trial, the cabling could be routed through appropriate trunking, or indeed alternative wireless solutions might prove to be equally effective such as those deployed by Triantafyllidis et. al, [12]. Wireless enabled devices that would not interfere with clinical facilities and satisfied the medical devices governance requirements³⁴ would enable data collection within complex clinical environments and work has been carried out in this area [13].

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 behaviour (see Figure 5).

The Ubisense approach is just one of a number of possible approaches to indoor location tracking that could be deployed in such a setting.

A comprehensive survey of position location technologies

³<http://www.mhra.gov.uk/Howweregulate/Devices/index.htm>

⁴<http://www.nres.npsa.nhs.uk/>

can be found in Manesis and Avouris [14] comparing by functionality, specifically in terms of their approach to tackling location and motion of people and objects. Other specific approaches such as RSSI range based estimation [15], Zigbee location tracking [16], near field imaging [17] and spinning beacons based on motes [18] have also been used in healthcare settings.

Our project focuses on ward settings however, work in other healthcare settings, such as outpatient monitoring in the “Big Nurse” project [19] or the pervasive augmentation of operating theatres [20] suggests a growing role for monitoring and tracking systems in this field.

Our aim is to combine location with other data to semantically annotate the space. In the same way that Coyle et. al, [21] have demonstrated the benefits of aggregating location information from multiple sources, we propose the integration of semantic information about the space obtained from multiple sources of which location tracking is one.

Other location based systems deployed in Healthcare settings have had more specific aims such as indoor wayfinding for people with cognitive impairments[22] or tagging objects using RFID to aid information tracking [23].

D. Semantic Web

Building on earlier work which added hyper-structure to video collaborations [24], mechanisms for capturing annotations from the skills-based session have been developed using Semantic Web technology and techniques [1]. These annotations can then be used to generate an ‘index’ into the video structure through which the detail-rich record can be more effectively used.

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 semi-structured data, and in fact can be seen as a universal format for data on the Web, providing greater interoperability and re-use than XML alone.

Fundamentally RDF data describes ‘things’, even if they cannot be directly retrieved on the Web (they just need to be identified on the Web). 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, co-authorship, etc.) accumulates on those people to describe communities of practice. We have previously extracted annotations from collaborative tools (instant messaging, knowledge mapping, and intelligent task lists) to create replay structure for distributed meetings [24]. 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. So for example, we can describe objects or people as being in specific locations at specific times and these can link to

annotations made about those objects or people. In this way we would aim to have an interlinked body of knowledge that can be queried spatially and temporally in order to understand what is happening with a space over a period of time.

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 [25]. A variety of notions of context may be expressed, including location and user tasks [26]. Ontologies can also be used to describe device capabilities, for example to facilitate content delivery to devices with diverse characteristics [27]. In this work we aim to look at using geographical location to link together semantic annotations (made by observers of the activities) in order to describe the context in which tasks were carried out. For example we might want to be able to describe that at a certain point in time (time marker) a user was carrying out a certain activity (observers annotation or SimMan logging information) whilst stood next to a certain piece of equipment (spatial annotation). It is the expressiveness, inter operability 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.

An ontology was co-designed with the nursing staff to capture naturalistic time sequenced observations typical of the scenarios, with a clustering of themes according to discipline specific relationships. For example “taking a blood pressure reading” was clustered under a heading of “taking and recording vital signs”. The individual activity of the reading then breaks down into further components such as “putting the cuff on”. The ontology was not intended to be in any way comprehensive nor to encompass all pervasive activities as has been attempted with other taxonomies[28].

A simple interface was built so that a mentor in the control room can, whilst monitoring a session, quickly capture events as they occur using the ontology; an event is time stamped when selected using a mouse in the tool, after which further details can be recorded. Further details of the text annotation system and its replay can be found elsewhere [8].

Although the observational tool has been designed to be used simply and quickly – thus not 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 therefore 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.

III. CASE STUDY

As has been stated previously, the initial drivers for this work were to develop technologies to help academics better understand how students learn in skills-based learning scenarios and to provide feedback tools that will allow the students to reflect on how they performed during the task [8]. By adding the location tracking as an annotation generator it was hoped that more detailed analysis could be generated from the data

which might in turn facilitate new forms of questioning, some of which could extend beyond the learning framework and instead form questions of the space itself, i.e. a hospital ward.

A. Location Tracking in the Nursing Lab

Previous deployments had found that tight grouping of people can cause significant degradation to the radio signal; the teaching scenarios meant bunching of students around a patient was almost inevitable, so we attached the tags to the epaulettes on the shoulders of the student uniform (positioning of the sensors at ceiling height also helped minimise interference).

Once a session had started it would not be appropriate to interrupt for technical adjustments, so we instrumented two of the students with a tag on both shoulders – both for redundancy, and to expand the data set to enable the exploration of location consistency (two sensors moving together a fixed distance apart) and orientation of the student.

Other actors in the scenario, such as the ward sister and doctor, wore the tags more conventionally using a lanyard around the neck. In addition to the participants, two pieces of mobile equipment in the lab, namely the dressing trolley and crash trolley, were also tagged. Both were expected to be moved by the students during the exercise giving additional contextual information for tasks taking place.

While the Ubisense sensors must be calibrated before use, this only fixes the position of the sensors in relation to each other, and the tags in relation to the sensors. The room was measured for future mapping activities and the Ubisense tags were used to ‘trace’ paths around the ward and specific objects within it whilst the cameras were recording. This data would then allow the coupling of the measurements of the physical space with the abstract co-ordinate space of the sensor system.

The Ubisense software logs a time-stamped record of updated tag positions; the sensor cell allocates time slots for the tags it can ‘see’ based upon their motion/activity up to a maximum of 10 updates per second. 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, forms another source of information for the subsequent connection of the different information sources for inference.

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 synchronisation purposes. To achieve this we used the buttons on a Ubisense tag to register an event in the log. The button press was synchronised to a verbal countdown on the video/audio recording (in the spirit of a clapperboard). For a more permanent installation one might expect to synchronise all the capture machines to a common reliable clock source and have the different infrastructure systems more closely coupled.

IV. DISCUSSION

A. The Trials

The nursing students found the Ubisense tags easy to wear and were able to ignore their presence; similarly the sensors

and knowledge of data collection didn't cause any apparent distraction beyond that caused by the cameras and microphones. It helps, of course, that the students are engrossed in the scenario and engaged as if it were 'real life' – the additional technology quickly recedes into the background.

The Ubisense product performed adequately, though not perfectly: on numerous occasions there was some lag before the sensors reported a tag in motion, particularly after the tag had been at rest, or when a nurse moved quickly across the room; in some instances it seemed that a tag had become 'stuck', although no single tag was motionless for the entire session.

While waiting to perform further trials it is not clear whether these discrepancies were caused by problems with the tags, with the particular sensor arrangement and calibration in the ward, or a more general weakness of the technology. A comparative analysis of the tag movement data has not yet been performed, but we expect this to be an interesting test in dealing with incomplete and unreliable data – after all, it is unlikely that sensor data will ever be perfect, and inaccuracies will always be present in human sources such as the observational text annotations.

In the debrief on the trial, the students suggested ways to improve the attachment of tracking devices to their uniforms so that they would be unobtrusive and safe to wear. Work by Tentori and Favela [29] had encountered similar problems with devices carried by nursing staff, where bracelet devices displaying patient information raised issues of potential intrusion and practicality.

B. Data Collection

Having collected semantic information from the manual annotations, location tracking system and SimMan logs, we might imagine a number of possible ways of relating this data to suggest student activity within the scenario.

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. This example is chosen because the data seems to support the process and provide the material for constructing more complex inferences to be made and more complex semantic information to be constructed.

In Figure 6, we can see one of the participants in the trial washing their hands at the sink. A manual annotation to this effect was created by the lecturer annotating from the video feed however the real-time annotation constraints do not enable the person performing the hand washing to be easily identified.

From the manual annotation file we get the annotation:

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



Fig. 6. A participant at the sink hand washing.




Fig. 7. A plot of the captured location data around this time.

</annotation>

which can be represented in tuple form as:

(person?, performs_activity, "hand washing")

In order to identify the person performing the task we need to combine this with the location information.

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 Figure 7). (The second plot in this case corresponds to the student who is visible in the foreground of the video.) This location processing was not automated in our trial and for a full system we would imagine a polygon matching module that might be able to provide a range of automatic matching techniques such as 'within

region', 'nearest', 'within x distance of' etc. Further detailed analysis of the location data is required to establish the likely effectiveness of such processing. Here we have the location data taken from tag 136 plotted in Figure 7:

```
2008-07-31 10:40:12Z,,136,-6.61,-2.65
2008-07-31 10:40:13Z,,136,-8.21,-2.73
2008-07-31 10:40:13Z,,136,-8.23,-2.73
```

From this 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 tuples above 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 concrete semantic descriptions of the activities from the disparate sources of semantic information in the system.

We have information from our measuring of the fixed lab that could extend this further by comparing with the tuples containing information about the fixed views of cameras in the system. We might identify that:

```
(Camera6, has_region, polygon)
```

The system has the information available to be able to establish which camera can now see the activity 'hand washing' by participant 'Eva'.

In this way, the fusion of the Semantic Web information from the different sources allows us to make inferences across the data to produce new semantic understanding of the activity, which in turn can be used to provide additional functionality across the data. This is simply one fairly straightforward walked example. Although supported by the data, this was not performed automatically and the consistency of the sensor data suggests that more work is needed on processing and filtering of the data. The mechanisms for performing the spatial queries described above would also need formalisation 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.

With such technologies in place however, we would hope that the data could be questioned automatically in other more complex ways, for example:

- proximity to an area of the ward with a specific function implies the tagged student is performing a particular task or class of tasks – e.g. washing hands at a sink, using the phone at the nursing station.

- movement of mobile equipment indicates at least the expectation of a specific class of activity.
- the position around a bed (at the head or foot) and in relation to other participants may imply which student is performing an annotated task. Such as attaching an oxygen mask or reading the chart.

Clearly at this stage the equipment is not developed sufficiently to be suitable for real world use in health care environments.

V. CONCLUSION

More detail analysis of the data is required before we can reassure ourselves that this process could be fully automated and that the data quality is suitable to produce useful inferences in general, but we believe that these initial trials suggest that the Semantic Web approach we have adopted may be potentially fruitful in the development of automated techniques for constructing usable semantic annotations in ward environments.

The development of ontologies for both annotation and tracking offer interesting potentials for future modelling in complex environments (e.g. our own work and the social interaction ontologies outlined by Chen et al [30] when they conducted audiovisual analysis of elderly people in a nursing home).

As outlined earlier, 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. The outcomes of such work could offer:

- insights into new and better ways of working,
- reinforcement of effective patterns of behaviour and activity,
- different venues and arrangements for the location of people and fixed or mobile equipment (ergonomics)
- tools to train and educate staff to be more effective and self reflective
- enhancements to after action review procedures
- strategies and tools to measure, collect and analyse different data streams
- evidence for teamwork analysis
- modelling of clinical environments to better reflect the activities within the environments
- research into, or associated with, the above activities.

In our recent work we have trialled the use of location tracking technology to capture and augment semantic annotations in the context of a skills-based teaching lab for nurses. After an initial analysis, we believe the location data - albeit imperfect - to be a useful bridge between observational text annotations and the full video record of the session. 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.).

The location tracking data has further potential as a primary data source. When synchronised 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.

It is the expressiveness, inter operability 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.

ACKNOWLEDGMENT

This work was in part funded by NCeSS grant RES-149-25-1056 and EPSRC "Semantic Media" grant (EP/C010078/1), and constitutes a case study within EP/C010078/1.

REFERENCES

- [1] 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," in *Workshop on Semantic Web in Ubiquitous Healthcare in conjunction with ISWC2007*, Nov. 2007.
- [2] J. Favela, M. Tentori, L. A. Castro, V. M. Gonzalez, E. B. Moran, and A. I. Martinez-Garca, "Estimating hospital work activities in context-aware healthcare applications," in *1st International Conference on Pervasive Computing Technologies for Healthcare 2006*. IEEE press, 2006, pp. 1–10.
- [3] J. E. Bardram and C. Bossen, "A web of coordinative artifacts: collaborative work at a hospital ward," in *GROUP '05: Proceedings of the 2005 international ACM SIGGROUP conference on Supporting group work*. New York, NY, USA: ACM, 2005, pp. 168–176.
- [4] 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. [Online]. Available: <http://www.bmj.com>
- [5] S. C. Brailsford, "Advances and challenges in healthcare simulation modeling: tutorial," in *WSC '07: Proceedings of the 39th conference on Winter simulation*. Piscataway, NJ, USA: IEEE Press, 2007, pp. 1436–1448.
- [6] M. Gobbi, E. Monger, G. Watkinson, A. Spencer, M. Weaver, J. Lathlean, and B. S., "Virtual interactive practiceTM: A strategy to enhance learning and competence in health care students," in *Medinfo 2004*, M. e. a. Fieschi, Ed., 2004, pp. 874–878.
- [7] T. Blackburn and C. Sadler, "The role of human simulators in healthcare training," *British Journal of Hospital Medicine*, vol. 64, no. 11, p. 677, 2004.
- [8] 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," in *4th International Conference on e-Social Science*, Jun. 2008.
- [9] M. Seropian, K. Brown, J. S. Gavilanes, and B. Driggers, "Simulation: not just a manikin," *Journal of Nursing Education*, vol. 43, no. 4, pp. 164–169, 2004.
- [10] 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.
- [11] T. D. of Health., "Delivery 21st century it support for the nhs: national strategic programme." The Stationary Office, London, Tech. Rep., Jun. 2002.
- [12] A. Triantafyllidis, V. Koutkias, I. Chouvarda, and N. Maglaveras, "An open and reconfigurable wireless sensor network for pervasive health monitoring," in *2nd International Conference on Pervasive Computing Technologies for Healthcare 2008*. IEEE press, 2008, pp. 112–115.
- [13] P. Turner, G. Milne, M. Kubitscheck, I. Penman, and S. Turner, "Implementing a wireless network of pdas in a hospital setting," *Personal Ubiquitous Comput.*, vol. 9, no. 4, pp. 209–217, 2005.
- [14] T. Manesis and N. Avouris, "Survey of position location techniques in mobile systems," in *MobileHCI '05: Proceedings of the 7th international conference on Human computer interaction with mobile devices & services*. New York, NY, USA: ACM, 2005, pp. 291–294.
- [15] J. Wang, X. An, R. Venkatesha Prasad, and I. Niemegeers, "Rf-opt: A range-free online person tracking system," in *1st International Conference on Pervasive Computing Technologies for Healthcare 2006*. IEEE press, 2006, pp. 1–5.
- [16] S. Takahashi, J. Wong, M. Miyamae, T. Terada, H. Noma, T. Toriyama, K. Kogure, and S. Nishio, "A zigbee-based sensor node for tracking people's locations," in *CASEMANS '08: Proceedings of the 2nd ACM international conference on Context-awareness for self-managing systems*. New York, NY, USA: ACM, 2008, pp. 34–38.
- [17] R. Henry, L. Matti, and S. Raimo, "Human tracking using near field imaging," in *2nd International Conference on Pervasive Computing Technologies for Healthcare 2008*. IEEE press, 2008, pp. 148–151.
- [18] H. Iin Chang, J. ben Tian, T.-T. Lai, H.-H. Chu, and P. Huang, "Spinning beacons for precise indoor localization," in *SenSys '08: Proceedings of the 6th ACM conference on Embedded network sensor systems*. New York, NY, USA: ACM, 2008, pp. 127–140.
- [19] R. Bader, M. Pinto, F. Spenrath, P. Wollmann, and F. Kargl, "Bignurse: A wireless ad hoc network for patient monitoring," in *1st International Conference on Pervasive Computing Technologies for Healthcare 2006*. IEEE press, 2006, pp. 1–4.
- [20] Jakob E. Bardram and Niels Nørskov, "A context-aware patient safety system for the operating room," in *UbiComp '08: Proceedings of the 10th international conference on Ubiquitous computing*. New York, NY, USA: ACM, 2008, pp. 272–281.
- [21] L. Coyle, S. Neely, P. Nixon, and A. Quigley, "Sensor aggregation and integration in healthcare location based services," in *First Workshop on Location Based Services for Health Care (Locare'06), 1st International Conference on Pervasive Computing Technologies for Healthcare 2006*. IEEE press, 2006, pp. 1–4.
- [22] 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," in *1st International Conference on Pervasive Computing Technologies for Healthcare 2006*. IEEE press, 2006, pp. 1–5.
- [23] J. Bravo, R. Hervás, C. Fuentes, G. Chavira, and S. W. Nava, "Tagging for nursing care," in *2nd International Conference on Pervasive Computing Technologies for Healthcare 2008*. IEEE press, 2008, pp. 305–307.
- [24] 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, no. 9-10, pp. 881–904, Nov. 2005.
- [25] T. Strang and C. Linnhoff-Popien, "A context modeling survey," in *First International Workshop on Advanced Context Modelling, Reasoning And Management, held in conjunction with Ubicomp 2004*, Sep. 2004.
- [26] I. Millard, D. D. Roure, and N. Shadbolt, "Contextually aware information delivery in pervasive computing environments," in *Lecture Notes in Computer Science, Volume 3479*, Jan. 2005, pp. 189 – 197.
- [27] O. Lassila and M. Adler, "Semantic gadgets: Ubiquitous computing meets the semantic web," in *Spinning the Semantic Web*, D. Fensel, J. A. Hendler, H. Lieberman, and W. Wahlster, Eds. MIT Press, 2003, pp. 363–376.
- [28] J. A. Muras, V. Cahill, and E. K. Stokes, "A taxonomy of pervasive healthcare systems," in *1st International Conference on Pervasive Computing Technologies for Healthcare 2006*. IEEE press, 2006, pp. 1–10.
- [29] M. Tentori and J. Favela, "Monitoring behavioral patterns in hospitals through activity-aware computing," in *2nd International Conference on Pervasive Computing Technologies for Healthcare 2008*. IEEE press, 2008, pp. 173–176.
- [30] D. Chen, J. Yang, R. Malkin, and H. D. Wactlar, "Detecting social interactions of the elderly in a nursing home environment," *ACM Trans. Multimedia Comput. Commun. Appl.*, vol. 3, no. 1, p. 6, 2007.