

A Grid Services Implementation for a Virtual Research Environment

Gary Wills, Lester Gilbert, Quintin Gee, Hugh Charles Davis, Timothy Miles-Board, David Millard, Les Carr, Wendy Hall, Simon Grange

► **To cite this version:**

Gary Wills, Lester Gilbert, Quintin Gee, Hugh Charles Davis, Timothy Miles-Board, et al.. A Grid Services Implementation for a Virtual Research Environment. 1st International ELeGI Conference on Advanced Technology for Enhanced Learning, 2005, Vico Equense (Naples)², Italy. pp.6, 2005. <hal-00190557>

HAL Id: hal-00190557

<https://telearn.archives-ouvertes.fr/hal-00190557>

Submitted on 23 Nov 2007

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

A Grid Services Implementation for a Virtual Research Environment

G Wills¹, L Gilbert², Q Gee², H C Davis², T Miles-Board¹, D Millard¹, L Carr¹, W Hall¹, and S Grange³

1. Intelligence Agent Multimedia Group

2. Learning Technologies Group

Electronics and Computer Science, University of Southampton, UK

{gbw, lg3, qg2, hcd, tmb, dem, lac, wh}@ecs.soton.ac.uk

3 Royal College of Surgeons of England, London, UK

simon@grange.me.uk

In the past, virtual learning and research environments have typically been implemented as monolithic systems. Current research focuses on creating such environments from selected Grid and Web services. In this paper we describe the Virtual Orthopaedic European University, which was created to support Higher Surgical Trainees to provide evidence for their learning contracts by carrying out experiments and publishing the results. We use this environment as a case study for the re-implementation of this system as a set of Web services, and we discuss the advantages of taking such an approach.

Virtual Research Environments, Web services, Grid Services.

1. INTRODUCTION

The current vision for the Grid focuses only upon the immediate aspects of e-science: the experiments, analyses and meetings which occur over the duration of a project. As well as these synchronic aspects, any scientific effort (and e-scientific efforts in particular) will have diachronic features. These are collaborative activities extending through time, enabling the influence of the project to carry on beyond its funded timescale and disseminating its knowledge beyond the boundaries of the original collaboration. These activities are a well-known part of the scientist's profession (publishing papers, publishing data, re-running experiments and checking others' results, comparing approaches from different projects, generalising or specialising the work of others, and teaching). The current deployment of Web technologies increases the effectiveness of this loosely-coupled collaboration and the role of the digital library is to focus the various channels (archives, publishers' Websites, and aggregation agents) into a single portal which mediates these strands of diachronic collaboration. There is also a looser coupling that exists between researchers-as-educators and their students, particularly in the context of higher education. The accepted picture of e-science could be enlarged (see Figure 1) from its current focus on experimentation and analysis to feature these processes of wider significance, since without these aspects of diachronic collaboration there would be no ongoing science and, indeed, no scientists.

Marchionini and Maurer [11], in their proposed future of digital libraries, envisaged the sharing of resources becoming an important factor in supporting teaching; this includes the ability to share raw scientific data and other datasets. Many e-science projects have collected a vast amount of data. If the next generation of scientists is to go beyond the present position it is essential that they have access to the raw data in their research and training. These early visions are slowly being realised, for example McGrath et al. have developed a system that will locate, browse and retrieve astronomy data across several databases [13], but there is still a need for those that have the technology skills, and librarians, and users, to work together to provide appropriate tools for handling, manipulating and analysing these large datasets [15]. There are projects beginning to do this; for example the Digital Library for Earth Science Education (DLESE) project allows students to explore geospatial materials and Earth data sets. Groups of students can then manually create reports using this data, and discuss them [12]. Weatherley et al. have proposed a model that will aid reviewers in reviewing complex material or a digital collection [21]. The peer review of collections and peer comment is a significant part of the dissemination process, which adds value to any collection. Lyon sees the digital library, in the context of an information grid, as consisting of a collection of resources for learning and teaching, data repositories for research purposes, or as archives of diverse cultural heritage materials [10]. While this is only a proposed scenario, Lyon recognises the need for researchers to undertake experiments, deposit raw data, and produce pre-prints using web services.

Recent advances in Web technologies allow such concepts to be realised. The Virtual Research Environment (VRE) has been implemented as a Web-based environment for supporting a critical subset of the e-science cycle (Figure 1): the collation and analysis of experimental results, the organisation of internal project discussions, and the production of appropriate outline documents depending on the requirements of conferences and journals selected for dissemination.

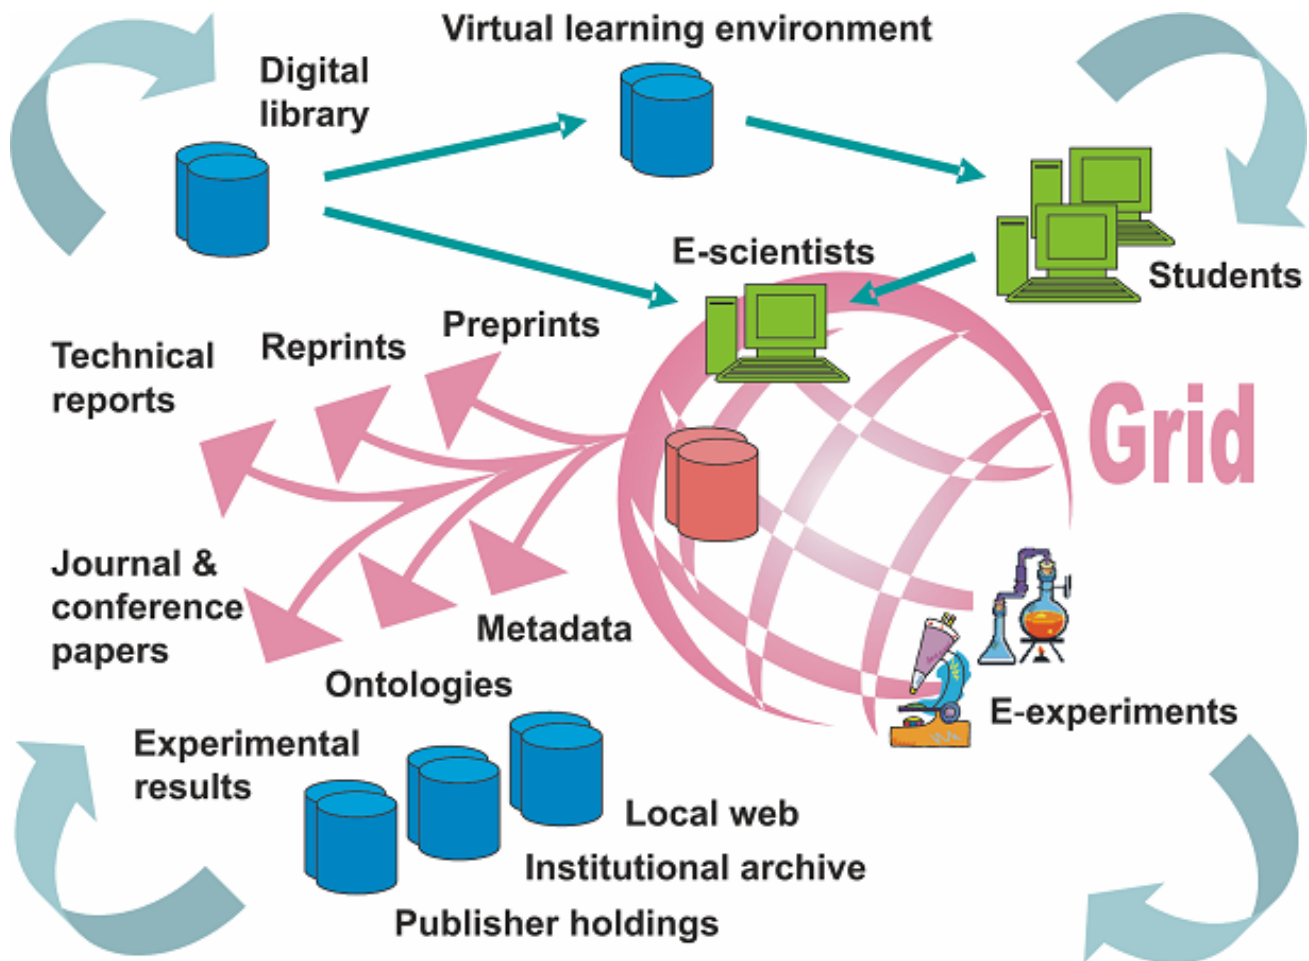


FIGURE 1 COMPLETE CYCLE OF E-SCIENCE

This paper reports on a Virtual Research Environment used in the education of surgeons.

The paper first describes a system built as a tightly integrated component of the Virtual Orthopaedic European University (VOEU). The paper then presents the current work in redeveloping the VRE as Grid/Web services. The paper concludes by discussing the potential impact of such systems and future work.

2. THE VOEU APPROACH

The Virtual Orthopaedic University provides an infrastructure for clinicians to manage the administrative, research and educational workload of the university within the digital domain.

The VOEU system included the following components, the last two being the focus of this paper.

- *Multimedia Educational Modules*, which provide the declarative (factual) base of material for the education of the users [3].
- A *Virtual Classroom* environment for exchange of views, and monitoring of progress [23].
- *Novel Modalities of Simulation* [6,7] for the emulation of surgical procedures for training and experimentation, focusing upon micro-surgery.
- A *Virtual Observatory* for the collection of data from simulation systems and the actual intra-operative data collection.
- A *Virtual Research Environment* will allow Higher Surgical Trainees (HSTs) and tutors to analyse data from existing journals, investigate hypotheses, comment on reviewed articles, and even prepare and submit articles for review. The VRE is a web-based archive of medical and technical material, which is peer-reviewed. In addition, tutors will be able to capitalise on these reviews by including the corresponding declarative and procedural knowledge in the educational modules. The implementation enables users to communicate using material tailored for their specific needs, which allows the presentation of media to be adaptive to the user experience and knowledge base. This combines declarative (factual) content with feedback from a clinical (procedural) case-based training and evaluation environment.

Combining the above disciplines within one working environment, the virtual university infrastructure [8] aims to meet the needs of clinicians by combining clinical, educational and research duties. The project builds on current thinking in educational research on pedagogy in terms of learning being situated and authentic, with learners adopting an active and constructive approach. In particular it builds on the problem-based learning literature [16], constructivism [17,18], communities of practice [22], situated learning [3,20,9] and activity theory [4].

The end users for this VRE will be the Higher Surgical Trainees, who are qualified surgeons training to be consultants. They are not computer specialists, their study is work based, and they are rarely co-located with other HSTs. During the six years of training they usually move post twelve times, and they have to keep a logbook (e-Portfolio). Therefore they typify both the average scientist trying to collaborate on a project, and a group of e-learners studying in a collaborative partner institution, i.e. they require tools that are easy to use for non-

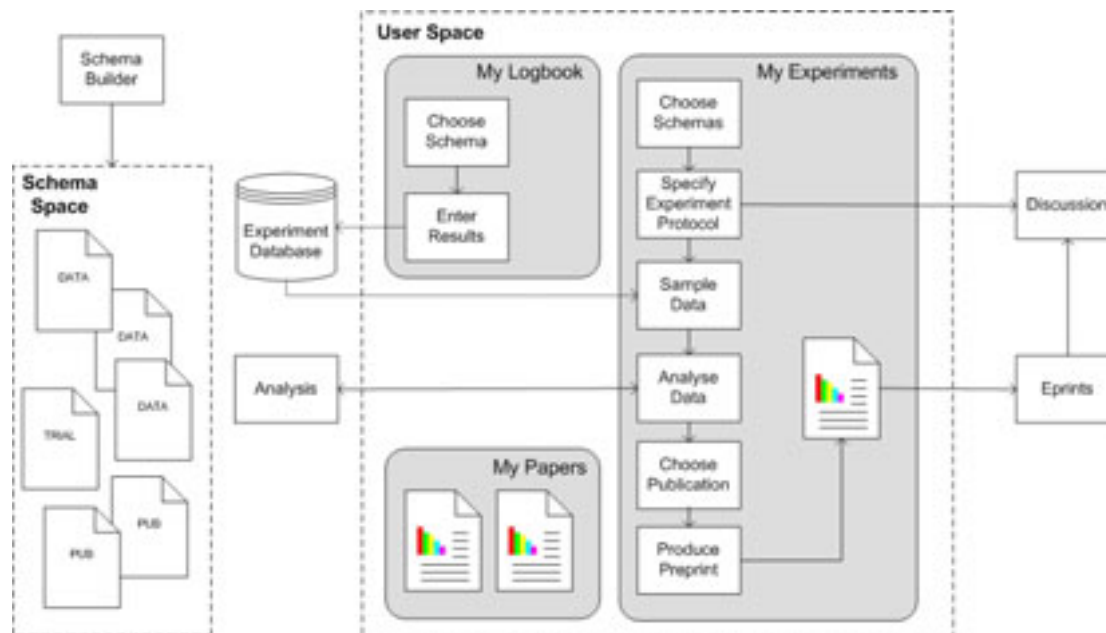


FIGURE 2 WORKFLOW IN THE VRE

computer science specialists. An objective of the project was to provide integrated computer support across the research and educational cycles because these activities are intrinsically coupled as a part of the requirements of a surgeon's Continuing Professional Development. They must undertake research and publish papers to achieve goals under the learning contracts with their Professional Colleges. The use of Grid technology provides an opportunity to improve the efficiency of both the teaching and the learning in the context of lifelong learners.

3. VIRTUAL RESEARCH ENVIRONMENT

The Virtual Research Environment in the VOEU project had two main functions: to support the educational process and to aid surgeons in preparing findings for publication. Orthopaedic surgical trials typically run for extended periods (up to 2 years), with postoperative assessment results being collected regularly. The collated results are then analysed and discussed by a team of surgeons before being disseminated to the wider orthopaedic community, see Figure 2.

3.1 SCHEMA SPACE

The schema space is the mechanism by which the Virtual Research Environment is configured to a particular e-learning community, through the formal specification of e-experimentation procedures relevant to that community. This configuration is currently achieved using three different types of schema.

- **Data schemas** describe the exact nature of the experimental data (for example, specification of variable names, types, and possible values). In VOEU there are a number of data schemas for collecting orthopaedic clinical trial data.
- **Experiment schemas** describe experimental procedures or *protocols*. For example, a protocol could specify that any e-scientist conducting an experiment of type *X* needs to record an *experiment description*, *statement of purpose* and an *outcome hypothesis*. Human-readable guidelines are also included, to help HSTs meet the requirements of the protocol and to help reviewers to ensure that the requirements have been met. As a simple example, the HST guidelines for the *experiment description* may state "summarise the content of the experiment", whereas the reviewer guidelines ask "does the experiment description adequately summarise the content of the experiment?"
- **Publication schemas** describe the required format for submitting experimental results to relevant journals/conferences (for example: *Abstract*, *Introduction*, *Background*, *Experimental Methods*, *Results*, and *Conclusions*). As with experiment schemas, human-readable guidelines are also included in publication schemas. In VOEU there are currently two publication schemas presenting the submission formats for the Journal of Bone & Joint Surgery (JBJS) and the British Medical Journal (BMJ). Where possible, the publication schema also describes any mappings between the experiment protocol (for example, specifying that the experiment *hypothesis* should appear in the *Experimental Methods* section of the article). This allows outline preprint 'previews' to be generated automatically without requiring the HST to copy and paste information between protocol and preprint.

3.2 USER SPACE

The user space is where the HST uses the schema space to organise practical data entry and collation, e-experimentation, and dissemination. The user space is further subdivided into three personalised areas *My Logbook*, *My Experiments*, and *My Papers*.

My Logbook or *e-portfolio* is an experiment logbook, in which experimental results can be entered (in accordance with a selected data schema). Logbook entries are subsequently added to the Virtual Research Environment community database, making data available (anonymously) to other community members.

My Experiments is a workspace for e-experiments, which the e-scientist works on. An e-scientist may be involved in an experiment in the capacity of *lead investigator* (initiates the experiment and acts as co-ordinator and contact for the duration of the experiment), *associate investigator* (assistant), or *reviewer* (monitors the progress of the experiment and reviews its outcomes according to guidelines). Reviewers have read-only access to the experiment protocol and set-up. When a new experiment is initiated, a discussion facility is automatically set up to facilitate and record communication between the e-scientists involved. This is also the means by which reviewers can give feedback to the practitioners.

3.3 EXAMPLE OF USE

To illustrate how a HST may use the VRE, this section outlines the process of managing e-experiments from the perspective of a fictional trainee surgeon, Sam. In Sam's view of the VRE user space, Sam is shown to be currently working on three trials, undertaking a different role in each. Sam is the co-investigator in the "charcot joints" trial, is writing a systematic review of experiments in the "rotator cuff" trial, and is also a peer reviewer of the "tear size" trial. Sam has also entered several experimental records in the personal logbook (patient details, operative procedures, and assessment results), part of Sam's e-portfolio.

Formalising Trial Protocol – To initiate a new trial, Sam first selects the experiment protocol from the available experiment schemas. The VRE then uses this schema to generate a number of data entry forms into which Sam enters specifics of the experiment. Guidelines for completing these forms are presented as "stretch text links" [14], which can be viewed/hidden as required. A tutor specifies the associate investigators and peer reviewers who will assist Sam on the trial. When created, the new trial will appear in Sam's VRE user space, and also in the user spaces of the associate investigators and peer reviewers.

Selecting a Dataset – To create a dataset for the new trial, Sam searches the data repository for appropriate data used in previous trials for suitable cases or uploads new cases from Sam's logbook. Since Sam has already specified the experiment schemas, only those cases matching this schema will be searched. Sam and associates subsequently add 42 different experimental results to the trial, which can be viewed in tabulated form for visual comparison.

Analysing the Dataset – To perform analyses on the dataset, Sam and associates choose from statistical methods offered by a distributed Analysis Engine. Using the experiment schemas and metadata from the Analysis Engine, the VRE tool is able to generate an entry form for each statistical method, which Sam can use to fine tune the analysis (specify test variables, groupings, etc.) The Analysis Engine queues the requested analysis and notifies Sam when the results are available. These results appear in Sam's VRE user space, and can be viewed.

Discussing the Results – Having obtained some significant results from the statistical analyses, Sam then decides to create a pre-print for discussion by the co-investigators on a discussion board created for the e-print. Sam, co-investigators, or a tutor, can specify who is able to see the pre-print and comment on it. For the trial, the pre-print can be taken on to publication. In this case Sam selects the JBJS publication schema, and the VRE toolkit generates a pre-print template using the information Sam entered in the trial protocol. Sam fleshes out this template, following the JBJS guidelines provided, and specifies which analysis results should be included in the preprint. After previewing the pre-print, Sam submits it; behind the scenes the VRE tool submits the pre-print and its associated metadata to the community e-prints server (where it subsequently becomes available to the members of community), and makes the paper available in the user space of Sam and associates.

4. GRID SERVICES FOR A VRE

The VOEU project is based on an integrated and tightly coupled architecture, making it difficult to expand as the user requirements change. The follow-up project to VOEU is the Collaborative Orthopaedic Research Environment (CORE) This takes as its foundations the VRE established in VOEU, but develops the modules into services within a Service Oriented Architecture (SOA) [19,0].

The aim of CORE is to provide integrated computer support across the research and education cycles. For surgeons these activities are part of their *Continuing Professional Development*. Surgeons undertake research and publish papers to achieve goals according to the learning contracts that they have with their Professional organisations.

These activities constitute a critical subset of the e-science cycle (Figure 1): the collation and analysis of experimental results, the organisation of internal project discussions, and the production of appropriate outline documents depending upon the requirements of conferences and journals selected for dissemination.

The CORE project will develop a VRE demonstrator based on a SOA for supporting all of these activities. The advantages of using a SOA are:

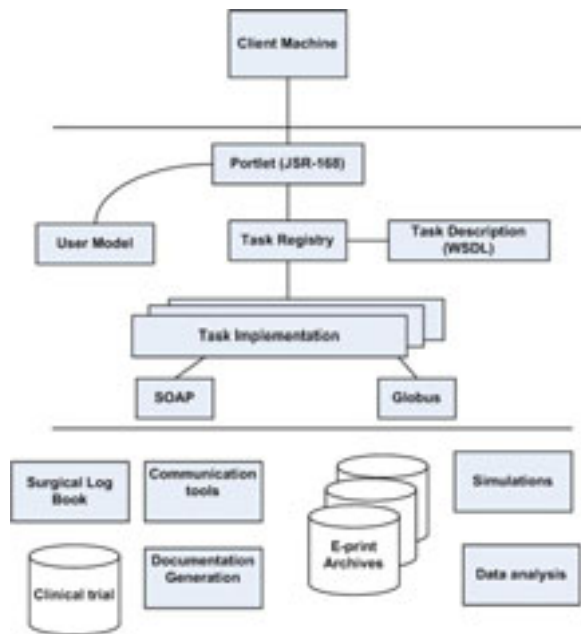


FIGURE 3 OVERVIEW OF THE PROPOSED INFRASTRUCTURE

- **Modularity:** As appropriate services are dynamically coupled, it is relatively easy to integrate new services into the framework.
- **Interoperability:** Due to standardisation of the communication and description of the services, third party services can easily be incorporated into the VRE as required.
- **Extensibility:** Due to the relative ease with which services can be incorporated into the VRE, there is less danger of technology 'lock-in'.

The CORE project is being implemented as a toolkit of generic components based on the Grid SOA infrastructure. Figure 3 illustrates the CORE Framework concept, in the context of an e-science community Web site and integrated Grid services.

Grid services were originally conceived as a method of distributing high performance computation, and it is true that the Grid will enable researchers and learners to run simulation and data analysis on a scale not possible today. However, Grid services also provide benefits in the areas of distributed knowledge and information management, offering functionality that is essential for serious e-learning applications, such as security and state awareness.

The Semantic Grid, using Semantic Web technologies to describe Grid services, promises further benefits in the areas of automation and organisation. For example, it may be possible to automate a feasibility check of possible trials by looking at the availability of appropriate resources, including lab time, experimental datasets and staff.

In an extended Grid environment, such as that described by the European e-Learning Grid infrastructure (ELeGI) project [4] it would be possible to search for and utilise alternative services and resources on the wider Grid. This is not only more robust, but offers the exciting possibility of new services being automatically incorporated into the local learning environment and new communities of remote students and teachers evolving naturally.

CORE is less ambitious in the short term, but still seeks to exploit this flexibility and dynamism. The focus of the generic toolkit components in CORE will be a Virtual Bone Biology Laboratory, incorporating the basic science disciplines of molecular and tissue biology, pharmacology, prosthetics, trials management, and the clinical disciplines concentrating on musculoskeletal applications in rheumatology and orthopaedics.

The CORE Bone Laboratory will provide a post-graduate training experience for researchers and clinicians with a background in bioengineering, initially treating musculoskeletal problems, with emphasis on training the next generation of clinicians and scientists. The CORE will need to provide Grid services for the simulation and modelling of bone and soft tissue biology, analysis of large scale experiments and the modelling of nanometric tissue units. The use of the Grid for distributed computation means that powerful analysis and modelling tools can be made available to students.

The Bone Laboratory also aims to be a testbed resource for scientists developing new technologies for information management, and constructing technologies on the nanometric scale. This will rely on the flexibility of the SOA and the secure accessibility of the Grid.

5. CONCLUSIONS AND FUTURE WORK

The development of a Virtual Orthopaedic European University described in this paper is underpinned by a pedagogical approach built on current education research. This paper has focused on a specific framework, the Virtual Research Environment, which supports the development and dissemination of documents by assisting authors in collating and analysing experimental results, organising internal project discussions, and producing papers.

The system provides a distributed architecture for institutions to manage multiple centres, advancing surgical standards through education and research. The tools are generic, and applicable across surgical and medical training. The system has been deployed in a local hospital and HSTs validated the trials.

With the advent of virtual infrastructures, there is the possibility of managing most of the administrative, research and educational workload of the university within the digital domain. This has potentially huge benefits for surgeons, providing greater access to information without the friction associated with traditional infrastructures.

The CORE project is taking the VOEU VRE and decoupling the processes by using a SOA and Web services. It will further allow the development of Grid services to support the virtual experiments on Tissue Growth in a virtual Bone Laboratory. CORE will enable HSTs to collect and analyse experimental results from their own or other people's experiments, organise internal project discussions, and produce appropriate documents. The project should have a major impact on a number of areas which include:

- being able to keep track of the research administration: trial protocol, ethical approval, and workflow as the trial progresses
- enabling access to research data from various trials and in formats that allow analysis of the data
- allowing easier meta-analysis or thematic reviews
- monitoring the effectiveness of surgical interventions
- enabling a consortium to write appropriate documents for dissemination (medical reports, journal articles, etc.)
- producing up to date learning and teaching material

6. ACKNOWLEDGEMENTS

Part of this work was funded by The European Commission under the Virtual Orthopaedic European University project, IST-1999-13079, Information Society Technology Program.

REFERENCES

- [1] Brown, J.S., Collins, A, Duguid, P. (1989) Situated learning and the culture of learning. *Educational Researcher* 18 (1): 32-42.
- [2] Carr, L., Miles-Board, T., Wills, G., Power, G., Bailey, C., Hall, W. and Grange, S. (2004) Extending the Role of Digital Library: Computer Support for Creating Articles. In *Proceedings of The Fifteenth ACM Conference on Hypertext and Hypermedia*, University of California, Santa Cruz, USA.
- [3] Conole, G., Wills, G., Carr, L., Vadcard, L., Hall, W., and Grange, S. (2003) Building a Virtual University for Orthopaedics. *EdMedia* 2003(1), 22-27. Available at <http://dl.aace.org/12697>
- [4] ELeGI: European learning Grid Infrastructure Project <http://www.elegi.org/index.php>
- [5] Engestrom, Y., Miettinen, R., Punamäki, R-L., Pea, R., Brown, J.S., Heath, C., Eds (1999) *Perspectives on activity theory. Learning in doing: social, cognitive and computational perspectives*. Cambridge, Cambridge University Press.
- [6] Grange, S., Bunker, T., Cooper, J. (1996) Networking virtual reality for shoulder arthroscopy. *British Journal of Healthcare Computing* 13[10], 26-28.
- [7] Grange, S., Bunker, T., Cooper, J. (1997) *Exeter Virtual Worlds Shoulder Arthroscopy Simulator*. London: Brunel University Press.
- [8] Hazemi, R. and Hailes, S. (1998) Reinventing the Academy. In: Wilbur S, Hazemi R, Hailes S, editors. *The Digital University*. London: Springer, 7-24.
- [9] Lave, J. and Wenger, E. (1990) *Situated Learning: Legitimate Peripheral Participation*. Cambridge, Cambridge University Press.
- [10] Lyon, L. (2002) Emerging information architectures for distributed digital libraries. In *International Conference of Digital Library—IT Opportunities and Challenges in the New Millennium*.
- [11] Marchionini, G. and Maurer, H. (1995) The roles of digital libraries in teaching and learning. *Communications of the ACM*, 38(4).
- [12] Marlino, M., Sumner, T., Fulker, D., Manduca, C., and Mogk, D. (2001) The digital library for earth system education: building community, building the library. *Communications of the ACM*, 44(5):80-81.
- [13] McGrath, R.E., Futrelle, J., Plante, R. and Guillaume, D. (1999) Digital library technology for locating and accessing scientific data. In *Proceedings of the fourth ACM conference on Digital libraries*, 188-194.
- [14] Nelson, T.H. (1987) *Computer Lib/Dream Machines*. Tempus Books.
- [15] Nevill-Manning, C. (2001) The biological digital library. *Communications of the ACM*, 44(5):41-42.
- [16] O'Dowd, J.K., Spencer, J.D. (1992) An audit of university education in trauma and orthopaedic surgery in Great Britain. *J R Soc Med* 85(4):211-213.
- [17] Papert, S. (1980) *Mindstorm*. New York, Basic Books.
- [18] Piaget, J. (1954) *The construction of reality in the child*. New York, Basic Books.
- [19] Smythe, C., Evdemon, J., Sim, S. and Thorne, S. (2004) *Basic Architectural Principles for Learning Technology Systems*. Available at <http://www.imsglobal.org/>
- [20] Suchman, L. (1988) *Plans and Situated Actions: The Problem of Human/Machine Communication*. Cambridge, Cambridge University Press.
- [21] Weatherley, J., Sumner, T., Khoo, M., Wright, M. and Hoffmann, M. (2002) Partnership reviewing: a cooperative approach for peer review of complex educational resources. In *Proceedings of the second ACM/IEEE-CS joint conference on Digital libraries*, 106-114.
- [22] Wenger, E. (1998) *Communities of practice - learning, meaning and identity*. Cambridge, Cambridge University Press.
- [23] Wills, G., Woukeu, A., Bailey, C., Ong, A., Carr, L., Conole, G., Hall, W. and Grange, S. (2004) *Ontological Driven Learning Agreements*. In *Proceedings of EDMEDIA*, 217-222.
- [24] Wilson, S., Blinco, K. and Rehak, D. (2004) *Service-Oriented Frameworks: Modelling the infrastructure for the next generation of e-Learning Systems*. Available at http://www.jisc.ac.uk/uploaded_documents/AlttilabServiceOrientedFrameworks.pdf.