

Loughborough University Institutional Repository

An e-learning framework for assessment (FREMA)

This item was submitted to Loughborough University's Institutional Repository by the/an author.

Citation: WILLS, G. ... et al, 2007. An e-learning framework for assessment (FREMA). IN: Khandia, F. (ed.). 11th CAA International Computer Assisted Assessment Conference : Proceedings of the Conference on 10th and 11th July 2007 at Loughborough University. Loughborough : Loughborough University, pp. 451-472

Additional Information:

- This is a conference paper.

Metadata Record: <https://dspace.lboro.ac.uk/2134/4596>

Version: Published

Publisher: © Loughborough University

Please cite the published version.

This item was submitted to Loughborough's Institutional Repository (<https://dspace.lboro.ac.uk/>) by the author and is made available under the following Creative Commons Licence conditions.



CC creative commons
COMMONS DEED

Attribution-NonCommercial-NoDerivs 2.5

You are free:

- to copy, distribute, display, and perform the work

Under the following conditions:

BY: **Attribution.** You must attribute the work in the manner specified by the author or licensor.

Noncommercial. You may not use this work for commercial purposes.

No Derivative Works. You may not alter, transform, or build upon this work.

- For any reuse or distribution, you must make clear to others the license terms of this work.
- Any of these conditions can be waived if you get permission from the copyright holder.

Your fair use and other rights are in no way affected by the above.

This is a human-readable summary of the [Legal Code \(the full license\)](#).

[Disclaimer](#) 

For the full text of this licence, please go to:
<http://creativecommons.org/licenses/by-nc-nd/2.5/>

AN E-LEARNING FRAMEWORK FOR ASSESSMENT (FREMA)

Gary B. Wills¹, Christopher P. Bailey¹, Hugh C. Davis¹, Lester Gilbert¹, Yvonne Howard¹, Steve Jeyes², David E. Millard¹, Joseph Price¹, Niall Sclater³, Robert Sherratt², Iain Tulloch³, Rowin Young³

An e-Learning Framework for Assessment (FREMA)

Gary B. Wills¹, Christopher P. Bailey¹, Hugh C. Davis¹, Lester Gilbert¹, Yvonne Howard¹, Steve Jeyes², David E. Millard¹, Joseph Price¹, Niall Sclater³, Robert Sherratt², Iain Tulloch³, Rowin Young³

1 University of Southampton,

2 University of Hull,

3 University of Strathclyde,

Abstract

The paper reports on the FREMA (Framework Reference Model for Assessment) project that aims at creating a Reference Model for the Assessment Domain and delivering it via a heavily interlinked Web site. Because the resulting network of resources (standards, projects, people, organisations, software, services and use cases) is so complex, we require a method of providing users with a structured navigational method that does not require them knowing at first what they might want to find. This led us to look at how overviews of e-learning domains have been handled previously, and work towards our own concept maps that plot the topology of the domain. FREMA was never intended to be a static resource and therefore we converted the original site to use a semantic Wiki, thereby allowing the Assessment Community to use the Knowledgebase to record their own projects, services and potentially new reference models.

Introduction

It is possible to characterise e-learning in terms of a number of domains that group related activities, such as managing e-portfolios or constructing learning content. The assessment domain is one of the most mature of these in terms of software and standards. Numerous commercial and academic tools are available, supporting a wide range of assessment activities, from assembling and running tests or exams to managing feedback and detecting plagiarism.

This raises problems when it comes to building new tools and creating new standards for the domain, as these must be correctly situated with existing work if they are to be successful. This problem is increasingly important in the world of Service-Oriented Architectures (SOA), as new services only become effective when they extend or support existing services. SOAs are an attempt to modularise large complex systems in such a way that they are composed of independent software components that offer services to one another through well-defined interfaces.

The service approach is ideally suited to more loosely coupled systems, where individual parts may be developed by different people or organizations. Wilson *et al.* (2004) discuss in detail the advantages of using SOA:

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

With SOAs there is a need to design complementary services that can be used together to some end. Sometimes these are known as composite services, but in larger cases could represent the infrastructure for an entire domain. Large sets of services that have been designed to work together are often known as service frameworks.

In the UK, the Joint Information Systems Committee (JISC) is financed by all the Further and Higher education funding councils and is responsible for providing advice and guidance on the use of Information and Communications Technology (ICT) for learning and teaching. Part of its strategy is the development of a SOA framework for e-learning (Oliver *et al.*, 2004, Wilson *et al.*, 2004b). JISC call this initiative simply 'e- Framework'.

The e-Framework is based on a service-oriented factoring of a set of distributed core services (Smythe *et al.*, 2004), where flexible granular functional components expose service behaviours accessible to other applications via loosely coupled standards-based interfaces. The technology used is Web Services and the intention is to extend the SOA programming model into a vast networking platform that allows the publication, deployment, and discovery of service applications on the scale of the Internet. However, the e-Framework suffers the same problem as all other service frameworks; mainly that it is difficult to coordinate the development of so many inter-related services by so many people and groups, and disseminate them to the communities that the frameworks serve.

In this paper we present our efforts to develop a Community Reference Model for the development of services within a large Service Oriented framework. Our work has been aimed at the e-Framework, and its development within the domain of e-learning in particular, but the approach is applicable to any service framework that has similar characteristics: i.e. is evolutionary rather than tightly designed, and is being driven forward by distributed, independent developers and users.

Web services in the assessment domain

In this Section we attempt to give some context to our Reference Model design, by explaining how it is based on concrete problems, faced by real users. In our case this is within the domain of e-learning, and in particular services related to assessment. We show how these real examples, or personas, can be translated into use cases that apply to service frameworks in general.

Description of the Assessment Domain

Conole and Warburton (2005) have recently presented a detailed review of the issues facing computer assisted assessment, and conclude by saying “The role of technology and how it might impact on assessment is still in its infancy and we need to develop new models for exploring this”. Reference models can be thought of as partially filling this need. The e-learning assessment domain has been classified in a number of ways in the past. For instance Bull and McKenna (2004) classify it into four broad categories based on purpose (summative, formative, diagnostic and self-assessment) backed up by a number of taxonomies. JISC themselves have developed a simple map of the assessment domain, using a single test as the connecting thread (Kassam, 2004).

There is a move in learning and teaching to use learning outcomes to define what is to be taught and therefore what is to be assessed. The skill levels defined in the learning outcomes and assessment are often set within Bloom’s (1956) taxonomy of learning objectives. Chang *et al.* (2004) have developed an assessment metadata model (taxonomy) to aid teachers in authoring examinations, which explicitly models the cognition aspect of an assessment in addition to the types of questions.

The e-learning domain is underpinned and sometimes driven by the use of technology. Sclater and Howie (2003) have defined the requirements for the ‘Ultimate’ assessment engine. In presenting these requirements they view the assessment domain from the perspective of the roles people have in the assessment process and how they interact with the resources. Whilst the scope of the assessment domain is open to interpretation, it is likely that core services will include item banks (question databases), delivery applications (that retrieve and render questions) and automatic assessment tools. If the interpretation is broad then services such as peer group formation and plagiarism detection might also be included.

As services within the domain are being developed by a wide variety of institutions for a number of purposes, it is necessary to focus the activities of the assessment community in order that they create interoperable web services and exploit their widest possible use (and re-use). What is required is not just a common repository for services, but a community wide understanding of the domain, and how independently authored services fit within it. If a reference model is to be a community focus point for service design within a framework then it is necessary for it to describe services in the context of well-defined domain processes and also relate them to existing

standards and software. This is a complex challenge due to the many existing e-learning standards, projects, and software.

Personas

We used an agile modelling technique known as 'Personas' in order to investigate the requirements of different members of the assessment community (Cooper *et al.*, 2003). To place personas in a modelling context: if actors and use cases may be considered as abstract classes, then personas and scenarios may be considered instances of those classes where an actor is characterized in detail.

The following are personas that represent the breadth of users that we might expect to interact with a Reference Model:

Persona 1, Will

'Will is an e-learning tool and web services developer in an academic institution. He is a 30-something post-graduate. He has a good knowledge of the assessment domain and has java and web services technical skills. He is developing an open-source application in the assessment domain focusing on feedback methods.

Scenario:

'I want to look up use cases and scenarios to help me design my application. This will help me to define my footprint in the assessment domain. I see there are some web services I could download but some are missing. What standards can I use when writing my own web services to ensure that I can interoperate with the web services I've chosen?'

Persona 2, Yvonne

Yvonne is a learning resource manager at a higher education institution with a background in academia and education. She is planning the institution's five year strategy for e-learning. She is responsible for ensuring that new systems meet quality assurance standards. She has a strategic grasp of the importance of e-learning but she is not an expert in the assessment domain.

Scenario:

'I want an overview of what this domain is all about. I want to know what standards are applicable in the domain to ensure that we comply with quality assurance requirements. I want to examine use cases and scenarios to understand the available footprints. I also want to know who the key players are and what the key projects are.'

Although these are just two personas from the assessment community of interest; they have widely different needs and levels of technical expertise and show the range of the spectrum of interaction. Access to resources within the Reference Model should therefore be at different levels of abstraction to match the different characteristics of interest identified.

Use Cases

From these and other Personas in the assessment context, we can generalize to three different Reference Model use cases: Domain Web Service Developers, Early Adopters, and Institutional Resource Managers. These are shown together in Figure 1

Domain Web Service Developers are actors within the domain that are in the business of creating working software services for a particular framework. They are interested in using the framework to place their own work in the context of the domain (for example, to learn domain vocabulary, and to discover where effort in the domain has been spent), in existing software and standards, and also in domain use cases and service profiles (abstract descriptions of services) that might be related to them.

Early Adopters are the primary actors within the domain that want to use emerging technology from the service framework. They are interested in orientating themselves within the domain and also in retrieving software that may help them in their own work.

Institutional Resource Managers are actors within the domain who are in charge of institutional policies and direction. They want to use the Service Oriented Reference Model (SORM) to ensure that their institution is using relevant standards in its business processes.

These actors are interested in different technical layers of the reference model. But these layers must be related in order to help the actors orientate themselves and to create an audit trail of decision making throughout the model.

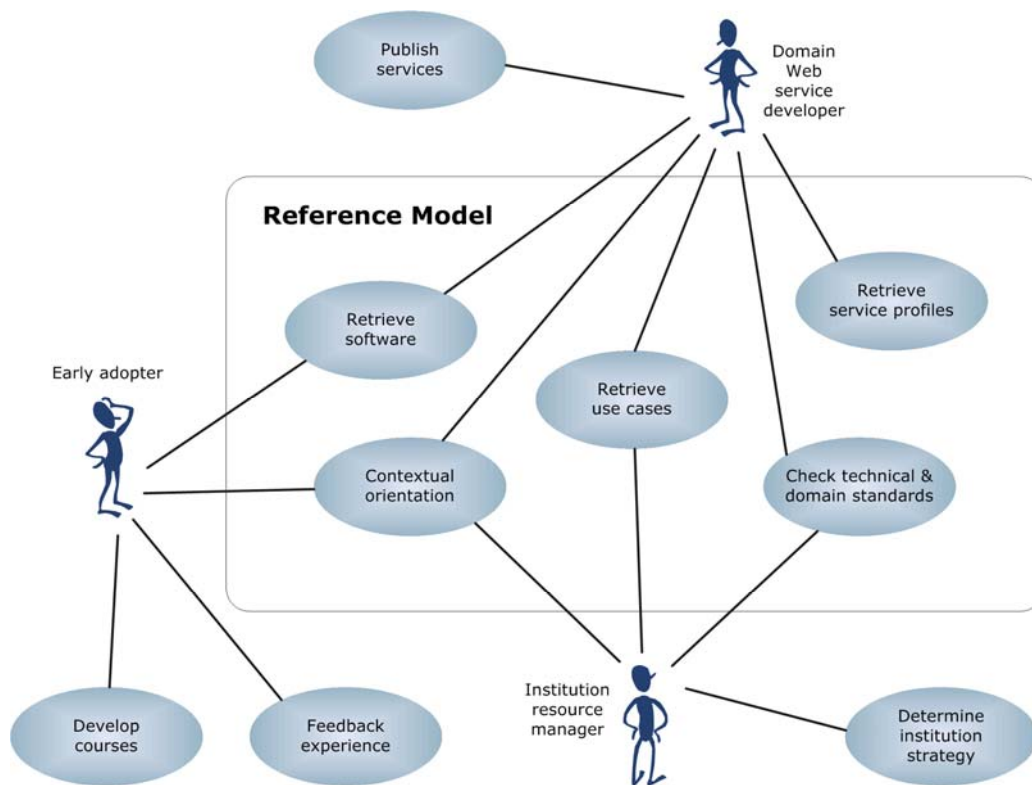


Figure 1 : Use Cases for a Reference Model

Anatomy of a reference model

If a Reference Model is to address the needs of such a broad spectrum of users it must contain a wide range of resources, such as descriptions of standards, existing software, use cases, projects, organizations, service profiles, and existing services. However, to be considered a model it must place these in relation with one another, so that it describes the real-world situation. To be an effective model it is necessary for users to be able to understand the model and draw more advantage from it than by examining the real world that is being modelled.

To enable this we have conceptualized a Service Oriented Reference Model as a number of layers, and defined the relationship between each layer. Each layer contains a different set of resources. We have chosen to model these resources ontologically so that the schema of relationships can be shared and understood across the domain. It has also allowed us to create a more dynamic model, which has an extensible set of relation types.

In this section we explain the purpose and content of each layer and describe how, for our Community Reference Model, we exposed the semantic web of resources through a dynamic and heavily interlinked Web site, described at the top level via complementary concept maps.

Layered Architecture

A Service Oriented Reference Model can be thought of as a series of layers. For tightly constrained domains, it may be possible to define a vertical slice through the layers, such that each layer exactly maps onto its vertical neighbours. For broader domains where each layer is smaller in scope but more concrete than the one below it, a Community Reference Model approach is more appropriate.

It is imagined that as a community uses and further develops a Reference Model its higher layers will cover more and more of the lower. Figure 2 shows the layers of the Community Reference Model and the processes that lie between them:

Domain Definition: This layer is an overview of the domain that the reference model covers. The definition contains instances from the ontology of domain resources (such as standards, people, and projects) and also the ontological relationships between them. Each of these instances and relationships have narrative descriptions associated with them. In addition each instance is placed in one or more concept networks, so that they may be found by users graphically browsing the domain.

Identifying Common Usage Patterns: This is the process of scoping the domain into a manageable subset. Manageable may mean areas that lie unarguably within the domain (according to the views of domain experts), or it may be a reflection of the resources available to create the higher level, more concrete layers of the model. In either case the patterns should include all key activities.

Use Cases: This layer formalises the usage patterns into use cases: formal descriptions of user activity in both diagrammatic and narrative form. The Use Cases become new resources, linked to each other and the rest of the Domain Definition through new ontological relationships and narrative descriptions.

Gap Analysis: This is the process of mapping the Use Cases to atomic services within a given framework and identifying which ones are missing a formal definition. Not all use cases will necessarily be mapped, although core activities should be covered.

Service Profiles: This layer contains the descriptions of those services identified in the gap analysis. *Service Profiles* are abstract descriptions of a service that may be fulfilled by several different *Service Implementations* that potentially expose different concrete interfaces. We therefore needed to model Service Profiles in a high level way that does not prescribe a data model or dictate explicit methods. To do this we created Service Resource Cards (SRCs), based on an existing agile technique called Class Responsibilities/Collaborations first described by Beck and Cunningham (1989). Our SRC models the capability of a service to realise a specific use case. The responsibilities of a service describe at a high level the purpose of a service: what it is for, what it does, and what it can provide to other

components. Collaborations with other services indicate where a service might consume another service to fulfil its own specific use case.

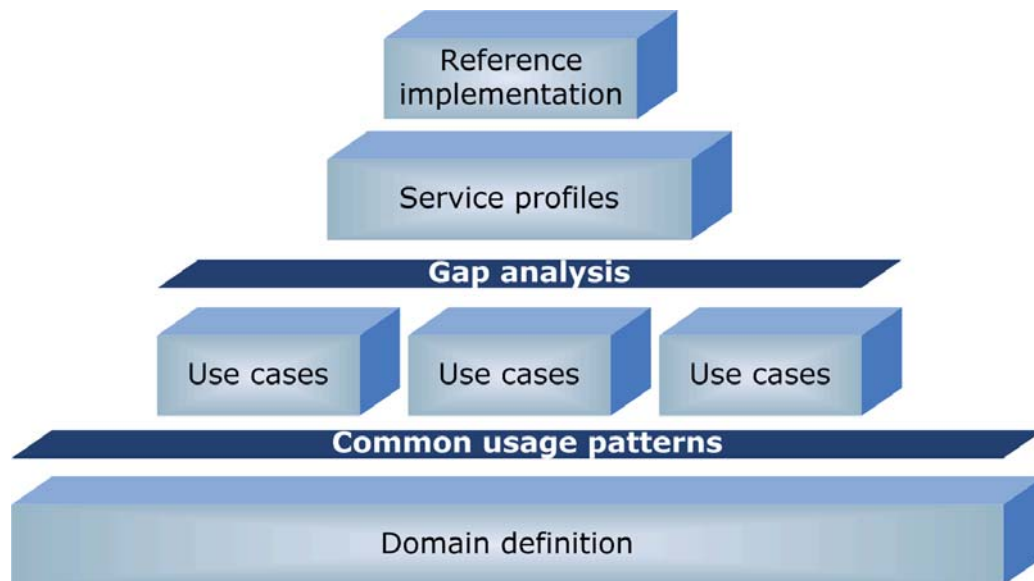


Figure 2: The Abstract Layers of a SORM

The Service Profiles and Service Implementations become resources in the model and are interlinked in the same way as other resources. In some cases the functionality of the identified service will be encompassed by existing software systems in the Domain Definition layer, in which case they should be linked together using the appropriate ontological relationship.

Reference Implementation: The most concrete layer is an actual reference implementation of the service profiles. Not all services will necessarily be implemented, and some may be wrappers around existing software. The implementations are not intended as definitive enterprise level pieces of code, but as exemplars that validate the service profiles and demonstrate any interoperability (although in open source cases they may also act as an actual software resource). These implementations become the final resources in the Reference Model, and are linked down through the profiles and use cases to the domain definition. This chain of links forms an audit trail that describes exactly why and how the software was conceived. The implementations may also be linked more directly (for example, they may draw on standards, or use software systems that have been described in the domain definition).

Each layer of the reference model is useable in its own right to achieve the use cases from Figure 1:

- *Domain Definition:* This might be used to develop a context for one's own work, to understand how existing work fits together, and to identify standards and locate experts.

- *Use Cases:* These can be used to help understand usage patterns within the domain and to help developers create new Service Profiles and thus Services.
- *Service Profiles:* Developers that wish to build new services that work within the domain framework will need to use the service profiles to ensure interoperability. They might also wish to create alternative versions of existing services, either to improve on the existing implementations, or for commercial reasons.
- *Reference Implementation:* Finally the actual software implementations are available to those developers that wish to build on some, or all, of the developed services.

We can demonstrate how the reference model supports service discovery and evolution from the perspective of one of our actors, instantiated as a persona, Will, (the domain developer) and follow his activities revealed in the Community Reference Model as he enacts his scenario.

Will's goal is to create and publish new services. He will use the domain definition layer to understand the scope of the domain and follow links into the use case layer to locate where his own expertise lies in the context of use cases and scenarios.

He will use the gap analysis to identify where competition for service delivery is high (many links from use cases to service implementations) and also where there are opportunities for him to create innovative new services (no or few links to service implementations).

For the opportunities he has discovered, Will is able to view service profiles where they have been specified in the service profile layer. Finally, he can follow links from the service profile to a reference implementation for his new service in the implementation layer which shows how it should interoperate with other services in the framework. Will is able to follow the chain of links back to the domain layer to check what domain and technical standards support the service profiles he is interested in. He may also follow the links back to the service profile layer to locate some existing services that he can re-use in the architecture of his new service.

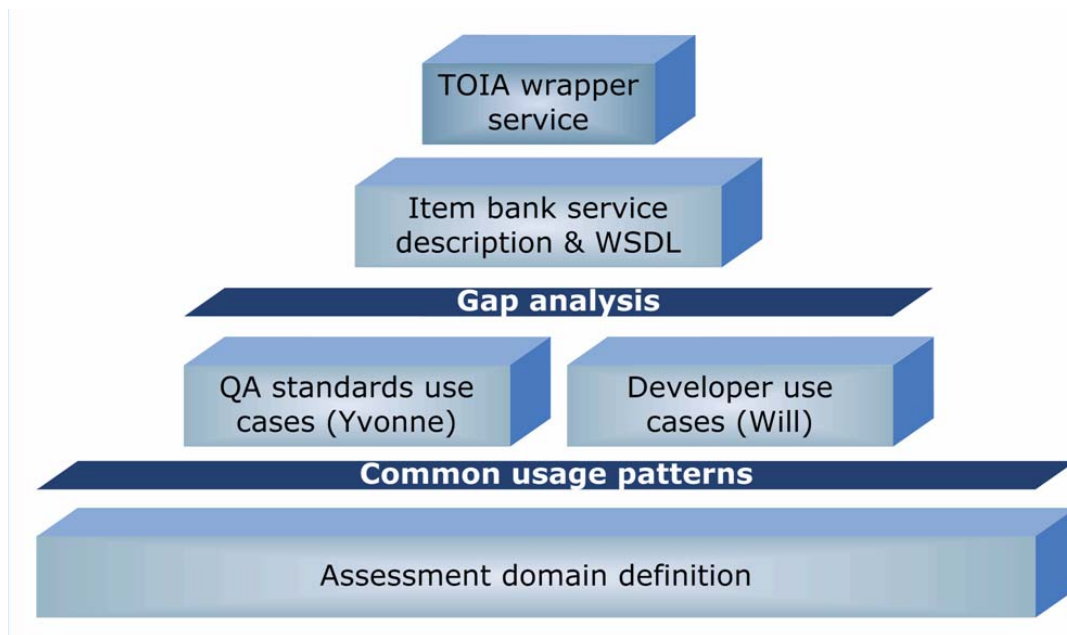


Figure 3 : Layers of a Community Reference Model from the Assessment Domain

Figure 3 shows a possible form for our instantiated SORM with our example personas (Yvonne: the Quality Assurance manager interested in standards support, and Will: the eFramework developer). From these uses cases a gap analysis will show which core services need to be profiled, in this case an Item Bank service to support Will’s development. Finally there are reference implementations of these services. In this case there is one, supporting the item bank service and providing a wrapper around TOIA, an existing item bank system.

The full version of the FREMA Community Reference Model is much more complex than this simple diagram is able to convey, with many use cases that map onto many services. It is likely that the FREMA Community Reference Model will both create new service implementations and wrap existing systems (sometimes to reveal more than one service interface).

Structuring the Reference Model

Since the Community Reference Model is designed to be a community resource it is important that it is accessible to all its users and reveals itself at many levels to them. Because of this required flexibility it is impossible to create a static representation of the resources, and instead we have opted for an ontologically modelled set of resources that are combined dynamically at the time of viewing, allowing different users to see the full domain, from base definitions to final service implementations, from a variety of views.

An ontology is simply the collection of classes and relations that are permissible for any given domain (it is called an ontology since it restricts and defines which parts of the world may be understood by entities conforming to it). The advantage of ontological modelling over database schemas is that it enforces a finely grained, and thus flexible and extensible, set of relationships.

It also means that the resources in the Reference Model could be described on the Semantic Web, which would enable interoperability between different Reference Models and reasoning about the described resources.

Figure 4 shows a graphical representation of the ontology that we have used to model our SORM. For simplicity we have not shown the attributes of each class, but have instead concentrated on the relationships between classes.

While the aim of this rich semantic storage layer is to enable users to come into the reference model from many different perspectives, there is a problem in that some users will not know where to find the resources that they are interested in within the model. To this end our ontology also includes a Concept class. Every instance in the reference model can have a conceptual relationship to one or more concept instances. We can then use graphical concept maps to help users orientate themselves and find resources. By investigating various alternative concept maps we hope that we have provided a non-expert means of navigation. We are thus using several kinds of information structure to encode and present the Reference Model.

This ontology is instantiated into a semantic network of resource instances and specific relationships. This is analogous with a *Topic Map* (although topic maps are normally presented visually, and our semantic graph is instead used to generate interlinked web pages). Within the ontology we also model concepts. The semantic graph of these concepts is a *Concept Map* which we do reveal graphically (concept maps can be considered a simple form of topic map that are intended specifically for human viewing and clarity).

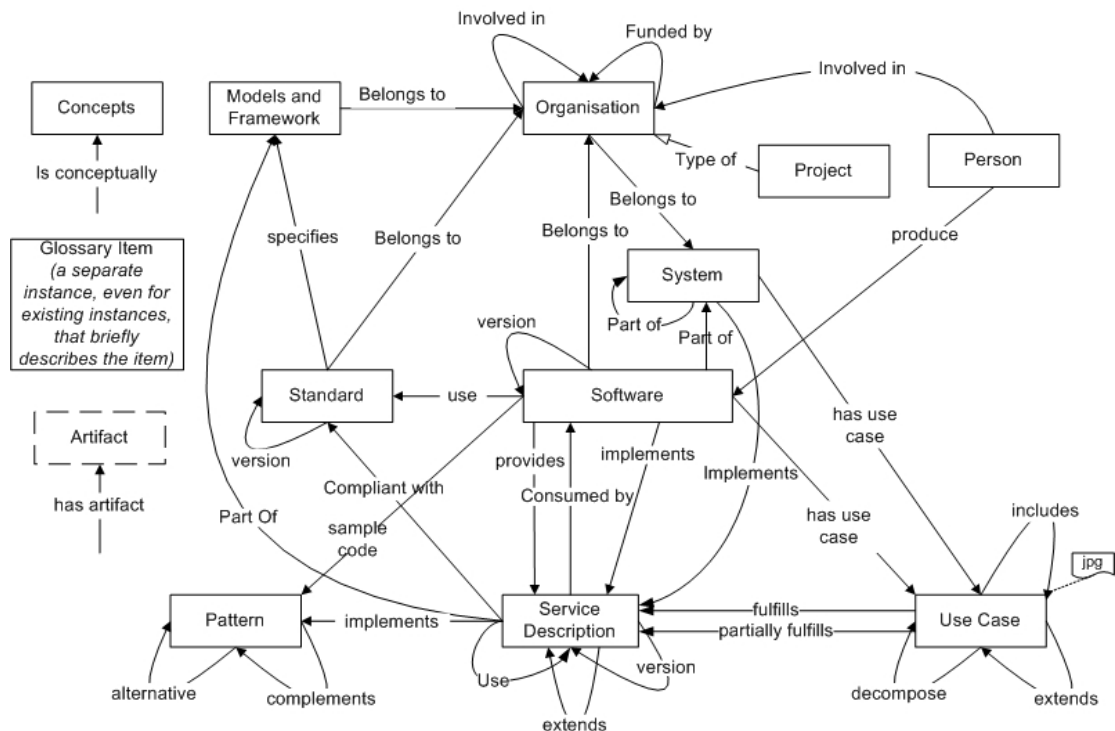


Figure 4: The Reference Model Ontology

Navigating the Reference Model

For the FREMA Assessment Reference Model we wanted to use a structure for the domain that could be used by human users of the model to orientate themselves and navigate around the resources. While the underlying resource types are modelled using an ontology, we did not want to expose users to this complexity and we also wanted to avoid the rigidity of a taxonomy. So we chose to create concept maps that described the domain in familiar terms, but which were not explicitly typed or restricted. Every resource in the reference model is associated with at least one concept. Users of the reference model can explore the maps and click through the concepts to the resources that are associated below.

The FREMA concept maps evolved over a period of several months through a series of consultation exercises. We visited a number of community events within the UK and interviewed a number of practitioners with the aim of extracting common terms and perspectives. These informed an initial, informal set of terms and relationships, which we then took back to the community for validation.

Our initial efforts at creating an overview map were a little too complex to be universally understood. We therefore broke down this map over several workshops in an effort to extract a simplified view of the domain. The result was a map of resource types that are considered important within the assessment domain, and a map of the common processes. We refer to the

resource types version as the Noun Map (Figure 5) and the processes version as the Verb Map (Figure 6).

The Noun Map draws heavily from the Ultimate Assessment Engine in that it contains stakeholders and roles (Sclater and Howie, 2003), however because it does not show workflow it does not connect these, or associate them with the types of resources they manipulate. The Noun Map is intended to allow users who deal with specific types of resources to find those resources in the map and discover what other resource types might be relevant.

The verb map shows what people do, but it does not group these activities according to any stakeholders, or relate them to any notion of resource types. There is an implicit clockwise order that follows a common view of how assessments are constructed and executed. The Verb Map is intended to allow people who are interested in a particular activity to find that process, and thus the resources underneath, and also find what other processes are related.

Semantic Wiki

The World Wide Web is the most popular hypertext system, yet it suffers a number of problems when evaluated alongside other hypertext systems. In particular, it has a very clear separation of author and reader, which means that web users cannot change the pages they are viewing. Creating web pages requires specialist skills, and collaborative authoring of a Web site is difficult. One general solution is a WikiWikiWeb (Wiki for short), a type of Web server (or application running on a traditional Web server) that allows any reader of its pages to alter those pages, or create new ones, by using simple web forms (Leuf and Cunningham, 2001). Crucially this allows non-specialist users to contribute to the hypertext.

Semantic Wikis (Völkel, *et al.* 2001) are an attempt to use the Wiki concept to make semantics accessible to ordinary users in the same way that ordinary Wikis make hypertext accessible. In Semantic Wikis users are able to type pages and links, forming a semantic network that can be queried. Semantic Wikis make semantics accessible because they are inherently freeform in nature and are non-restrictive, allowing the creation of *semantics-on-demand*, without a complex ontological design process beforehand.

Rather than construct our own Semantic Wiki system, we wanted to exploit an existing system that had typed links, nodes, and first-class types. We looked at a number of existing Semantic Wikis, including IkeWiki (Schaffert *et al.*, 2006). Kaukola (Kiesel, 2006), WikiSar (Aumueller and Auer, 2005) and Semantic MediaWiki (SMW) (Völkel, *et al.*, 2001).

In the end we chose Semantic Media Wiki (SMW) as it is relatively mature (as it is based on MediaWiki), has a large user base, offers a number of Wiki features (such as image and user management) and fits our key criteria.

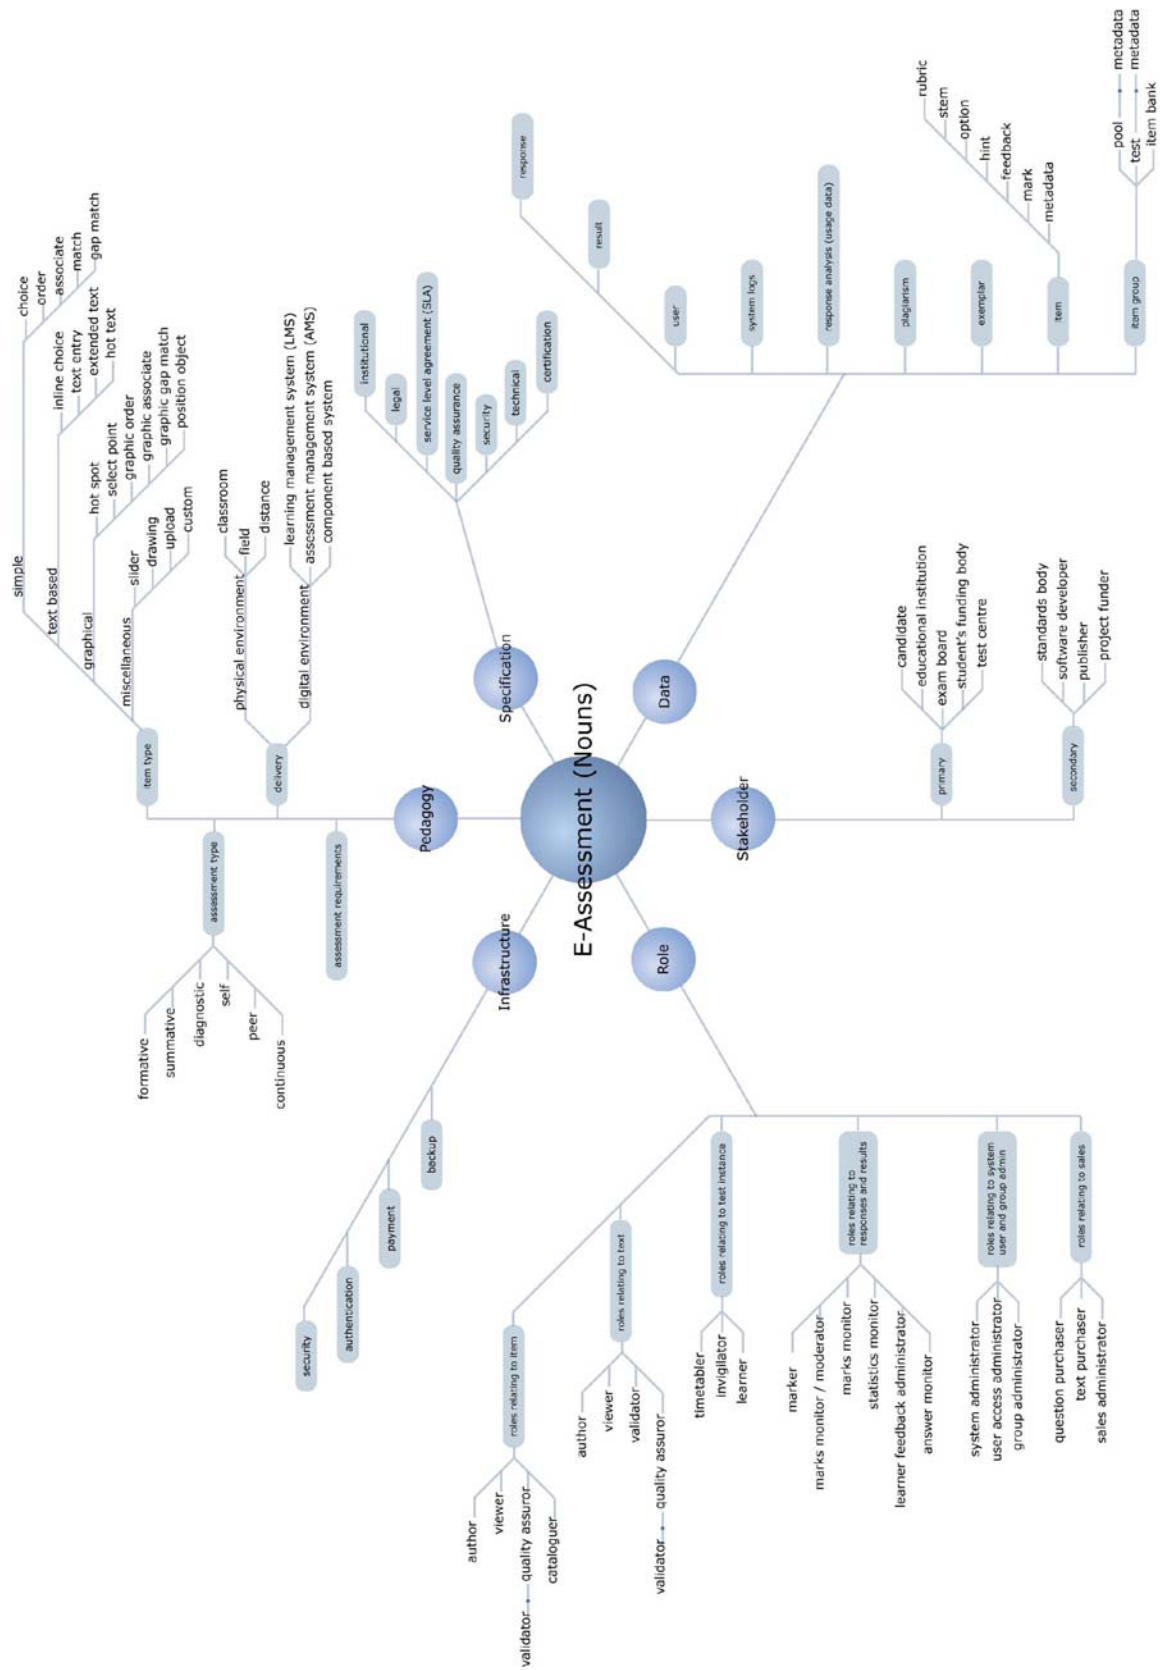


Figure 5: FREMA Noun Concept Map

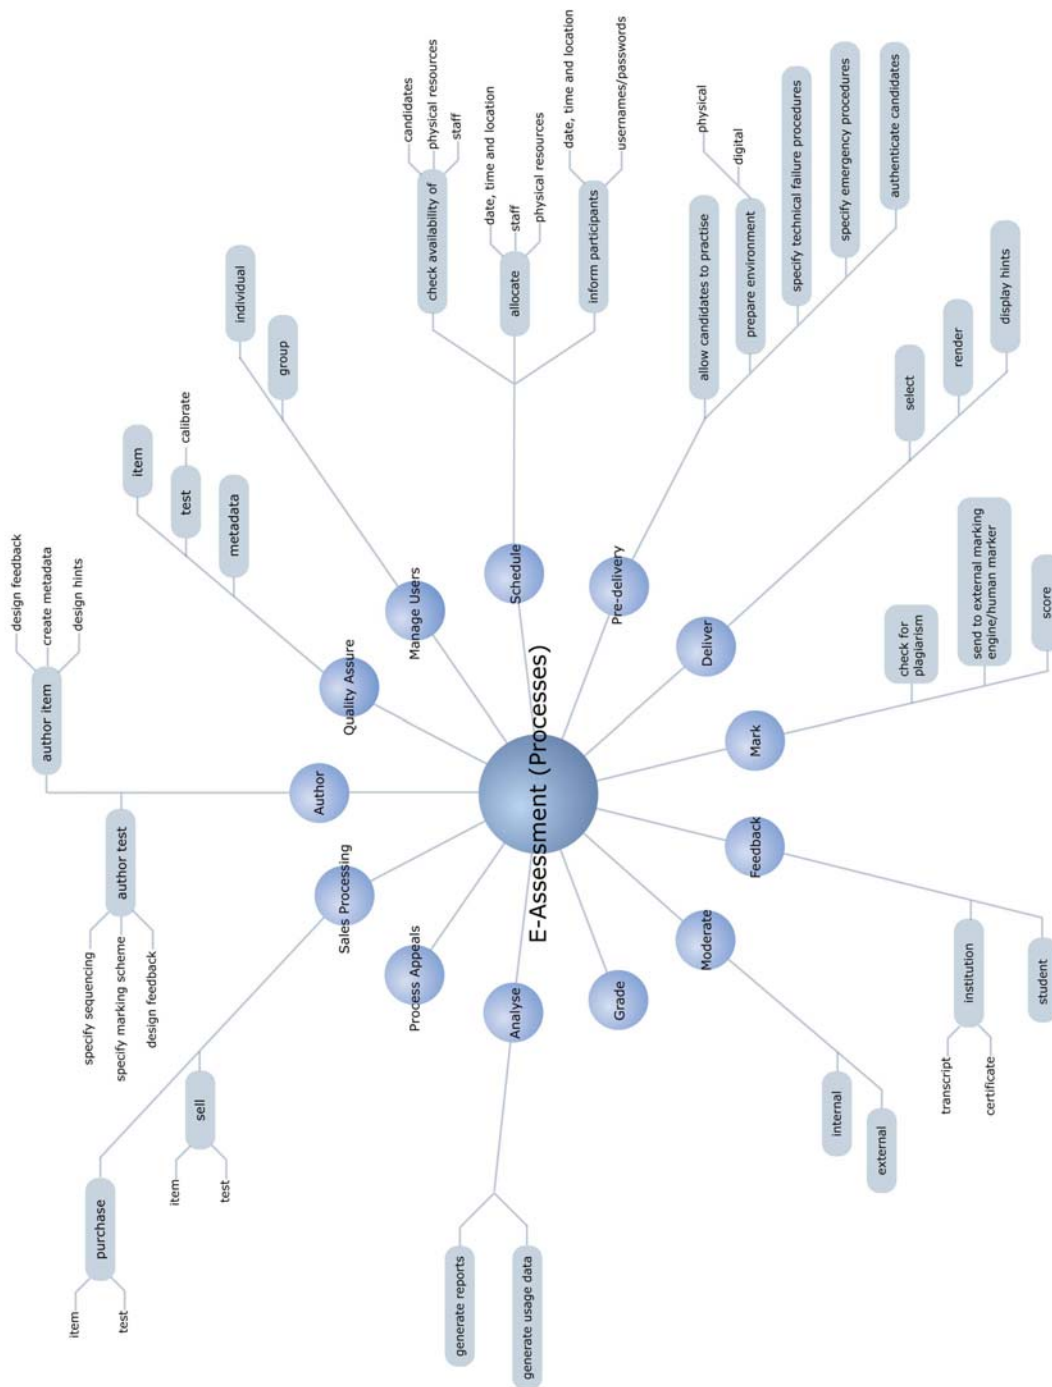


Figure 6: Verb Map

Converting the FREMA knowledgebase into a Semantic Wiki was not the trivial process that we hoped for. However, by sacrificing some of the functionality of the original site, and writing limited SMW extensions, it was possible to replicate most of the original website while gaining all the advantages of using a Wiki: open editing, administration, discussion, file management, etc. Figure 7 shows a list of resources in the FREMA SMW and Figure 8 shows a resource page.



Figure 7 FREMA Semantic Wiki List of Organisations



Figure 8 FREMA Semantic Wiki Resource Page

Conclusion

In any complex domain where disparate people are required to work together to achieve a common aim it can be helpful to provide a mechanism for overview. Once people delve into more specific areas then this becomes a need to support navigation and orientation. We believe that while a complex domain itself may be best modeled with an ontology or a topic map, the

overview is best presented as a concept map. In this paper we have presented two complementary concept maps of the assessment domain that we have used with the FREMA reference model for Assessment. On the FREMA website we use these concepts maps to help users discover resources, orientate themselves within the domain, and discover new areas that might also be of interest.

Moving from a knowledgebase-driven web site to a Semantic Wiki means a change in thinking. One must release control of the structuring ontology, and place one's faith in the wisdom of the user community. But it is also liberating, and the potential advantages are many: a familiar editing paradigm, co-ownership of content, and evolution rather than stagnation of structures and terms.

References

Aumueller, D., and Auer, S. (2005). Towards a semantic wiki experience – desktop integration and interactivity, in WikSAR, *Proc. of 1st Workshop on The Semantic Desktop*, Galway, Ireland.

Beck, K. and Cunningham, W. (1989). A laboratory for teaching object oriented thinking, *ACM SIGPLAN Notices*, 24(10), 1-6.

Bloom, B.S. (1956). *Taxonomy of Educational Objectives*, Longman.

Bull, J., and McKenna, C. (2004). *Blueprint for Computer Assisted Assessment*, Routledge Falmer.

Chang, W.-C., Hsu, H.-H., Smith, T., and Wang, C.-C. (2004). Enhancing SCORM metadata for assessment authoring in e-Learning, *Journal of Computer Assisted Learning*, 20, 305.

Cooper, A. and Reimann, R. (2003). *About Face 2.0: The Essentials of Interaction Design*, John Wiley & Sons.

Conole, G. and Warburton, B. (2005). A review of computer-assisted assessment, *ALT-J Research in Learning Technology*, 13, 17-31.

Kassam, S. (2004). A summary of the assessment domain, in *Proceedings JISC/CETIS conference*, Oxford, UK.

Kiesel, M. (2006). Kaukolu - Hub of the semantic corporate intranet, *Workshop: From Wiki to Semantics*, ESWC.

Leuf, B. and Cunningham, W. (2001). *The Wiki way: quick collaboration on the Web*, Addison-Wesley.

Olivier, B. (2005). *The eFramework: an Overview*, JISC.

Olivier, B., Roberts, T., and Blinco, K. (2005). The e-Framework for Education and Research: An Overview, retrieved from www.e-framework.org, July 2005.

Reich, S., Will, U. K., Nuernberg, P. J., Davis, H. C., Groenbaek, K., Anderson, K. M., Millard, D. E. and Haake, J. M. (2000). Addressing Interoperability in Open Hypermedia: the Design of the Open Hypermedia Protocol, *New Review of Hypermedia and Multimedia*, 5, 207-248.

Schaffert, S., Bischof, D., Buerger, T., Gruber, A., Hilzensauer, W., and Schaffert, S. (2006). Learning with semantic wikis, *SemWiki 06*, Budva, Montenegro.

Sclater, N. and Howie, K. (2003). User requirements of the “ultimate” online assessment engine, *Computers & Education*, 40, 285–306

Smythe, C., Evdemon, J., Sim, S., and Thorne, S. (2004). Basic architectural principles for learning technology systems, IMS Global Learning Consortium. .

Wilson, S., Blinco, K., and Rehak, D. (2004). *An e-Learning Framework: A Summary*, JISC.

Wilson, S., Blinco, K., and Rehak, D. (2004b). *Service-Oriented Frameworks: Modelling the infrastructure for the next generation of e-Learning Systems*, JISC.

Völkel, M., Krötzsch, M., Vrandečić, D., Haller, H., and Studer, R. (2006). Semantic Wikipedia, in *Proceedings of WWW 06*, Edinburgh, 585-595.

