Research community knowledge portals

Vladan Devedžić

Department of Information Systems and Technologies
FON – School of Business Administration, University of Belgrade
POB 52, Jove Ilića 154, 11000 Belgrade, Serbia and Montenegro
E-mail: devedzic@fon.bg.ac.yu

Abstract: This paper discusses principles and kinds of research community knowledge portals. It also presents a case study – experiences of a research group in establishing a knowledge portal, the kinds of practical efforts needed in such a project, the major hurdles and pitfalls on the way, and the lessons learned. Although the case study discusses a research community knowledge portal in a specific research field – artificial intelligence – the experiences it describes are relevant for developing portals to other knowledge areas as well.

Keywords: knowledge portals; portal technology; online communities; research community portals; semantic interoperability.

Reference to this paper should be made as follows: Devedžić, V. (2005) 'Research community knowledge portals', *Int. J. Knowledge and Learning*, Vol. 1, Nos. 1/2, pp.96–112.

Biographical notes: Vladan Devedžić is an Associate Professor of Computer Science at the Department of Information Systems, FON – School of Business Administration, University of Belgrade, Serbia and Montenegro. He has received all his degrees from the School of Electrical Engineering, University of Belgrade, Serbia and Montenegro (BS, 1982; MS, 1988; PhD, 1993). His main research interests include software engineering, intelligent systems, knowledge representation, ontologies, intelligent reasoning, and applications of artificial intelligence techniques to education and medicine. So far, he has authored and co-authored about 230 research papers published in international and national journals and conferences. His major long-term professional goal is to bring close together ideas from the broad fields of intelligent systems and software engineering. He has developed several practical intelligent systems and tools, and actively participates to several ongoing projects, both nationwide and international.

1 Introduction

A web portal can be defined as a website or service that provide a single point of access to aggregated information, and usually offers a broad array of resources and services, such as e-mail, forums, search engines, personalisation, and online shopping malls (Bailey and Treloar, 2001; Mitchell, 2004; Strauss, 2000). In other words, a portal is a gateway to web access, a hub from which users can locate all the web content they commonly need. Portals can focus on a specific topic (or on a specific set of related topics), or can be quite general, such as the one shown in Figure 1.

Figure 1 The aboutportals.com web portal

Massa (2002) stresses some practical aspects of web portals from the user's perspective. A portal is a system of integrated programmes (typically things like free e-mail, fora, classified ads, and a search engine) that not only makes it easier for a user to find information – it also provides a consistent experience and a sense of community to him/her. Moreover, it represents a good starting point for his/her journey as opposed to trying to go out and spend a lot of time to find all the pieces of the puzzles for himself/herself.

The primary goal of most portals is ease-of-use (Mitchell, 2004):

> "Besides having a single point of access – a virtual front door – portals generally try to provide a rich navigation structure. Portals using web pages for their user interface will, for instance, often include numerous hyperlinks on the front page."

There are different categorisations of portals. For example, one can talk about:

- Internet portals – portals targeting general audiences on the web; internet gateways or libraries; they typically offer general-purpose services such as featured content, numerous hyperlinks, search capability, stock quotes, and customisation based on user locale.

- Intranet portals (also called Enterprise portals) – organisation-private websites, whose tends to be restricted to the information most relevant to the organisation. An important issue with intranet portals is that of information security and support for groups and group administration (adding and removing groups and members, maintaining access rights, auditing, and so on). For example, access to certain intranet documents may be restricted to certain individuals or project teams with a "need to know" (Mitchell, 2004). Unlike internet portals that are usually produced by third parties, much of the content of an Enterprise portal is generated by the community users, hence there is few commercial contents on such portals.
- Horizontal portals (Horizontal Enterprise Portals, or HEPs) – a kind of MegaPortals, covering a wide range of topic hierarchies (e.g., Excite, Yahoo, AltaVista, Netscape's Net Center, AOL.com, Infoseek…).
- Vertical (VEPs or Vertical Enterprise Portals or Vortals) – pertaining to specialised online communities of practice, focusing on topics of interest for such communities (for example, CNET.com (shopping mall), animalhouse.com (college), MP3.com (music), pets.com (pets), women.com (women's issues), sportsline.com, etc.).

A growing trend in the practice of web portal deployment is using portals as the solution to an organisation's knowledge management and knowledge access needs. Such *knowledge portals* make available to knowledge workers all the pieces of information they need to access, and all the knowledge applications they need to use. Even if an organisation's intranet contains much of the resources that less experienced or novice knowledge workers need, the time and effort required by new hires to locate them is often costly and training and support resources are limited. Knowledge portals provide means to capture and share the expertise of more experienced knowledge workers and integrate it into a single source for learning, performance support, and ongoing knowledge sharing needs of novices.

A *community web portal* (or just *community portal*) is an internet platform for communication and information services of interest to a particular web community and possibly to a more general public as well (Staab et al., 2000). Conceptually, community portals can be seen both as a kind of knowledge portals and as a kind of VEPs. The differences are minor, and are related to the breadth of focus and the degree of restrictions for users. Like all the other VEPs, community portals weave loose pieces of information into a coherent presentation adequate for sharing knowledge with the user. However, the focus of interest on community web portals is usually slightly narrower. Community portals are also like other knowledge portals, but are generally more open to public since they may not be related to a specific enterprise or institution and hence information access restrictions can be somewhat relaxed.

Research portals are community web portals that provide research groups, university boards, managers of higher education institutions, and other decision makers in science and government with the latest information on important topics in research and public policy. They can also ensure for quick and efficient channels of communication by supporting and sustaining personal relationships and contacts that arise during the yearly conferences and to build a growing database of information for conference participants. Thus, research portals can serve as ongoing virtual conferences and platforms for dialogue among the scientific community, government, and the general public.

This paper discusses important issues related to portals and portal technology, and describes experiences of a research group who undertook an effort of setting up, popularising, and maintaining a research portal. The research domain and interests of the group is Artificial Intelligence (AI). However, in this paper AI is used only as an example domain – an illustration of building a research portal for a specific community. The focus of the paper remains on the prerequisites, the development and maintenance processes, and usability issues of the portal technology when deployed to support activities of a research community. The ultimate goal of the paper is to share experiences of the entire process of building a research portal, from realising the needs for it, to developing the initial version and modifying it subsequently, to what it takes for a research community knowledge portal to succeed.

The paper is organised as follows. Section 2 briefly covers some important concepts from web portal technology. Section 3 talks more about knowledge, learning, and research portals, providing the rationale for developing them. Section 4 introduces the GOOD OLD AI research portal and describes the ideas driving its setup and use. In Section 5, important conceptual issues are covered pertaining to design and structure of research community portals. Section 6 discusses ontological issues, commenting on the frequently asked question "What are the contents/elements that might be on a portal?" Sections 7 thru nine present various practical aspects and lessons learned during the development of the GOOD OLD AI portal.

2 Web portal concepts and technology

Portals are largely based on existing web application technology, such as web servers and Java 2 Platform Enterprise Edition (Wege, 2002). There are now a wide and increasing number of portal development toolkits that simplify the process of building portals. They range from expensive, document-focused repositories like Lotus Domino, through cheap out-of-the-box solutions like Zope, through to Java-based solutions providing rich functionality from a sophisticated set of servlets which can be added to (Bailey and Treloar, 2001). Many such toolsets are open source portal frameworks; for example, see PHP nuke – Open Source Web Portal System (http://phpnuke.org), or DotNetNuke (http://www.dotnetnuke.com/).

In order to build an Enterprise portal, one needs four components (Bailey and Treloar, 2001):

1 authentication (for customisation and information access control)

2 directory services (account information, user's role information, and departmental affiliations; ensuring that applications such as mail, news or web-access can quickly check access-control and other user profile information without introducing significant lookup latencies)

3 database (to manage information components and access to these)

4 content management (to provide the rich range of content the users want, accurate, up-to-date, and from multiple sources).

These four components are generally provided by *portal servers*, specific applications that provide portal implementers with a common set of services such as (Wege, 2002):

- content syndication – gathering content from different sources (generally, the syndication service talks to every attached back-end system via the appropriate protocol)
- content aggregation – preparing content from different sources for different users
- customisation – offering specific content tailored to the user's needs, i.e., content, layout, presentation, and profile information customisation
- portal administration and user management – defining user groups, interaction channels, and authorisation information.

Customised information is shown to users by means of *portlets* – pieces of code that run on the portal server and provide content to be embedded into portal pages. The *portal engine* is responsible for receiving requests for information, invoking appropriate portlets, aggregating the possibly multiple portlet response, and returning it as the response to the user.

Typically, many portlets are local to the portal server – many are domain specific, and using local resources is more efficient. However, there is also a need for deploying portlets in one portal server that were developed in a different one. Until recently, no standards for portlets existed, and thus consuming such remote portlets in a generic way has been impossible. The things have largely improved in late 2003, when the Organization for the Advancement of Structured Information Standards (Oasis) released the first version of the Web Services for Remote Portlets (WSRP) standard, which specifies the interfaces a remote portlet producer must implement to allow another application (typically a portal) to consume its portlets, regardless of the technology the producer and consumer use (Bellas, 2004). Another standard, the Java Community Process' Java portlet specification, defines a Java API for implementing WSRP-compatible portlets that can be deployed in any standard Java portlet container.

3 Knowledge, learning, and research portals

The major objective of knowledge portals is to provide a comprehensive knowledge management solution to the problem of immediate access to organisations' crucial knowledge assets, reference information, and other resources, as well as to ensure for efficient user orientation and navigation. Essentially, knowledge portals leverage the expertise of experienced knowledge workers by capturing and sharing their knowledge efficiently, thus minimising the new hires' time-to-proficiency and foster ongoing learning and communication with other team members.

The steps of establishing a knowledge portal are as follows (adapted from Schneble (2002)):

1 define the portal's purpose and target users

2 define, classify, and organise the portal's knowledge content

3 identify and develop knowledge objects

4 define organisational and navigational schemes

5 create a portal maintenance plan

6 create and rapidly deploy a prototype knowledge portal

7 validate the portal design by conducting a usability test

8 make necessary changes based on the usability test

9 continue to maintain, cultivate, and migrate knowledge

10 conduct periodic value-add measurements to ensure continued effectiveness.

Roughly speaking, the first five steps are related to portal conceptualisation, analysis, and design. The next three steps pertain to implementation and validation, and the last two steps represent the portal's evolution. The entire process is not sequential as it may seem to be, because in many steps and at many points in time it becomes necessary to reconsider and modify decisions and actions made in one or more of the previous steps.

Knowledge is one way or another connected with learning, hence there is the concept of *learning portals* – specific knowledge portals that stress the learning issue. The New Measure Group defines learning portal as 'a website used by all members of a Professional Learning Community to build and share knowledge. It is a 'place' where information is posted, constantly processed, and is transformed by people into improved practices.' (http://www.newmeasure.com/learningportals.php). It is also the doorway to the capabilities provided by a Learning Management System (LMS), i.e., it often uses an LMS as its basis. Specifically, a learning portal is the interface that allows learners to locate content, track their progress toward training goals, and perform a variety of other related functions (NextDelta Solutions, 2004). Essential value-added services of learning portals include learning management capabilities, learner assessment and tracking, communities of interest, learning content authoring and upload, and the extension of learning to an organisation's value chain (Barron, 2000; see also http://www.sumtotalsystems.com/). And, yes, there is much more than that. The higher-education side of the e-learning spectrum has its own parallel universe of learning portals, providing high-quality courseware.

Since the rest of the paper focuses on research portals only, and research activities are tightly coupled with both knowledge and learning, it is necessary to make two subtle distinctions between research and other knowledge and learning portals. First, the typical category of users of research portals is research communities. Industry, governmental institutions, and other organisations are often interested in using information from some research portals. They are rather infrequent content providers for such portals. Moreover, there is a whole myriad of *research community portals* of different sizes that are mostly of interest to narrow research communities only. The size of such a research community portal is most often determined by its breadth of focus and the community size. Home pages of some of the larger-size and well known community research portals are shown in Figure 2 (the NNR research portal, the gateway to neural network resources) and Figure 3 (the KDnuggets research portal for the field of data mining and knowledge discovery).

Figure 2 The NNR research portal (http://www.neoxi.com/NNR/index.html)

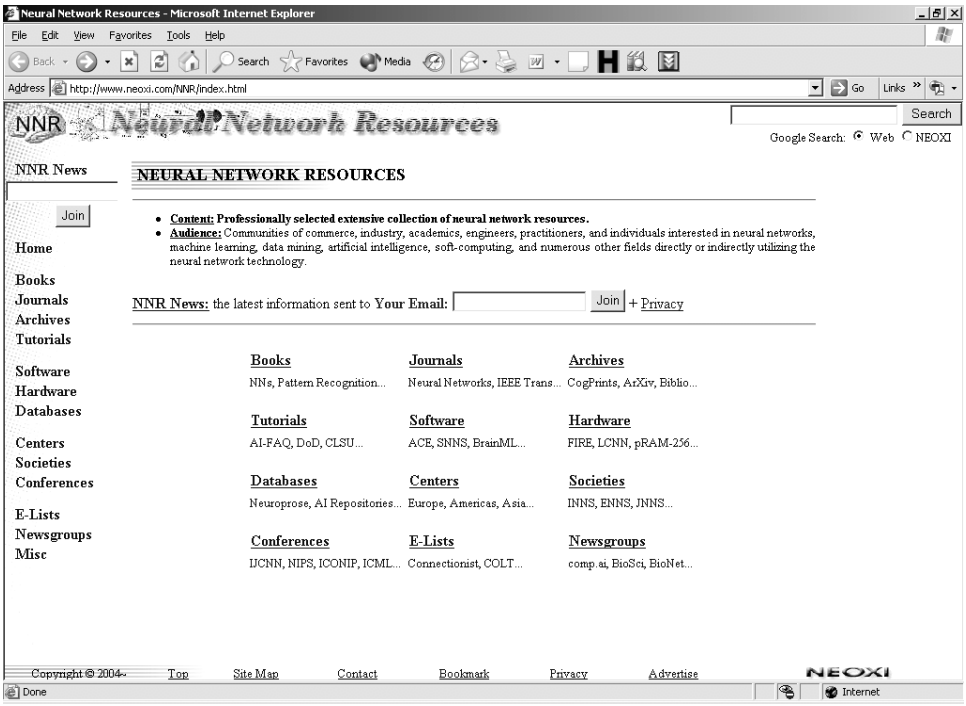

Figure 3 The KDnuggets research portal (http://www.kdnuggets.com/)

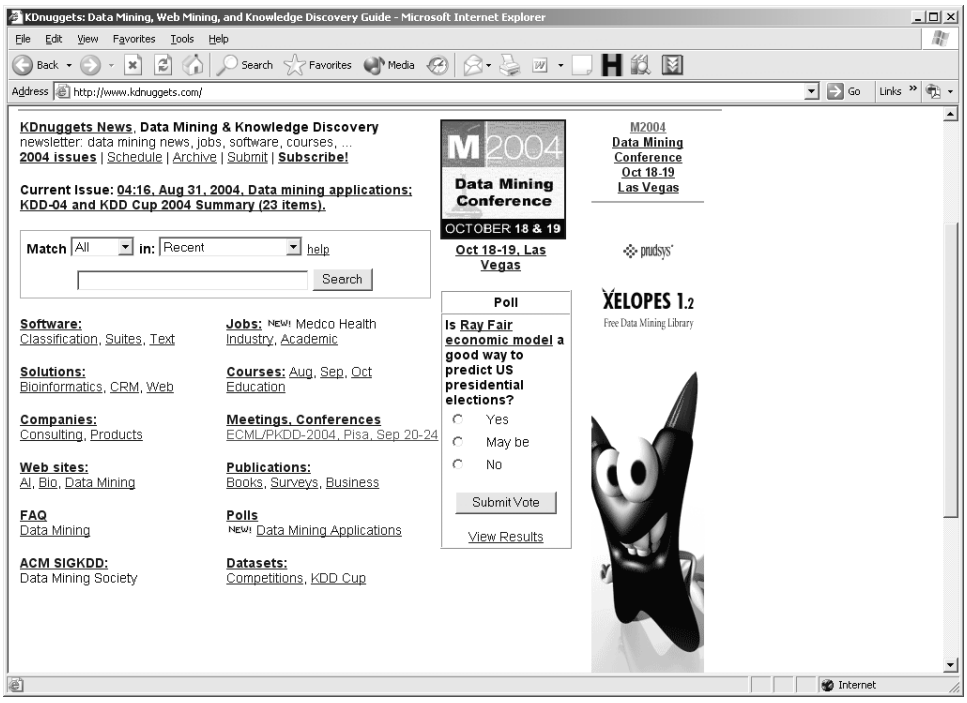

Second, both large-scale research portals and smaller research community portals provide (in addition to other common VEP services) a number of services specific to research activities and not typically found on other community web portals. These include access to:

- relevant research publications, online research journals and magazines, article databases and archives, usually with dropdown menus with a multitude of searchable options
- international information and public policy resources
- conference announcements, technical and professional exhibitions and contests, and information on other relevant meetings of the community members
- specific issues concerning communication and collaboration with industry, research funding agencies, and sponsors
- various benchmarking and best practices
- newsletter, daily postings and weekly bulletins, e-mail news services
- statistical reports, general reports, and idiosyncratic but relevant facts
- direct links to virtual libraries
- direct links to other web pages important for the community
- information for students and faculty members.

From the above list, probably the most important information that each researcher wants from a research community portal is a highly-structured hub of links to most of the relevant research resources. Hence as a rule, all research community portals include on their home pages 'Links', 'Interesting Links', 'Resources', and similar links denoting such hubs that provide direct access to the most frequently visited web pages of interest.

The case study presented in the following sections is a small-scale one – it relates to setting up, populating, maintaining, and using a research community portal by a small research community or group, such as the GOOD OLD AI group (http://goodoldai.org.yu), for daily communication, learning, information and knowledge acquisition and sharing, publishing, and other typical research activities.

4 The GOOD OLD AI portal

The GOOD OLD AI research group started in 1996 as a small group of MSc and PhD students, coordinated by their advisor, interested in various topics in AI and in Software Engineering (SE). Perhaps the closest analogy of the group's research interests and activities can be drawn from the following excerpt from the Aims & Scope of the *International Journal of Software Engineering and Knowledge Engineering, IJSEKE* (http://www.ksi.edu/ijsk.html): 'the interplay between software engineering and knowledge engineering: how knowledge engineering methods can be applied to software engineering, and vice versa.' More specifically, the group members' interests cover knowledge representation, ontologies, Semantic Web, intelligent reasoning, intelligent tutoring systems, e-Learning, intelligent agents, modern software architectures and tools

for developing intelligent systems, adaptive hypermedia, user modelling, software patterns, data mining, web mining, neural networks, and applications of AI and SE to the domains of education, medicine, and information processing.

Over time, the group grew to over 30 researchers, from both academia and industry, and now has its regular meetings and activities. Recently, the group also joined some international projects and organisations (e.g., the EC Framework 6 IST-507310 Network-of-Excellence project ProLearn (http://www.prolearn-project.org/), and the ARIADNE Foundation (http://www.ariadne-eu.org/), and IFIP Technical Committee on Artificial Intelligence, TC 12 (http://www.ifiptc12.org/)).

Until recently, the group had its ordinary website, Figure 4, with about a dozen different pages and links to the group members' personal pages, projects, publications, and other relevant facts. However, increase in the group's activities, abundant contributions from the group members, as well as purely organisational factors gradually led to the idea of organising the group's site in the form of a research community portal. Unfortunately, the idea was pending for about two years, because most of the group members were not familiar enough with the portal technology to allow for a quick implementation of the idea. In fact, out of being overloaded with other duties and activities, the group kept the idea constantly postponed until a talented and student well-knowledgeable of web portal technologies, Zoran Ševarac, joined the group and carried out the initial portal setup himself. From that time onwards, the group's portal at http://goodoldai.org.yu, Figure 5, is constantly evolving.

Figure 4 The earlier home page of the GOOD OLD AI research group

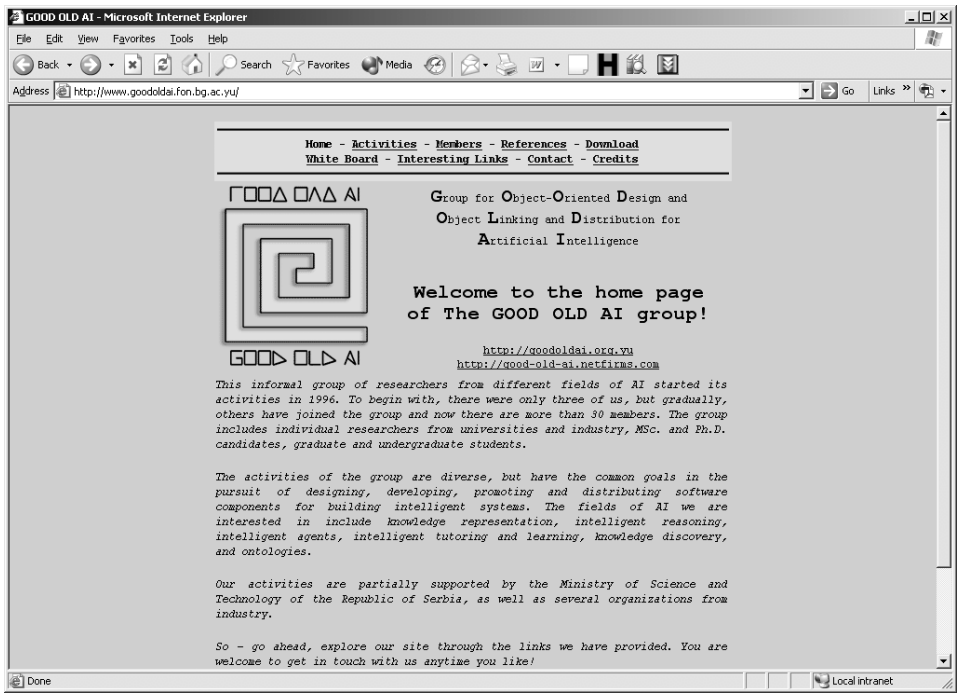

Figure 5 Present home page of the GOOD OLD AI portal

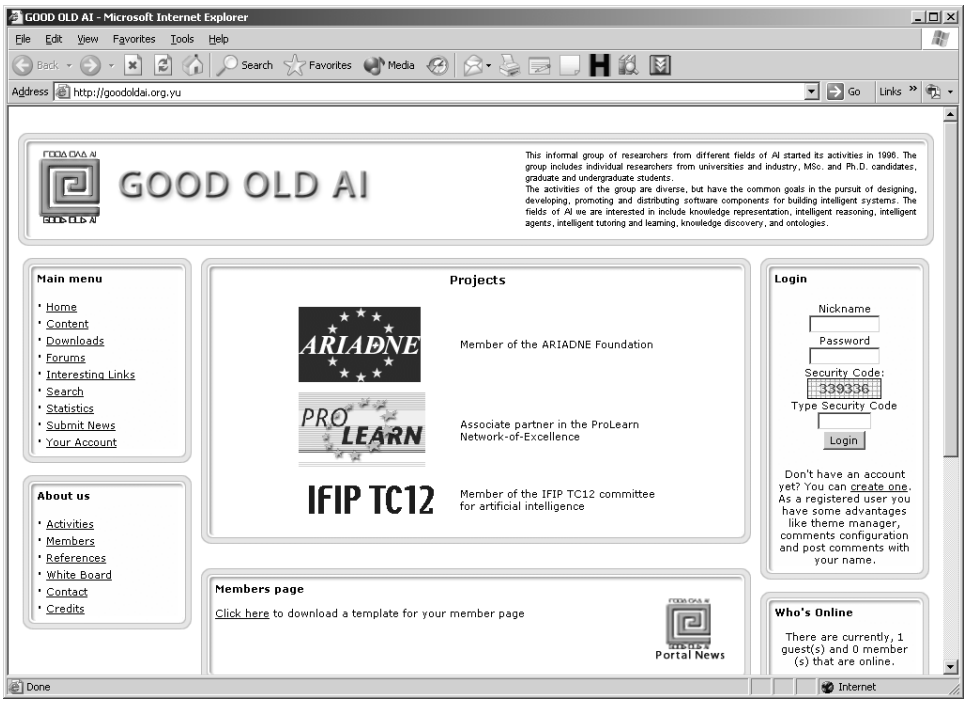

Recalling the steps of establishing a knowledge portal (Section 3), in the case of the GOOD OLD AI portal the first two steps happened to be more difficult than it was originally conceived. As for the portal's purpose and target users, the initial ambition to make a community web portal for AI and SE, or a learning portal, proved to be too-large-a-piece for the beginning. Hence it gradually decreased to the idea of using the portal for 'shaping up an AI research group', i.e., as a dynamic and evolving single point of access to important research information, collections of interesting links, presentation of the group's research results, its international cooperation, and the group members' communication and forums. The portal's initial knowledge content, Step 2, is focused on (but not limited to) a hub of interesting links (Figure 6), the group's research results and member pages, downloads, and international cooperation information and assets. Section 6 covers Steps 3 and 4 in more details, and Section 7 briefly covers the other steps.

Note that in designing the portal and deciding on its contents and knowledge objects, the group members already had quite a clear vision of 'what it should look like'. The vision came from an extensive use of other research community portals and other sites of interests for the group members, such as the Semantic Web Portal (http://www.semanticweb.org), AI Topics (http://www.aaai.org/aitopics/), AIED association and the *IJAIED* journal (http://www.ijaied.org/), Mizoguchi Lab (http://www.ei.sanken.osaka-u.ac.jp/english/), and the already mentioned KD Nuggets and NNR portals.

Figure 6 The *Interesting Links* hub at the GOOD OLD AI portal

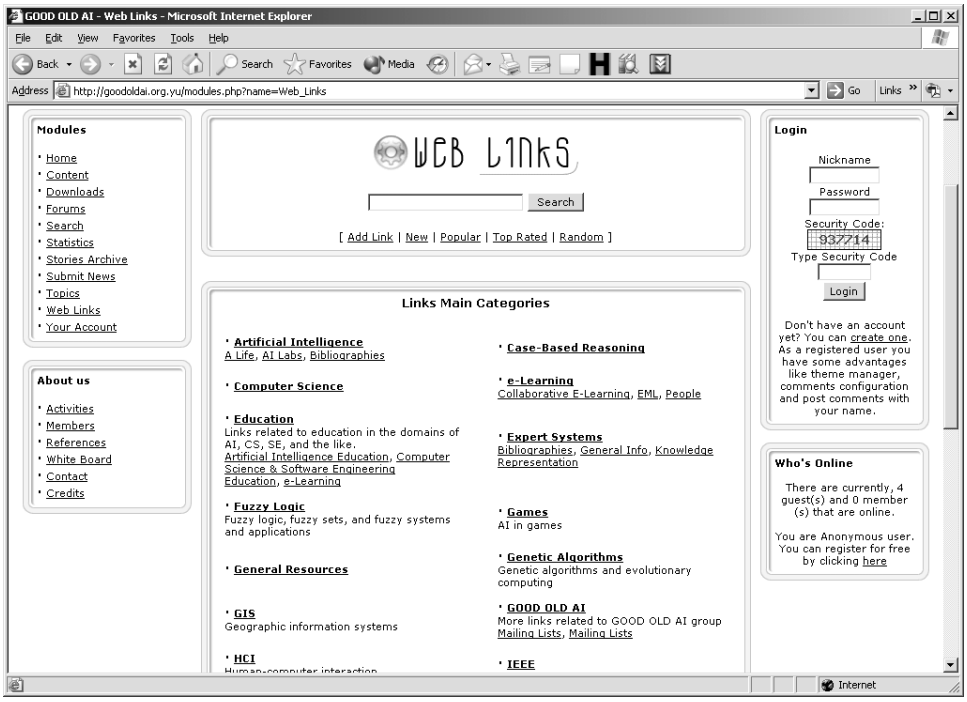

5 Conceptual issues

Why a portal for what is described in Section 4, and not an ordinary website? True, the contents currently in focus of the GOOD OLD AI portal may be well organised without the portal technology. However, there are four compelling conceptual reasons for preferring a portal:

- Storage of and navigation through knowledge assets – once the portal navigation, graphics, and initial content are developed, it is easier not only to browse the knowledge assets, but also to have them stored using well-proven database technology rather that just ordinary file storage.
- Ease of maintenance – adding new knowledge assets, deleting or modifying existing ones, reorganising the site's design, performing backup, and many administrative tasks are much easier with the portal technology and do not require specific skills, hence *any member* with administration authorities can do it.
- Evolution – as the knowledge contents of the site continue to grow, re-designing the site and its underlying databases becomes necessary and the portal technology largely alleviates that task.
- Communication – even in a small research community, day-to-day communication other than email is frequent; virtually all portal toolkits provide forums, chats, newsgroups and other communication tools.

Returning once again to the Step 1 of establishing a knowledge portal (Section 3), deciding on who a research community portal will serve, precisely, quickly leads to the decision of the kind 'general-purpose vs. focused portal.' In other words, should the portal's content cover interest of a broader research community, or should it stick only to a handful of research topics? Narrow-focused research community portals require less specialists to contribute to it and to maintain it, but the vision of the portal's evolution, the research community growth, and the desire to become recognised by more and more experts over time may be the reason for a broader focus. The GOOD OLD AI portal was initially designed focused, but with the idea of evolving into a more general one over time.

'Purely research-oriented vs. practice-oriented' is yet another similar decision to be made at the conceptual level. Purely research-oriented portals are likely to be organised much like hubs of interesting links. If the portal is intended to be used by both academics and industry, then more attention must be paid in the portal design on practical projects, sponsors' ads, links to relevant government agencies, concrete products like software and tools, and so on. Such portals are likely to at least partially evolve into repositories of documents and tools, such as the CMU AI repository (http://www-2.cs.cmu.edu/afs/cs.cmu.edu/project/ai-repository/ai/html/), Java.net (http://www.java.net/), or Source Forge (http://sourceforge.net/). Of course, there may be a number of different 'hybrid' designs between the two extreme categories. For the time being, the GOOD OLD AI portal is more research-oriented than practice-oriented.

Deciding on including the News section in a research community portal must be done bearing in mind the problem of automatic updates. No matter how handy the News section may be for the research community members, it must be automated or it requires some time, accuracy, and dedication. A potential solution to that and many other problems may be the semantic integration of multiple web portals (Hartmann and Sure, 2004). It is a new concept, but provides automated updates across different portals based on ontologies and semantic information. Initial version of the GOOD OLD AI portal did not include the News section.

Finally, it is important to understand and ensure from the very beginning that the research community portal provides a *service* to the members, without becoming a *burden* in terms of facilitating the community's needs for global access and management of research resources. To do their jobs, members of the community need only to learn how to use the portal, to add/delete/modify content, and can continue to use something they understand and are familiar with, without having to learn complex new systems (Bailey and Treloar, 2001).

6 Ontological issues

Identifying the knowledge (content) objects to be put on a portal, as well as defining organisational and navigational schemes for the portal (Steps 3 and 4 in developing the portal) involves some ontological engineering decisions. Moreover, the portal designers can control the results the users get when they do a search for a keyword or click on a link provided on the portal. By addressing such issues, the portal owners can become trusted experts on the portal topic over time (Massa, 2002).

The ontology of the portal topic(s) should not be confused with the portal development tool capabilities and the related portal design decisions (such as article publication and management, downloads, members list, reviews, story archives, and web links as tools provided by the portal development software), although there is some practical interdependency between the two issues. The ontology of the portal topic(s) is about the semantics and workflow of the portal content. To make the things more complicated, different domain experts have different opinions about the domain ontologies and topic hierarchies. Likewise, deciding on the granularity of the topics to be shown to the user at a single page is not only the question of the domain ontology; it also has to do with, e.g., human-computer interaction and the page GUI design. Moreover, no matter what specific portal development tool is used, some terminological inconsistencies with what the domain experts are used to are highly likely to occur – there is no tool that will provide the terminology exactly the same with what the portal users expect it to be. Hence adaptations and trade-offs to that end should be taken into account early.

To partially illustrate this discussion, it should be noted that designing the conceptual-level ontology of the GOOD OLD AI portal's *Interesting Links* was a truly challenging problem. Given the research orientations of the members of the GOOD OLD AI community, as well as the desire to increase external traffic and visits to the portal, resulted in the *Interesting Links* hierarchy as partially shown in Figure 6. One of the driving forces in opting for the portal topic hierarchy as it stands (though it is permanently evolving) was a user-centric approach: it was assumed that the target users are familiar with and used to the topic hierarchies from the other research portals relevant for the community, such as those mentioned in Section 4, as well as with those from general internet portals and HEPs, such as Yahoo, AltaVista, and Infoseek. Currently, the top level of the *Interesting Links* hierarchy includes the following topics: *Artificial Intelligence*, *Computer Science & Software Engineering*, *Education*, *General Resources*, *Related Areas*, and *Research*. In order to further illustrate the (quite large) hierarchy, one level down the tree from *Artificial Intelligence* are the topics such as: *AI Labs*, *AI Education & Research*, *Artificial Life*, *Case-Based Reasoning*, *Expert Systems*, *Fuzzy Logic*, *Games*, and so on. The reader is welcome to explore the other details of the hierarchy at http://goodoldai.org.yu/modules.php?name=Web_Links.

The multidisciplinary research interests in the GOOD OLD AI group introduce further levels of complexity in the topic hierarchies. For example, what if something can be put under different subcategories? Should *AI Bibliography* be put under *Artificial Intelligence*, or under *Research*? Or both? Also, some portal development tools may prevent duplications in the categories/hierarchies. Issues like this one must be expressed explicitly as *constraints* in the domain-ontology/topic-hierarchy design.

There were a number of other ontological decisions of the same kind to be made in the conceptualisation of the GOOD OLD AI portal. The same problem arose with the organisation and design of the other content elements, such as publications, member pages, etc. Furthermore, it is necessary to anticipate the common queries the users put and the typical workflows they follow when using the portal. Finally, the design-for-change philosophy is adopted early, assuming that the problem of conflicting and changing requirements will arise as the contributions from the users accumulate.

All these issues lead to the questions of semantic interoperability, metadata, and content annotation. These are still hot research topics rather than well-established approaches with a stable support by development tools, and only a few standards are

published officially. However, a sound ontological engineering approach in the research community portal design makes it necessary to consider these issues seriously.

7 Technological and practical issues

GOOD OLD AI portal was developed using the open source web portal framework PHP nuke (http://phpnuke.org). It was a cheap solution. Commercial frameworks would require much more initial investments (initial build is expensive, the maintenance burden is large, some training may be needed, as well as tech support and new skills (Strauss, 2000)). However, it also introduced some practical problems that the group did not anticipate (see below).

On the positive side, Steps 6 through 8 in establishing the portal (Section 3) were easy. It took one man-week effort to build an 'empty' prototype version, starting from the group's pre-existing website (Figure 4), and to initially populate it with some content. It took another 3–4 easy man-months to further populate the content pages, validate the portal design starting from the conceptual-level domain ontology (Section 6) by conducting usability tests, and make some changes.

It was only then that the group has made Step 5, i.e., a portal maintenance plan was adopted informally by the group members. It is a long-term effort for the GOOD OLD AI group to further maintain, cultivate, and migrate the knowledge to/from the portal, as well as to periodically evaluate its usability and evolution (Steps 9 and 10).

The group members who initially populated the portal with content have got used to the underlying PHP nuke tools easily. It is a totally automated and advanced content management system to create dynamic websites, written in PHP and uses the most popular SQL database servers. From the user side, PHP nuke's features that are of primary importance for the GOOD OLD AI as a research community portal include online status of guests and members, search functionality, user-customisable web pages, statistics module for information about page views and visitors, article publication and management, and modules like feedback, download, members list, reviews, story archives, and web links. On the admin side, it is possible to take backup of the database at one click, manage blocks and content, add a file for download, administer users, FAQ, fora, newsletters and messages, as well as reviews, sections, polls/surveys, topics, modules and add-ons, and also to conduct the website configuration. There can be different ranks of the portal administrators. Ordinary users can make their comments and suggestions to the administrators.

The downside of using PHP nuke as an open-source tool for developing the GOOD OLD AI portal was related to the problems of administration and security. Although ordinary tasks such as content management seem to be quite straightforward at the first glance, it takes a considerable time to learn all the details and tricks of PHP nuke administration tools. A painful experience of that kind for the GOOD OLD AI group was when one morning the group members discovered that the portal was attacked by hackers, in spite of the careful security level setup. Luckily, the damage done was quickly repaired due to regularly performed backup. Still, the problem of secure administration and content protection remained for a while and it took additional effort to mitigate it.

8 External references to the GOOD OLD AI portal

Quite naturally, each research community strives to increase the traffic on its portal and get recognised. Personal contacts and advertising the portal help – for example, it was through personal contacts that resulted from international cooperation that the GOOD OLD AI portal is referenced from the IFIP TC12 committee. However, it requires a constant inflow of content contributions and a lot of maintenance efforts if the portal is to remain popular.

A good sign of increasing popularity of a research community portal are external references to it that do not come from advertising it. For example, the GOOD OLD AI portal is now referenced by the Encyclopedia of the INSEAD's Centre for Advanced Learning Technologies (http://www.insead.fr/CALT/Encyclopedia/ComputerSciences/AI/) as a starting point in the area of AI, and by the University of Colorado at Denver, School of Education, (http://carbon.cudenver.edu/~mryder/reflect/theory.html), among others. Note, however, that the group has received a number of external requests for referencing the portal elsewhere for commercial purposes that had nothing to do with the group's mission and have been thus rejected.

9 Lessons learned

The following bullet points summarise the GOOD OLD AI group's experiences with setting up, using, evaluating, and maintaining its research portal:

- Although the ten steps of establishing a knowledge portal are listed in Section 3, (adapted from Schneble, 2002), these should not be taken as "The Ten Commandments" of developing a research community portal. They provide a good development framework and useful guidelines.
- How should a research group get started on building its portal? In brief, by recognising the benefits of the portal technologies to the group's daily work, activities, and research, then deciding who it will serve, deciding on the content, and finally getting the right people knowledgeable of portal technologies to set up the initial version quickly. Evolving the group's homepage into a portal could be a good strategy, but it should be done as part of an overall plan.
- The portal should evolve gradually, a few features at a time (Strauss, 2000). It is easy to set up the initial version. However, it is a long-term effort for the research group/community to maintain and evolve the portal.
- Portals are vulnerable and their security is critical, hence authentication and other forms of protecting the portal is necessary even for research community portals (passwords, session keys, different Public Key Infrastructure techniques, and the like).
- Usability tests must be conducted intensively in early development stages. They ensure that the ontological issues are understood correctly and that the portal is likely to attract attention in the future as well.

- Portal technology changes and improves quickly, especially in the domain of open-source tools. An eye should be kept on new developments – they may improve the portal's usability and security at very low cost.
- Each research community portal should support the community members' activities, communication, resource exchange, innovation, and learning. Throughout the portal lifetime, it is necessary to periodically measure and assess its value to the community. Researchers have already suggested methodologies of how to do it; for example (see Liebowitz, 2003).
- It takes time for a research portal to get recognised and accepted by the community. Advertising the portal to interested individuals, institutions, and professional organisations is important.
- Portal technology is already advanced. Further advancements will follow as well, and to that end an important trend to follow is that of semantic integration of content from multiple related portals (Hartmann and Sure, 2004; Staab et al., 2000). Luckily, new trends like that one usually develop using exactly research community portals.
- It is about knowledge, not about technology. Technology is important, but knowledge is more important.

10 Conclusion

Building a research community portal is a long-term effort, although initiating it and using the portal technology is easy. Even small research communities and groups may want to consider establishing their portals, because of a number of benefits portals can bring them at a low cost. However, it should be noted that it takes a lot of effort for a research community portal to succeed, in terms of getting recognised as an important authority for the research community and possibly other interested parties. Only large research groups and communities build their portals quickly and efficiently.

The members of the GOOD OLD AI group are ready to share their experiences in setting up their portal with other researchers. Contacts, comments, and enquiries from other people will help the group improve the GOOD OLD AI portal. The group has already put a lot of efforts in establishing it, but feels that it is only a beginning. The group members hope the lessons they learned about knowledge portals and that they shared in this paper will be useful to all other potential developers of research community knowledge portals, especially to those who feel the need to introduce some structure, organisation, and personalisation into the abundant online research material they have access to.

References

Bailey, N. and Treloar, A. (2001) *What's Under the Threshold? Portals and Portal Infrastructures*, last visited: September 2004, Online, Available: http://ausweb.scu.edu.au/aw01/papers/refereed/treloar/

Barron, T. (2000) 'A portrait of learning portals', *Learning Circuits – ASTD's Online Magazine for e-Learning*, May, last visited: September 2004, Online. Available: http://www.learningcircuits.org/2000/may2000/Barron.htm

Bellas, F. (2004) 'Standards for second-generation portals', *IEEE Internet Computing*, Vol. 8, No. 2, pp.54–60.

Hartmann, J. and Sure, Y. (2004) 'An infrastructure for scalable, reliable semantic portals', *IEEE Intelligent Systems*, Vol. 19, No. 3, pp.58–65.

Liebowitz, J. (2003) 'Keynote paper: measuring the value of online communities, leading to innovation and learning', *International Journal of Innovation and Learning*, Vol. 1, No. 1, pp.1–8.

Massa, B. (2002) *What is a Portal Anyway and What Can It Do For Me?*, last visited: September 2004, Online, Available: http://www.ksinclair.com/Article72.htm

Mitchell, B. (2004) *What is a Portal, Really?*, last visited: September 2004, Online, Available: http://compnetworking.about.com/library/weekly/aa011900a.htm

NextDelta Solutions (2004) 'The zen of e-learning', *White Paper*, last visited: September 2004, Online, Available: http://www.nextdelta.com/zen.doc

Schneble, J. (2002) 'Case study: knowledge portal at Cisco', *Learning Circuits – ASTD's Online Magazine for e-Learning*, May, last visited: September 2004, Online, Available: http://www.learningcircuits.org/2002/may2002/schneble.html

Staab, S., Angele, J., Decker, S., Erdmann, M., Hotho, A., Maedche, A., Schnurr, H.-P., Studer, R. and Sure, Y. (2000) 'Semantic community web portals', *International Journal of Computer and Telecommunications Networking*, Vol. 33, Nos. 1–6, pp.473–491.

Strauss, H. (2000) 'What is a portal, anyway?', *Syllabus Magazine*, January, last visited: July 2004, Online, Available: http://www.syllabus.com/techtalks/events/000120portal.asp

Wege, C. (2002) 'Portal server technology', *IEEE Internet Computing*, Vol. 6, No. 3, pp.73–77.