AKTuality:

An Overview of the Aims, Ambitions and Assumptions of the Advanced Knowledge Technologies Interdisciplinary Research Collaboration

Nigel Shadbolt*

Kieron O'Hara†

1. Introduction

Knowledge Management (KM) is not only about Knowledge Technology. The organisational debris from failed attempts to impose new technical infrastructures that are either inappropriate to their work environments, or where people are not willing to share knowledge is ample evidence. However, even if we accept the slogan, "80% people, 20% technology", large scale KM initiatives will invariably depend on some form of technology to assist communication and storage; moreover, certain forms of KM depend on new, emerging technologies, and technology can (as in many fields) also be a driver for innovation. AKT brings together knowledge technology specialists from both the computing and human sciences, with an acute recognition of the interplay between people, organisations and technology in effective KM.

In this paper, we will discuss six challenges of knowledge management, and show how any integrated KM approach will need to meet these challenges. The AKT research programme is built around these challenges, and the technologies that result will be essential for anyone who needs to deal with knowledge and to maximise its value. The research is being carried out with a strong focus on real-world testbed scenarios, ensuring applicability. There is no presumption that an organisation will need to reconfigure itself completely just to run advanced knowledge systems; we aim to produce advanced and illustrative prototypes of software and methods that will slot easily into a business environment.

This is an ambitious programme, but with the AKT consortium's resources and impressive track record in collaborative research, we are confident that we will achieve our goals.

2. From Information Overload to Actionable Knowledge

In our wired world, data and information have never been so easy to access and store. Globally, it's estimated that 1-2 exabytes of data are now being generated each year, almost all of it in purely digital form (Varian 2000, O'Hara 2002, pp.24-27). Not surprisingly, this

* Intelligence, Agents, Multimedia Group, Dept. of Electronics and Computer Science, University of Southampton, Highfield, Southampton, SO17 1BJ, United Kingdom, nrs@ecs.soton.ac.uk.

† Intelligence, Agents, Multimedia Group, Dept. of Electronics and Computer Science, University of Southampton, Highfield, Southampton, SO17 1BJ, United Kingdom, kmo@ecs.soton.ac.uk.

info-bonanza is experienced all too frequently not as a boon, but as a burden. We are not just swimming but drowning in data. Managers, researchers or workers have to grope their way through the 'infosmog' before reaching their decisions, confident of one thing alone: they have not taken everything properly into account.

We find ourselves in this situation because whilst accessing and storing data and information is technically cheap and relatively straightforward, *interpreting* it requires intellectual investment: *attention, time, expertise and experience* are now the commodities in highest demand and shortest supply. Veterans and other experts are so valued because they *make sense* of information, seeing patterns, implications and connotations that others miss. Knowledge, like beauty, is thus in the eye of the beholder: 'one man's data is another man's knowledge'.

So, in contrast to information, let's be pragmatic and define *knowledge* as *information sufficiently interpreted to enable action.* Everything is information until it is interpreted and enables some form of action (even if only to decide that it is irrelevant). Even if accessible one click away, until it is interpreted – by a human or a machine – it is meaningless, and thus useless. When information is interpreted, then it can be matched with, and brought to bear upon, the particular problems your business or organisation is addressing. To turn your information into actionable knowledge, it is necessary to understand the connections between it and your business processes. This brings us to the kinds of pressing questions now facing organisations: what information to use when, how to find it, who may know, who may already have rated its significance, how to present it to the right people in the right form at the right time, etc.

This is, of course, easy to say and hard to do. Most information is not currently stored in ways that assist interpretation, whether by people or machines. People can always interpret information given enough time and skill, but these are scarce, valuable resources. Put bluntly, *there's plenty of information technology in use, but barely any knowledge technology* (O'Hara 2002).

If knowledge assets are to be managed effectively, then Advanced Knowledge Technologies that render information interpretable in order to enable effective action must be a part of the solution. AKT is very much about 'the economics of attention.' Our goal is to release people from coping with the information deluge, in order to focus on the really pressing, hard, interesting problems. AKT consortium technologies already enable machines to do some of the basic interpretive work: filtering, classifying and recommending information. We can intelligently present hyperlinks to assist browsing, and can link people to relevant experts. We have a suite of tools for enriching content with conceptual information, automatically, semi-automatically, and where human intervention is required or desirable, manually.

Knowledge Context:
Knowledge as a Source of Competitive Advantage

Intellectual resources – knowledge – have fast become the major source for the creation of value. Knowledge has overtaken labour and capital as the basis of competitive advantage in developed economies.

Pharmaceuticals, software, foods, telecommunications – examples abound of products and markets where raw materials are of negligible value compared to the know-how that goes into design, assembly and marketing. Knowledge and innovation are "a matter of life and death" in today's competitive markets

Baumol, 2002, p.viii

3. Challenges for Managing Knowledge

Dealing with knowledge is easier said than done (witness the dreams, and failure, to codify expertise as rules in any but extremely well understood domains). Most often it lies within an organisation implicitly, out of sight, undervalued and underused. Often, it leaves the building each night at 5pm. When people move from department to department, or to a competitor, they take their knowledge with them. Expertise that has been built up over years can be lost overnight.

Managing the flow of knowledge around an organisation is a challenge. More precisely, it is *six* challenges. Failing to meet any one of these challenges can derail an organisation's ability to use its knowledge assets to its best advantage.

We see the six generic challenges for knowledge management as follows:

- **Acquiring it**: Although we often suffer from an overdose of data – infosmog – all too often the problem is that the knowledge available is insufficient or poorly-specified. The challenge here is to get hold of the information that is around, and turn it into knowledge by making it *usable*. This might involve, for instance, making tacit knowledge explicit, identifying gaps in the knowledge already held, acquiring and integrating knowledge from multiple sources (e.g. different experts, or distributed sources on the WWW), acquiring knowledge from unstructured media (e.g. natural language or diagrams).

- **Modelling it**: Modelling bridges the gap between the acquisition of knowledge and its use. Knowledge model structures must be able to represent knowledge so that it can be used for problem-solving. One important knowledge modelling idea is that of *ontologies*, which are specifications of the generic concepts, attributes, relations and axioms of a knowledge base or domain. Ontologies can act as placeholders and organising structures for acquired knowledge, while also providing a format for understanding how knowledge will be used.

- **Reusing it**: One of the most serious impediments to cost-effective use of knowledge is that often knowledge bases or systems are constructed afresh. It is unusual for problem-solving experience or domain content to be acquired and then reused, partly because knowledge tends to require different representations depending on the problem-solving that it is intended to do. Understanding the use and application of knowledge would enable more leverage to be gained from the knowledge already at hand, thereby increasing the returns on the investment in those knowledge assets.

- **Retrieving it**: When a knowledge repository gets very large, finding a particular piece of knowledge can become very difficult. There are two related problems to do with knowledge retrieval. First, there is the issue of finding knowledge again once it has been stored, understanding the structure of your archive in order to navigate through it efficiently. And second, there is the problem of retrieving the subset of content from the repository that is relevant to the matter in hand. This second problem, the dynamic extraction of knowledge from a repository, may well set problems for a knowledge retrieval system that alter regularly and quickly during problem-solving.

- **Publishing it**: The challenge of publishing or disseminating can be described as getting the right knowledge, in the right form, in the right place, to the right person, at the right time. Different users will require knowledge presented and visualised in different ways, and the quality of such presentation is not merely a matter of preference, but can radically affect the value of the knowledge to the user. Getting presentation right will involve understanding the different perspectives of people with different agendas, while an understanding of knowledge content will help to ensure that important related pieces of knowledge get published at the appropriate time.

- **Maintaining it**: The last challenge is to keep the knowledge repository functional. This may involve the regular updating of content as content changes (e.g. as price lists are revised). But it may also involve a deeper analysis of the knowledge content. Some content has a considerable longevity, while other knowledge dates very quickly. If a repository of knowledge it to remain active over a period of time, it is essential to know which parts of the knowledge base must be discarded and when. Other problems involved in maintenance include verifying and validating the content, and certifying its safety.

Each of these six challenges requires a multidisciplinary strategy. They also implicitly define a seventh challenge: to produce an *integrated response* to the six challenges. "Integrated" in the sense of embedded into an organization's existing infrastructure, knowledge flow and working practices. And "integrated" in the sense that we need a seamless way to manage knowledge throughout its lifecycle, where a system would facilitate, for example, both the acquisition of knowledge and its retrieval. The links between the challenges are related to issues such as, for instance, a common infrastructure, and a life cycle model that treats the various challenges as elements of a global process.

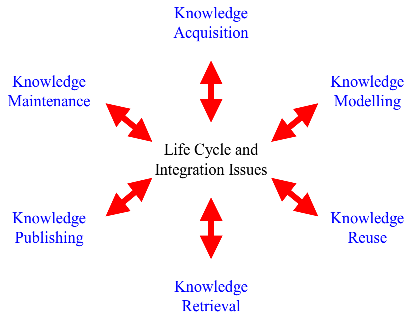

4. AKT: Advanced Knowledge Technologies

AKT is an interdisciplinary research collaboration, funded by Britain's EPSRC, to address the six knowledge challenges in an integrated fashion, developing technology and an understanding of the dynamics of knowledge within an organisation or firm to facilitate the management of knowledge and the extraction of value from it.

The six challenges described above – which every knowledge manager will recognise – are the organising principle behind AKT. The major technology development workpackages within the project are each focused on one of these challenges. Of course, an integrated approach is essential, and other workpackages are devoted to ensuring integration between the approaches to each challenges, and to ensuring underlying, lightweight, flexible infrastructure support.

To this end, the second focus of the AKT project is on a number of testbed scenarios. These are real-world problems, which exist independently of the project, to which we will apply our ideas, ensuring relevance, practicality and applicability. The challenges will drive our theory; the testbeds will drive the applications of that theory.

It can be seen from the discussions of the knowledge context that technology is already part of the problem. The Internet is a technological

AKT's Roots

The original impetus for AKT came from two workshops held under the aegis of the EPSRC on Software Assisted Knowledge Acquisition (SAKA). SAKA I was held at Abingdon in November 1996, under the coordination of Professor Derek Sleeman. It brought together leading academics and industrialists from a number of major blue chip (e.g. Unilever, BA, BT, CCN, Rolls-Royce) and high tech (e.g. Cogsys, Integral Solutions, Telictech, CCL) firms to discuss the requirements that industry would wish to place on the knowledge technologies of the future. SAKA II, held at Nottingham in April 1997 under the coordination of Professor Nigel Shadbolt, involved leading knowledge technology academics from the UK, Europe and the US in discussions on the industrial requirements that SAKA I had elicited.

These meetings identified a felt need within industry for Advanced Knowledge Technologies and sketched the intellectual agenda that evolved into AKT.

creation, and it is technology that will be required to get the most out of the new developments of the semantic web. And information overload is caused chiefly by computational storage methods multiplying the amount of knowledge that can be stored enormously while diminishing the costs of the storage. Whether these are seen as technological problems or technological opportunities, it is clear that technology is an essential part of maximising the use of an organisation's knowledge in the knowledge context we have described.

Technology Context: The Next-Generation Semantic Web

The World Wide Web is like a large library. But searching through it is unintelligent and often painful. Furthermore, as it gets larger, the dangers of infosmog will increase.

The **Next Generation Semantic Web** is getting around these problems by developing ways of sifting *intelligently* through information. By providing different *layers* of representation for content, the Semantic Web allows increasingly sophisticated and more powerful inferences to be made about the content, as the details are abstracted away from the raw material.

For instance, imagine a news story as it appears in the Open University's *KMi Planet* ejournal. Displayed there, the content is 'rendered' – generated on the fly – for display on the web. The 'look' of the page is determined behind the scenes by a **hypertext markup language** (HTML) encoding. For instance, what makes the headline the headline is a pair of tags '<H1>' and '</H1>' which show where it begins and ends.

The next layer above is the **extended markup language** XML, which allows a more abstract – and hence more powerful – representation. Now the headline is delimited by the tags '<Headline>' and '</Headline>'. These tags are arbitrary, and their only meaning is by mutual agreement by specific communities of users. So the people behind *KMi Planet* have decided that '<Headline>' should be interpreted as '<H1>'. But mobile phone users, for example, may interpret the '<Headline>' tag differently, to allow it to appear in a very small font. Hence XML allows what is called 'repurposing' of content.

So far so good. But we want intelligent lookup. If we move up another layer, we can now get some *inferential capability*, via the **Resource Description Framework** (RDF). RDF represents relations as *triples*, two objects and a relation between them. Both objects and the relation are actually locations on the web. This is a simple but amazingly powerful idea, allowing much more complex searches than are possible on the World Wide Web, enabling principled, large scale indexing and retrieval by machines.

But it is not yet intelligent. The top layer of representation in the next-generation semantic web might be in the **Web Ontology Language** (OWL). This, and languages like it, provide the ability to *reason* about the relations stored in RDF triples, thus going beyond search and retrieval, giving the machine some capacity to understand the implications of particular relations – for instance if an article is *authored-by* a person, then that person *wrote* the article. It is representation languages like this that support the applications AKT is developing.

The challenge of the Semantic Web is to find a representation language powerful enough to support automated reasoning, but simple enough to be usable. **WebOnt**, a World Wide Web Consortium working group which includes two AKT members, is currently wrestling with this problem.

Only one **puzzle** remains. If it is so intelligent, why is the Web Ontology Language not called WOL?

5. Knowledge Dynamics, Knowledge Services

As we have emphasised, to say that technology is crucial is far from saying that it is 100% of the solution to KM problems. Technology lives within other contexts than the knowledge context. Technology has to be used within organisations to meet the problems to which those organisations need to respond. Therefore it is essential that we come to understand the particular requirements of the firms and organisations that will wish to use advanced knowledge technologies, so that we can provide genuine solutions to real problems.

The aim of AKT is to develop a deeper understanding of the *dynamics of knowledge* in a technological, industrial context, the ways that knowledge might be said to 'flow' around organisations, and how the use of knowledge should be integrated into the Semantic Web. This will enable us to become a *knowledge service provider*. AKT's research enables knowledge solutions to be tailored to particular knowledge problems, rather than requiring your whole knowledge management practice to be uprooted and replaced.

Put another way, AKT's research agenda is **not**: develop technologies to facilitate knowledge management. Rather, the agenda **is**: understand the *dynamics* of knowledge, and thereby the roots of knowledge management problems; understand the *properties* of knowledge, and thereby develop the most effective representations and technologies; understand the *capabilities* of knowledge technologies, and thereby match appropriate solutions with problems.

Understanding Knowledge Dynamics

What do we mean by the *dynamics of knowledge*? We engage with a wide variety of perspectives on issues such as the nature of individual and collective organisational knowledge, how 'it' may be said to 'flow', or what it means to say that we have 'captured', 'stored', or 'reused' knowledge. Do different technologies encourage the dominance of particular metaphors for knowledge? For instance, what assumptions do we make when we use formal ontologies to 'model knowledge'? We draw inspiration from relatively recent work that has articulated the nature of the *"practices"* from which the term *community of practice* derives its name, a concept that has gained enormous currency in recent years within both the KM research and business communities. *Practices* are the stuff of everyday office life, emerging through the interplay of *tacit processes* and *formal artefacts* (see left).

> …Such a concept of practice includes both the explicit and the tacit. It includes what is said and what is left unsaid; what is represented and what is assumed. It includes language, tools, documents, images, symbols, well-defined roles, specified criteria, codified procedures, regulations, and contracts that various practices make explicit for a variety of purposes. But it also includes all the implicit relations, tacit conventions, subtle cues, untold rules of thumb, recognizable intuitions, specific perceptions, well-tuned sensitivities, embodied understandings, underlying assumptions, and shared world views. Most of these may never be articulated, yet they are unmistakable signs of membership in communities of practice and are crucial to the success of their enterprise.
>
> Wenger, 1998, p. 47.

For instance, we see formal ontologies (regarded by many as a foundation for the Semantic Web) as symbolic tools within a community of practice. We have shown that when there is reasonable consensus, it opens the possibility of introducing knowledge services based on an ontology co-constructed by knowledge engineers with stakeholders. An ontology can reify a "shared world view", creating "well-defined roles", "specified criteria" and "codified procedures" as described by Wenger. We regard symbolic representations such as ontologies as *boundary objects* whose role is to support communication and negotiation over meaning between communities of practice.

Investigating what we have termed 'knowledge dynamics' draws on our expertise in cognitive science, philosophy and collaborative work processes, and as we build and integrate technologies, we maintain an ongoing dialogue with other leading KM research and industrial groups for whom this is of particular interest. The result is a much deeper and nuanced understanding of how knowledge technologies impact, and are moulded and sustained by, the everyday work practices of the communities using our tools.

Configurable Knowledge Services

In tandem with studying knowledge dynamics, the design and implementation of our technologies are shaped by a number of principles:

- Maximal flexibility.
- Incorporation of the best aspects of standard or well-known software (if Lotus Notes will do the trick, then why force people to learn new languages and interfaces?).
- Influencing standards (we have a number of representatives, for example, on World Wide Web Consortium working groups).
- The Internet is crucial. Our primary focus is on how knowledge can best be represented, used and managed on the Semantic Web.
- Avoiding a single overarching system.
- We do not assume that users will be willing or able to scrap their existing systems, learn new representation languages, or reorganise their knowledge management teams. Of course, solving a problem or responding to an opportunity will require change to corporate practice. Nevertheless, our research is intended to help determine the *appropriate* amount of change, and so avoid the extremes of mere cosmetic tinkering on the one hand, and overreaction on the other.
- A flexible approach to dissemination and product take-up, to provide all and only the knowledge services users need.

Our project is focused and structured around the six knowledge challenges outlined above. However, that does not mean that we are forcing knowledge managers and other professionals to 'see' or conceptualise knowledge in a particular, rigid, way. We maintain that the six challenges will present themselves in different guises to anyone who is concerned with the problems of maximising the value of knowledge.

Technology Context: Ubiquitous Computing

A trend that must be recognised is ubiquitous computing. The Internet is largely understood from the point of view of the currently dominant user interface, the PC. However, it is clear that Internet access will be much more widely available in future thanks to developments in communications and wireless technology. The networked computer is melting into the infrastructure: when you can access the Internet from mobile phones, handheld computers, or even your car, keys or credit card, it is clear that many of our ideas about the presentation and publication of knowledge and information using the Internet will need to evolve.

6. AKT: Technical Agenda

Given our remit to 'think out of the box', the technical agenda of AKT can be unpredictable! Nevertheless, at the half way stage of our research, the general tenets of our approach are clear.

Knowledge acquisition (KA) research is focusing on the development of KA techniques appropriate for the new information-rich environment. In particular, KA has to become ubiquitous; as the web becomes the information-seeker's first port of call, one is invisible without a web presence. To support the creation of knowledge, while preserving trust and privacy, we will need new techniques for representing and trawling the vast knowledge resources of the web (O'Hara forthcoming, chapter 5). AKT is currently experimenting with techniques for harvesting and integrating knowledge about British Computer Science, currently storing millions of items of information (Leonard & Glaser 2001). The information is already there: until now no-one had the resources to make sense of it.

Ubiquitous KA will depend on artefact enrichment, in other words, methods, languages and logics for annotating and enriching the content of objects such as web pages, KA materials, text fragments etc. Such content enrichment could be the new bottleneck in content creation. Again, a "killer app" here would be to make content enrichment a painless by-product of use (Ciravegna et al 2002, Vargas-Vera et al 2002).

One area of interest is that of multimedia KA, i.e. coming to understand (finding models and structures for) knowledge expressed in non-text media. Another is using natural language information extraction (IE) techniques (Maynard et al 2002), or integrating them with standard KA methods, for example applying domain ontologies to facilitate the IE from texts (Vargas-Vera et al 2001).

In **knowledge modelling**, the use of *ontologies*, powerful formalisms for expressing and structuring conceptual schemes for domains, will be a central idea to look at here (Compatangelo & Meisel 2002), but will not be the only focus of our work.

There will also be a drive to examine special-purpose modelling formalisms for particular important subsets of discourse. One example will be the modelling, dissemination and sharing of complex processes within organisations (Chen-Burger et al 2002). Another is the modelling of argumentation and design rationales (Buckingham Shum et al 2002). A third is the modelling of domains to support semantic markup (Ciravegna et al 2002).

In **knowledge reuse**, the aim will be to exploit the software engineering principle of reuse as much as possible. For instance, it is possible to identify user profiles, such as navigation strategies for getting round a complex website. Reuse of such strategies is very likely to be a valuable extension of our web technologies. Intelligent reasoning services can be provided by allowing users – including non-programmers – to assemble knowledge-based prototypes quickly from reasoning components stored in distributed libraries. Other examples include facilitating ontology reuse by supporting merging and sharing (Kalfoglou & Schorlemmer 2002), or multiple reasoning methods over them (Hu et al 2003), and enabling knowledge base reuse (Sleeman et al 2002).

In **knowledge retrieval**, one of our aims is to enable the retrieval of knowledge represented in natural language (Cunningham et al 2002). The use of natural language in retrieval has two aspects. First, there is the use of natural language as the basis for the interface to our knowledge services. For instance, ontologies could be used to interpret knowledge queries in natural language forms from users. Second, there is the extraction and retrieval of information from texts and other unstructured sources. Other focuses in this challenge include the identification of informal communities of practice (O'Hara et al 2002), and the location of other key resources within organisations (Gibbins et al 2003, Hasan 2003, Wills et al 2002). We are also very alive to the technological possibilities of the enriched markup of web pages.

For the challenge of **knowledge publishing**, our work includes the creation of representational forms – web pages, narratives, sets of hyperlinks – tailored to context, using ontologies, on the fly (Kalfoglou et al 2001, Kim et al 2002). Such an agenda, which exploits input from computational linguistics, hypertext theory, and the exploitation of semantic structures from ontologies, is highly promising. A related area is the creation of software to capture the structure of meeting discussions, facilitating communication both synchronously and asynchronously (Buckingham Shum et al 2002).

In **knowledge maintenance**, the challenge is to understand the quality of the knowledge, which of course is an essential factor in any decision to continue maintaining or scrapping some knowledge repository. This challenge includes assessing the fitness for purpose of knowledge bases and other resources, improving the representational efficiency (e.g. by ensuring that referring terms are used consistently – Alani et al 2002), and using technologies to determine the value of knowledge and other intangible assets. It also requires understanding the changes and evolution of knowledge over the lifecycle of systems, especially distributed systems (Schorlemmer et al 2002).

7. Conclusion

AKT is a project of extraordinary ambition, with the remit and resources to 'think the unthinkable'. AKT is creating technologies to take knowledge through the whole of its lifecycle from acquisition to obsolescence, technologies that are integrated and yet flexible enough to apply to the specific problems of individual users in their organisations. Our systems are grounded in other organisational contexts in order to understand how the dynamics of knowledge shape – and are shaped by – our tools. Our understanding of, and familiarity with, the technological context – particularly the knowledge-releasing potential of the Semantic Web – ensures that our work will have relevance and impact in years to come.

8. Acknowledgements

This paper was supported under the Advanced Knowledge Technologies (AKT) Interdisciplinary Research Collaboration (IRC), which is sponsored by the UK Engineering and Physical Sciences Research Council under grant number GR/N15764/01. The AKT IRC comprises the Universities of Aberdeen, Edinburgh, Sheffield, Southampton and the Open University.

The paper includes material from other sources. The authors would like to acknowledge Simon Buckingham Shum for material from the *AKT Manifesto*, and Marc Eisenstadt and John Domingue for material from the *Re:AKT* newsletter.

9. References

Harith Alani, Srinandan Dasmahapatra, Nicholas Gibbins, Hugh Glaser, Steve Harris, Yannis Kalfoglou, Kieron O'Hara & Nigel Shadbolt (2002) 'Managing reference: ensuring referential integrity of ontologies for the semantic web' *Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2002)*, and this volume

William J. Baumol (2002) *The Free-Market Innovation Machine: Analyzing the Growth Miracle of Capitalism* (Princeton University Press, Princeton)

Simon Buckingham Shum, David De Roure, Marc Eisenstadt, Nigel Shadbolt & Austin Tate (2002) 'CoAKTinG: collaborative advanced knowledge technologies in the grid' *Proceedings of 11th IEEE International Symposium on High Performance Distributed Computing, 2nd Workshop on Advanced Collaborative Environments*, and this volume

Yun-Heh Chen-Burger, Austin Tate & David Robertson 'Enterprise modelling: a declarative approach for FBPML' *Proceedings of 15th European Conference on Artificial Intelligence (ECAI) Knowledge Management and Organisational Memories Workshop*, and this volume

Fabio Ciravegna, Alexiei Dingli, Daniella Petrelli & Yorick Wilks (2002) 'User-system cooperation in document annotation based on information extraction' *Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2002)*, and this volume

Ernesto Compatangelo & Helmut Meisel (2002) 'EER-CONCEPTOOL: conceptual analysis of EER schemas and ontologies, Technical Report, Dept. of Computing Science, University of Aberdeen, and this volume

Hamish Cunningham, Diana Maynard, Kalina Bontcheva & Valentin Tablan (2002) 'GATE: an architecture for development of robust HLT applications' *Proceedings of the 40th Anniversary Meeting of the Association for Computational Linguistics (ACL '02)*, and this volume

Nicholas Gibbins, Steve Harris & Nigel Shadbolt (2003) 'Agent-based semantic web services' *Proceedings of the 12th International World Wide Web Conference (WWW 2003)*, and this volume

M. Maruf Hasan (2003) 'A spreading activation framework for ontology-enhanced adaptive information access within organisations' *Proceedings of the AAAI Spring Symposium on Agent Mediated Knowledge Management (AMKM 2003)*, and this volume

Bo Hu, Ernesto Compatangelo & Inés Arana (2003) 'Inference fusion: a hybrid approach to taxonomic reasoning' *Proceedings of the 16th International FLAIRS Conference*, and this volume

Yannis Kalfoglou, John Domingue, Enrico Motta, Maria Vargas-Vera & Simon Buckingham Shum (2001) 'MyPlanet: an ontology-driven web-based personalised news service' *Proceedings of IJCAI 2001 Workshop on Ontologies and Information Sharing*, and this volume

Yannis Kalfoglou & Marco Schorlemmer (2002) 'Information flow based ontology mapping' *Proceedings of the 1st International Conference on Ontologies, Databases and Applications of Semantics (ODBASE '02)*, and this volume

Sanghee Kim, Harith Alani, Wendy Hall, Paul Lewis, David Millard, Nigel Shadbolt & Mark Weal (2002) 'Artequakt: generating tailored biographies with automatically annotated fragments from the web' *Proceedings of 15th European Conference on Artificial Intelligence (ECAI) Semantic Authoring, Annotation and Knowledge Markup Workshop*, and this volume

Thomas Leonard & Hugh Glaser (2001) 'Large scale acquisition and maintenance from the web without source access' *Proceedings of K-CAP 2001*, and this volume

Diana Maynard, Hamish Cunningham, Kalina Bontcheva & Marin Dimitrov (2002) 'Adapting a robust multi-genre named entity system for automatic content extraction' *Proceedings of 10th International Conference on Artificial Intelligence: Methodology, Systems, Applications (AIMSA 2002)*, and this volume

Kieron O'Hara (2002) *Plato and the Internet* (Icon, Cambridge)

Kieron O'Hara (forthcoming) *Trust in the Twenty-First Century* (Icon, Cambridge)

Kieron O'Hara, Harith Alani & Nigel Shadbolt (2002) 'Identifying communities of practice: analysing ontologies as networks to support community recognition' *Proceedings of the IFIP World Computer Congress Workshop on Information Systems: the E-Business Challenge*, and this volume

Marco Schorlemmer, Stephen Potter, David Robertson & Derek Sleeman (2002) 'Knowledge life cycle management over a distributed architecture' *Expert Update* 5, and this volume

Derek Sleeman, David Robertson, Stephen Potter & Marco Schorlemmer (2002) 'Enabling services for distributed environments: ontology extraction and knowledge base characterisation' *Proceedings of 15th European Conference on Artificial Intelligence (ECAI) Knowledge Transformation for the Semantic Web Workshop*, and this volume

Maria Vargas-Vera, John Domingue, Yannis Kalfoglou, Enrico Motta & Simon Buckingham Shum (2001) 'Template-driven information extraction for populating ontologies' *Proceedings of IJCAI 2001 Workshop on Ontology Learning*, and this volume

Maria Vargas-Vera, Enrico Motta, John Domingue, Mattia Lanzoni, Arthur Stutt & Fabio Ciravegna (2002) 'MnM: ontology driven semi-automatic and automatic support for semantic markup' *Proceedings of the 13th International Conference on Knowledge Engineering and Knowledge Management (EKAW 2002)*, and this volume

Hal Varian (2000) *How Much Information?* http://www.sims.berkeley.edu/research/projects/how-much-info/summary.html

Etienne Wenger (1998) *Communities of Practice: Learning, Meaning and Identity* (Cambridge University Press, Cambridge)

Gary Wills, Harith Alani, Ronald Ashri, Yannis Kalfoglou & Sanghee Kim (2002) 'Design issues for agent-based resource locator systems' *Proceedings of 4th International Conference on Practical Aspects of Knowledge Management*, and this volume