

SERENA: A Multi-site Pervasive Agent Environment That Supports Serendipitous Discovery in Research

Jamie Forth¹, Thanasis Giannimaras¹, Geraint A. Wiggins¹,
Robert Stewart², Diana Bental², Ruth Aylett²,
Deborah Maxwell³, Hadi Mehrpouya³, Jamie Shek³, and Mel Woods³

¹ School of Electronic Engineering and Computer Science
Queen Mary University of London, Mile End Road, London, England, E1 4NS

² School of Mathematical and Computer Sciences
Heriot-Watt University, Edinburgh, Scotland, EH14 4AS

³ Duncan of Jordanstone College of Art and Design
University of Dundee, Perth Road, Dundee, Scotland, DD1 4HT

Abstract. We present SERENA, a multi-site, pervasive, agent environment that supports serendipitous discovery in research. The project starts from the premise that human users cannot be aware of all the research information that is relevant to their work, because of the compartmentalisation of research into fields around particular journals, and, simply, because there is too much to know. In particular, the Semantic Web provides a resource which can assist, but there is more to be discovered than the things that a user might deliberately search for. SERENA, then, attempts to assist researchers by presenting them with information that they did not know they needed to know about their research.

Keywords: Serendipity, semantic web, pervasive agents.

1 Introduction

We describe SERENA, an agent-based, semantic, pervasive, embedded personal assistant system, for academic researchers. SERENA is designed in response to the RCUK Digital Economy “Designing Effective Research Spaces” challenge¹, and is, in the first instance, intended for arts and humanities researchers.

The key idea of SERENA is to create a Serendipity Arena: a virtual space in which serendipitous discovery of several different kinds is more likely to happen than elsewhere. For example, two SERENA users with common research interests might be travelling to the same place, but not be aware of the fact; SERENA would alert them and, if its users desired, arrange a discussion meeting. However, the aim is not merely to create an academic dating service; rather, the idea extends to objects in the world too, so that, for example, a researcher arriving at King’s Cross station in London could be informed that a unique manuscript,

¹ www.epsrc.ac.uk/SiteCollectionDocuments/Calls/2009/DERSSandpitCall.pdf

of direct interest to their work, and borrowed temporarily from Egypt, is on show in a nearby museum on that day; what is more, they could be told this in advance, when they book their train ticket on line, so they can plan for a visit. Above all, SERENA is meant to find for its users things that they did not know they wanted.

These superficially simple tasks involve significant background reasoning, about users' interests and activities, and about real-world data, such as location and time, both current and planned. They require a substantial software infrastructure that can travel with the user, but also supply substantially more computational power than a standard mobile device; this is sometimes referred to as "the Mobile Dilemma". The infrastructure must also be capable of semantic analysis of a user's writings and actions, if not in real time, at least in a time frame that is practically useful in context (consider the examples above: the temporal requirements of the corresponding notifications are quite different). SERENA's target users cannot be assumed to be expert computer users, and therefore it must work with them in ways that are both easy to use and engaging.

Thus, under detailed scrutiny, the superficially simple idea of SERENA quickly expands from an engaging and useful tool to an archetypal general AI problem, including language understanding, proactive semantic reasoning, intelligent interaction and pervasive presence. The only factor to reduce the challenge is the nature of serendipitous discovery itself: unusually in computer science, and only within reason, the system can be *usefully ambiguous or even wrong*. This is so because, ultimately, SERENA forms a hybrid system with its users, and unexpected, off-the-wall information from the computational part of the hybrid can nevertheless usefully inform or stimulate the human part. Therefore, SERENA has reduced responsibility for correct reasoning: its suggestions need only be useful; they do not have to be logically correct.

Luck et al. [5, p. 40] identify seven key challenges for multi-agent system research, all of which are relevant to SERENA—though not all are currently implemented. Here, we exemplify challenges 4.2 (increase quality of agent software to industrial standard), 4.3 (provide effective agreed standards to allow open systems development), 4.4 (provide semantic infrastructure for open agent communities) and 4.6 (develop agent ability to understand user requirements). The remaining challenges are for future work. This said, the key contribution described in this paper and associated demo, on the practical applications of agents, is the *assemblage* of the system, not its components.

The funded SERENA project, as a whole, includes designers, HCI and usability specialists, and computer scientists, working together in a broad coalition. It has considered the epistemology of serendipity [8,6,7], the meaning of the concept of serendipity to researchers, and their reactions to it [13], and approaches to concept extraction from text [11]. Also work has been done on language processing with a view to discovering users' interests from their tweets [12] and their goals from their notes and email messages [10]. The current paper, though, takes a top-down perspective, explaining the overall conceptualisation of the system,

explaining how the primary challenges are met, and in particular focusing on the agent system that forms the core of the pervasive environment.

With this aim in mind, the rest of the paper describes a selection of the functions of SERENA, from a user’s perspective, identifying the engineering challenges, and outlines how they are integrated by SERENA’s core agent system. There is not space here to describe every aspect of SERENA in detail. Rather, we aim to give a flavour its capabilities, and to demonstrate how the agent system contributes, on multiple levels, to the elegant implementation of the whole.

2 Affordances and Constraints

We now describe SERENA from the user’s point of view, explaining how the various affordances of the software affect its design. We identify the points at which the agent approach is particularly useful.

2.1 Ubiquitous Intelligent Personal Support

SERENA is conceived as a ubiquitous computing environment, which should interact with its users via their mobile devices and desktop computers, and via public installations in public spaces. The requirements of these three *interactors* are somewhat different. The first two are *private* to the user, and authenticated, while the third is *public*: this difference impinges on the nature of the information that can be displayed. Mobile devices tend to have significantly smaller displays than desktop machines, and a public display is more likely to be large than small, so this dimension distinguishes the first category of interactor from the second and third. In the SERENA project so far, we have focused on the mobile private SERENA, with only one proof-of-concept public SERENA interactor currently planned. These are described in §2.2.

SERENA’s ubiquity, delivered through mobile and fixed devices, imposes constraints on its design, which agent-based designs are well suited to meet. First, the behaviour of the system must be consistent and persistent. That is to say, an interactor must behave uniformly, and must not lose information if, for example, it is switched off, or (more likely) if there is a network outage. A very neat way of addressing this issue is to maintain an agent that simulates the mobile device, and then to implement synchronisation of information between the mobile device and its agent, which can be done independently of the workings of the rest of the agent system, and of the user’s interaction with the device². We term this kind of agent, that echoes the behaviour of an entity in the physical world, a *shadow agent*³. Shadow agents of other kinds appear too, the most important being the user agency, a group of agents that shadow the user, supplying information about the user (e.g., location, research interests, privacy settings) to

² Of course, the user’s interaction with the device as a whole may be restricted by a temporary network outage, but this is not a soluble problem.

³ The term is borrowed from the UK government system, where the official Opposition forms its own cabinet, *shadowing* the ministerial (dys)functions of the government.

the rest of the system, but also autonomously acting on their behalf to make serendipitous discoveries. The user agency is described in §3.2.

A key advantage of the shadowing approach is that it allows us to meet SERENA's requirements for high-power computing and network access in static installations with high-powered servers with fast access to the Internet, both of which are *sine qua non* for the deep inference required of SERENA if it is to be helpful to its users. A further concomitant advantage is the stability and security of such managed systems. The agent community, running on these servers, can work uninterrupted, communicating internally with the shadows of the physical world interfaces, and those interfaces can be updated live when connected, or asynchronously when an interactor reappears on the network after being disconnected. What is more, the asynchronous nature of the agent community absolves SERENA's implementers of the need to manage the appearance and disappearance of interactors, and also the delays inherent in web services and sites; agents communicate via message passing, and, simply, when they are disconnected, or when there is a delay, no external messages appear (though of course the corresponding shadow agent may continue working independently on the basis of prior information). A final bonus of shadowing is in the use of shadow agents to represent external internet resources within the agent community. This means that issues of translating between languages and formats need not be spread throughout the system, but can be handled by a specialist shadow, so that the information is manipulated into a SERENA-friendly form, exactly once, as soon as it enters the system. Our use of FIPA ACL standards⁴ means that responses can be straightforwardly associated with their corresponding queries, on receipt, using conversation management.

SERENA is implemented using the Java Agent DEvelopment Framework⁵ [1], which supports two important SERENA requirements. In JADE, agents run in notional *containers*, one or more containers per machine, but multiple machines can be connected together into a JADE *platform*. This allows failover to be implemented at the agent level: duplicate containers of SERENA agents can be run at multiple sites, and one fails, another can take over.

2.2 Interactors

Private SERENA

One of SERENA's design constraints is that its users should not have to use specialist interfaces and equipment; it should work invisibly behind familiar interaction paradigms, making itself noticed only when a direct user response is required, or when the user's attention needs to be drawn to SERENA output. To this end, we have designed an interface, which is the current focus of our interaction work, as an Android app, conceptualised as a *Semantic*⁶ *Notebook*

⁴ <http://www.fipa.org/specs/fipa00061/SC00061G.html>

⁵ JADE; <http://jade.tilab.com>. It seems likely that Erlang (www.erlang.org) is a future candidate for such implementations.

⁶ This epithet is justified further in §3.

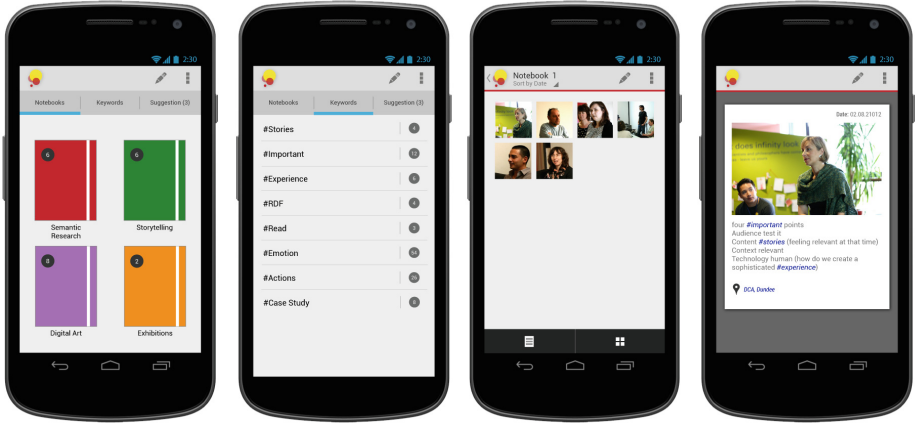


Fig. 1. The Semantic Notebook App. Left to right: Notebook view (creating a new notebook and personalising the cover); Goal list (keywords and goals); Sort By: Visual and date; Keywords and Multimodal notes with Goals (Tags in the body of the text)

[9]. The researcher-user is invited to make notes, add tagged images, keywords, and so on, all in free text. The notes can be organised in a predictable but useful way. Interaction with the user is then managed by SERENA processing the user’s text, and then adding annotations (e.g., items of text, web links), to the notes, making the distinction between the user’s notes and SERENA’s additions clear by means of typography. Example views are given in Figure 1. When SERENA adds a new annotation, the Android notification system informs the user, according to their preferences. The user can then follow up the suggestions at their leisure, or delete them if they wish. In some cases, time-sensitive notifications will fade away when they become stale, though a user might choose to override this.

This kind of interaction is advantageous for two reasons: first, SERENA must avoid the paperclip effect⁷; and, second, the reasoning required to suppress pointless information is often extensive (see §3), and cannot be performed on the fly, while the user waits. The concomitant advantage of asynchrony is that network outages do not degrade the experience: the user will come to expect SERENA suggestions at some time after they make their notes, but not immediately.

Public SERENA

The criteria for a public instantiation of SERENA are quite different. First, there are significant security issues concerned with a user’s personal or private information: in the current prototype, we address this simply by using information that we know to be public. There are major open opportunities in a public installation of this kind: it is not merely intended to be a terminal that individuals

⁷ The irritation produced when Microsoft’s Clippy character used to intrude unexpectedly, distracting the user from their task, with often incorrect information.

can use to access private SERENA, but something altogether more collective. One key affordance, given the right social context, is the ability to communicate with more than one user at once via the same channel. For example, at an academic conference, one might display connections between consenting delegates, in terms of the connections that SERENA has found between their work and interests: since SERENA is focused on finding *unexpected* connections, this approach might be expected to add value to the social interaction at the event.

Our first prototype public SERENA is conceived in the context of a major UK city library. It is architecturally melded with the building, in that its outputs are projected directly on to walls, using site-specific designs that integrate with the architecture. Its outputs consist of simple visualisations (the simplest being mere text) of documents that are ordered via the library's on-line order system. The information is filtered, so that no connection with the library user ordering it can be made: their name is not displayed, nor is the time at which they ordered it. In principle, however, interesting work could be done finding connections between documents ordered by different people, and displaying on public SERENA a summary of these: a sort of local, temporary Zeitgeist analysis. An example design is shown in Figure 2.

Communication with the Agent System

SERENA's interactors communicate with the Agent System via shadow agents as outlined in §2.1. Externally, the connection from the shadow to the mobile device is implemented directly in Java, running as a background process in Android, and connecting to the shadow via a persistent WebSocket^{8,9}. This process then communicates with front end UI managers, such as the Semantic Notebook App. Necessarily, some local storage of information (such as what is on the current display, what is in the notebook, both from user and from SERENA, and what has or has not been synchronised with the shadow agent) must be managed, and this is done using a local database, working in the same language as the agent system (see §3); other issues of synchrony are dealt with at the level of WebSockets or below, in TCP/IP.

3 Knowledge Representation and Inference to Support Serendipity

3.1 Formalism: The Semantic Web

In order to make interesting and unexpected connections, SERENA explores and combines information from multiple sources, using the growing array of Semantic Web resources. Increasing amounts of information from many different domains are being made available as Linked Open Data (LOD). LOD uses syntactic and semantic standards such as The Resource Description Framework¹⁰ (RDF) and

⁸ <http://www.w3.org/TR/websockets/>

⁹ <http://tools.ietf.org/html/rfc6455>

¹⁰ <http://www.w3.org/RDF/>



a)

<p>Thought-images : Benjamin's ghosts : interventions in contemporary literary Caroline drama : Language and stage in Medieval and Renaissance England / The language of space in court performance, 1400-1625 / The art of rhetoric made easy, etc. / John Holmes. Actors and acting in Shakespeare's time : The History of Jamaica: or, general survey of the ancient The history of Jamaica : reflections on its situation Virtues and Vices and other essays in moral philosophy. The elements of moral philosophy / James Rachels. Elements of the art of rhetoric : Shakespeare's freedom A Poetical Grammar of the English Language, and an Epitome Nine lives of William Shakespeare Shakespeare William Shakespeare</p>	<p>Vybledia? fotografie Shakespeare The Art of Rhetoric made easy, or the elements of oratory. Coping with global environmental change, disasters</p>
--	---

b)

Fig. 2. An example design for one instantiation of public SERENA. Books being ordered from a library are presented as part of the architectural structure of the library in near-real time. a) The live display. b) Direct view of the information displayed.

OWL¹¹ (the Web Ontology Language), and is available for query across the web. SERENA also uses Semantic Web ontologies currently being developed and integrated to express information in specific domains, such as FOAF¹² to describe people and relationships, DBpedia¹³ for general knowledge, GeoNames¹⁴ for geographic locations and DBLP¹⁵ and Dublin Core¹⁶ for publications. These ontologies and resources are being realised in individual web-accessible databases which can be searched by tools such as Sindice¹⁷ or merged into larger databases of machine-readable information such as FactForge¹⁸. As a result of these initiatives it is now possible to combine information from many different sources, at a general, domain-independent level, to link location data obtained from a mobile device to information about nearby places of interest, for example.

SERENA's domain knowledge and agent control commands are represented in RDF, a good choice for expressing highly structured knowledge-based information; together with the OWL-DL subset of OWL, this affords a Description Logic [2, ch. 9], a well-understood basis for knowledge representation and inference. SERENA's agents communicate in RDF, also, and its message envelopes meet the FIPA Agent Communication Language specification. This approach eliminates the need to translate domain knowledge acquired from the Semantic Web to internal agent knowledge representations. It can also support multiple levels of agent reasoning.

Because SERENA agents are themselves described in RDF, the possibility is open to build agents that reflect on and modify the behaviour of others. However, this is currently deferred to future work.

3.2 User Modelling: Agent.Me

Over time, SERENA builds a model of its user, expressed in RDF, including information given by the user, information inferred directly under the control of the user, and information inferred about the user by the system. This information is made available to a collection of agents, which act as the user's shadow in the agent community. The modelling process is kick-started by our Discover.Me.Semantically service¹⁹, which searches for information about a new user, and then consults with them to select what is relevant.

¹¹ <http://www.w3.org/TR/owl2-overview/>

¹² <http://www.foaf-project.org>

¹³ <http://dbpedia.org>

¹⁴ <http://www.geonames.org>

¹⁵ <http://dblp.uni-trier.de>

¹⁶ <http://dublincore.org/documents/dc-rdf/index.shtml>

¹⁷ <http://sindice.com>

¹⁸ factforge.net

¹⁹ Source code under GPLv3 license:

<https://github.com/robstewart57/discover-me-semantically>

Running instance:

<http://serena.macs.hw.ac.uk/serena/discover-me-semantically/>

Discover.Me.Semantically is a web-based tool that allows its user to author RDF representing their professional and personal interests, skills and expertise. The stand-alone implementation allows the user to download this RDF representation as a file to be hosted on their own web pages²⁰.

The standalone implementation (which the reader is invited to try) also offers a path to visualize this RDF on a linked-open-data visualization called LodLive²¹, to explore the paths along links away from their skills and interests. The tool knows of several Web resources, and records some aspects of equivalence between them, and these *sameAs* links can also be explored. The RDF representation of the user so generated can also be stored in SERENA's user model, and then be used by a user's shadow, or Agent.Me to assist in finding information of interest. Other readily available information sources also contribute to the user model, such as bibliographical information, obtained from a user's BIBTEX or EndNote file. In the longer term, we intend to repurpose ideas from Intelligent Tutoring Systems, in which the computer's model of a user is made visible to the user [3], so that the user can reflect on it, and correct it if necessary.

3.3 Supporting Serendipity with Inference: Goal Detection

It is in the nature of serendipity to be unpredictable: if one could create the effect to order, it would not be serendipity, by definition. Our aim, instead, is to enhance the conditions in which serendipity might take place. Ultimately, this process will be managed by the agent system, with some agents requesting information to give to the user, and others searching for answers according to the competence of the resource they are shadowing. The asynchronous, open nature of the agent system, and also of our interface method (see §2.2) mean that relatively little overhead needs to be expended on simple question-answer interaction. More interesting reasoning, however, can be carried out by generalised reasoning agents, but this is future work.

A key issue in supplying the user with useful information is to understand the research goals that they are expressing in their notes, files and email²². We have begun work with the GATE natural language processing system²³, with some success in detecting goals in natural language sentences [4], and an ontology for goals has been defined [10]. These ideas will inform the Agent.Me.

An example of these ideas in action can be found in our case study of connecting users within one institution. For this, we use information gathered for the UK Research Excellence Framework²⁴. The information is about publications of every academic in the institution (coverage is not universal, in fact, but this does not prevent the system from working), and about the academics themselves.

²⁰ Having a `foaf.rdf` file attached to one's academic webpage is becoming commonplace for Semantic Web based researchers.

²¹ <http://lodlive.it/>

²² Of course, SERENA does not access files or email without permission.

²³ <http://gate.ac.uk>

²⁴ <http://www.ref.ac.uk/>

This information is available in RDF format, delivered via a web service. Also included, for many publications, is the plain text of the abstract.

Our approach is to extract semantic annotations from titles and abstracts, using the OpenCalais web service²⁵. Once this is done, we run semantic web queries to deduce answers to such leading questions as:

- Which people in different schools (who therefore may not know each other) describe the same concepts in their papers?
- Are there more experienced specialists (e.g., professors) who often publish papers on concepts also of interest to (e.g.) early-career research associates?

Perhaps unsurprisingly, these queries produce many answers, and we are working on heuristics to filter them in a useful way, with respect to promoting serendipity.

4 System Overview and Example

Figure 3 illustrates the overall system architecture with some example agent types, and also shows the potential communication between agents in two specific tasks. These are: information flow resulting from the addition of a new note to the Semantic Notebook App (dashed arrows); and information flow resulting from the arrival of two users with common interests in a location with a public SERENA installation (dotted arrows).

New Notebook Annotations

When User 1’s device shadow agent announces to User 1’s Agent.Me that a new note has been added to the user’s Semantic Notebook, the associated information seeker agent broadcasts a request for information on the tags included in the note. Some time later, a web resource agent finds a semantic web resource that referring to some of the same concepts. The resource is returned to the user’s interactor via its shadow, and presented in the Notebook as a title with a clickable link.

The same broadcast request also reaches an inference agent, whose speciality is to make connections between SERENA users who publish on related concepts in different research fields. This agent can search for publications in DBLP, by concept; it then broadcasts a request to all Agent.Mes to ask for the research fields and institutions of SERENA’s users whose papers it finds. It then returns an answer to User 1’s Agent.Me listing the papers of those users in their institution, who publish on the concepts of interest but in different research fields.

Arrivals and Meetings

A more complicated example arises when two SERENA users who have enabled their location agents arrive at the same conference, where there is a public SERENA interactor. The users’ location agents broadcast their arrival at the

²⁵ <http://www.opencalais.com>

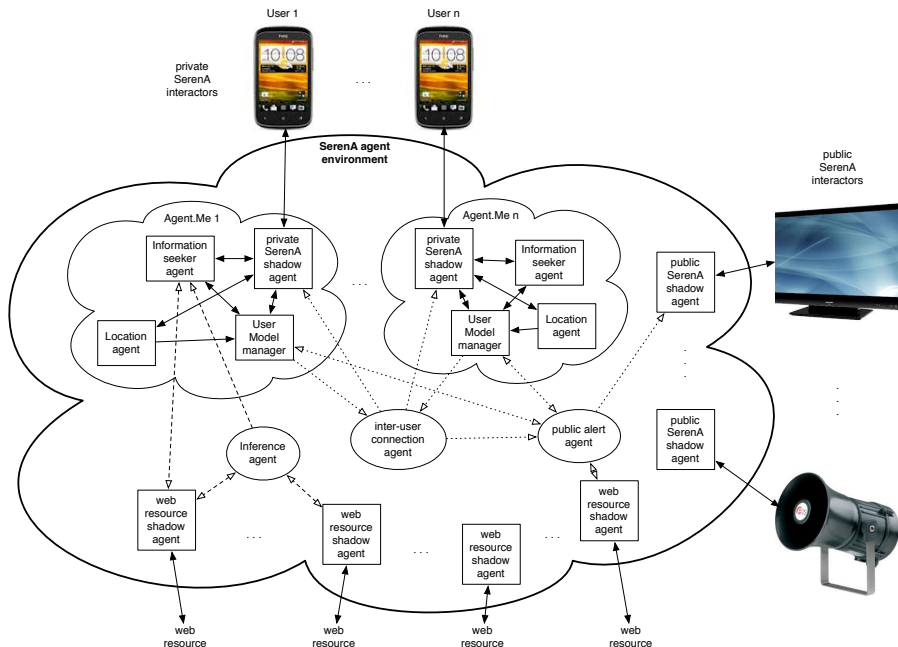


Fig. 3. Overview of SERENA, with two potential agent conversations. Solid arrows indicated on-going administrative information flow. Dashed arrows show communication resulting from the addition of a new note to the Semantic Notebook App on User 1’s phone. Dotted arrows show communication resulting from the arrival of two users with common interests in a location with a public SERENA installation. See §4 for details.

event (which their Agent.Mes infer is significant from their diary entries), and the inter-user connection agent notices their physical coincidence. It queries their respective user models, and learns that they are both interested in being introduced to other SERENA users. The inter-user connection agent broadcasts the opportunity to meet, which is relayed to the users’ interactors by their Agent.Mes, but is also picked up by the local public alert agent, associated with the public SERENA installation. This agent checks with both users that they allow their images to be used on public SERENA, and, if they do, tells the shadow agent to announce their presence and introduce them to each other.

5 Conclusion and Future Work

In a paper of this length, it is not possible to cover all the aspects of a system design as large and diverse as that of SERENA. Here, we have attempted to convey the basic ideas and *raison d’être* of the system, and to give enough detail to explain the contribution of the agent framework and communication style, and

the resulting interaction with users. We believe this to be a novel contribution to practical applications of multi-agent technology, in that to our knowledge, a distributed agent system of this scale has not previously been deployed.

There is, it is clear, a substantial amount of work to do before we can claim that SERENA has fulfilled its potential—although formal evaluation of various aspects of the work outlined in this paper is in progress, there are many more possibilities for emergent behaviour that have yet to be enabled. Detecting and enhancing such behaviours will be the focus of future work.

Acknowledgements. This work is supported by EPSRC Sandpit Research Grant EP/H042741/1 – SerenA: Chance encounters in the Space of Ideas.

References

1. Bellifemine, F., Bergenti, F., Caire, G., Poggi, A.: Jade — a Java agent development framework. In: Bordini, R., Dastani, M., Dix, J., Fallah Seghrouchni, A. (eds.) *Multi-Agent Programming. Multiagent Systems, Artificial Societies, and Simulated Organizations*, vol. 15, pp. 125–147. Springer, US (2005)
2. Brachman, R.J., Levesque, H.J. (eds.): *Readings in Knowledge Representation*. Morgan Kaufmann (1985)
3. Bull, S., Dimitrova, V., McCalla, G.: Open learner models: Research questions. *International Journal of Artificial Intelligence in Education* 17(2), 83–87 (2007)
4. Giannimaras, A.: Goal detection in natural language text. MSc thesis, School of Electronic Engineering and Computer Science, Queen Mary University of London (2012)
5. Luck, M., McBurney, P., Preist, C.: *Agent Technology: Enabling Next Generation Computing*. Agentlink (2003)
6. Makri, S., Blandford, A.: Coming across information serendipitously: Part 1: A process model. *Journal of Documentation* 68(5), 684–705 (2012)
7. Makri, S., Blandford, A.: Coming across information serendipitously: Part 2: A categorisation framework. *Journal of Documentation* 68(5), 706–724 (2012)
8. Makri, S., Erdelez, S.: Thematic issue on opportunistic discovery of information. *Information Research* 16(3) (2011)
9. Maxwell, D., Woods, M., Makri, S., Bental, D., Kefalidou, G., Sharples, S.: Designing a semantic sketchbook to create opportunities for serendipity. In: *Proceedings of BCS HCI Conference* (2012)
10. Piao, S., Bental, D., Whittle, J., Aylett, R., Makri, S., Sun, X.: A pilot study: Deriving a users’ goal framework from a corpus of interviews and diaries. In: *Proceedings of the Workshop on Semantic Relations II - Enhancing Resources and Applications* (2012)
11. Piao, S., Forth, J., Gacitua, R., Whittle, J., Wiggins, G.: Evaluating tools for automatic concept extraction: A case study from the musicology domain. In: *Proceedings of Digital Futures 2010* (2010)
12. Piao, S., Whittle, J.: A feasibility study on extracting twitter users’ interests using nlp tools for serendipitous connections. In: *Proceedings of the Third IEEE International Conference on Social Computing (SocialCom 2011)*. IEEE (2011)
13. Sun, X., Sharples, S., Makri, S.: A user-centred mobile diary study approach to understanding serendipity in information research. *Information by Chance: Opportunistic Discovery of Information* 16(3) (2011)