Weaving Personal Knowledge Spaces Into Office Applications

Heiko Maus, Sven Schwarz, Andreas Dengel

Knowledge Management Department, German Research Center for AI, Trippstadterstr. 122, Germany
{firstname.lastname}@dfki.de

Keywords: Semantic Desktop, Personal Information Management, PIMO, Knowledge Management

Abstract

The paper presents recent developments in our research on Semantic Desktop for personal knowledge management supported by an ecosystem of applications and plug-ins using the knowledge worker's Personal Information Model (PIMO) – a formal representation of his mental model for knowledge work – in everyday applications. We explain how the infrastructure enables the availability of the PIMO as one vocabulary throughout different applications as well as mobile access, the importance of the mental model in the PIMO, and how to get direct benefits from the PIMO in daily activities. We also address steps towards building a Group Information Model from individual PIMOs within the ecosystem.

1. Introduction

The modern working environment places high requirements on knowledge workers: they are confronted with various applications, are involved in several projects and processes, work in changing teams, are on the road with a mobile office, and finally, face an ever increasing flow of information. The resulting knowledge spaces are complex, dynamic, distributed over several applications, and use different vocabulary. It is hard to keep the overview in the resulting personal knowledge space.

This challenge is addressed with the concept of the Semantic Desktop [1,2]. It follows the strategy to embed the mental model of the knowledge worker in daily work by means of a Personal Information Model ("PIMO"). The user's mental model in the PIMO consists of concepts (called "Things" such as specific topics, projects, persons, tasks, …), associations between them (persons are *member of* projects, a task *has topic* "Semantic Desktop", ...), and finally, associated resources (documents, e-mails, web pages, pictures, …) (see [3] for a detailed motivation of the PIMO).

The PIMO serves as an easy to understand conceptualization of the knowledge worker's mental model, which can be used as a common vocabulary across different applications. Therefore, the PIMO provides the means required for a multi-criterial document classification considering the user's subjective view [4]. Fig. 1 shows an example of a PIMO graph with resources from the file system, the web, a task tool. The resources are associated with topics, events, tasks, etc. This graph then serves in different applications to find and access resources or things, to annotate and to relate them (this vocabulary is used in the upcoming figures as well).

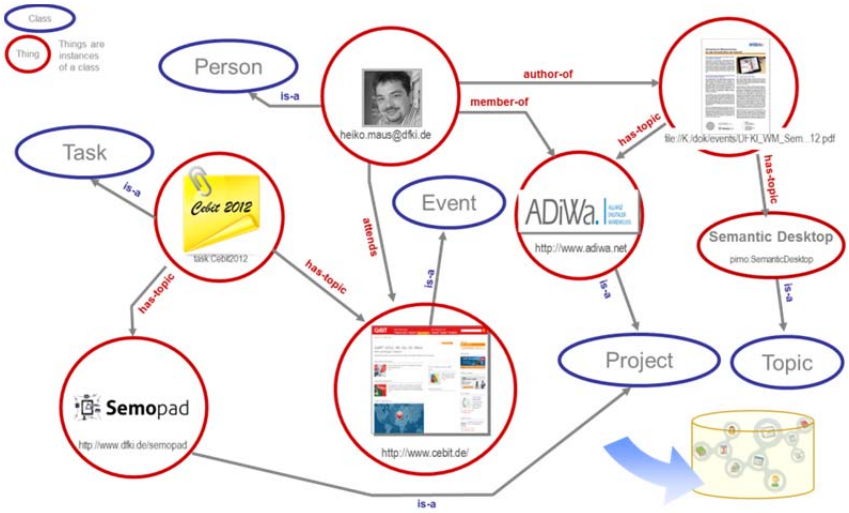

Fig. 1: A schematic excerpt of a Personal Information Model:
Things, classes, resources, and associations

The PIMO uses the semantic power of the formal representation of the PIMO ontology[1] [3], thus, introducing a knowledge representation layer on the user's computer. Besides enabling to annotate and interconnect resources over application borders, further semantic services are possible which make use of the semantic representation of the user's mental model in the PIMO.

Challenges of the Semantic Desktop are the initial bootstrapping of a user's PIMO to cover a relevant part of his current mental model of his knowledge work and the ubiquitous availability of the things and resources in the user's daily work activities. To support this, the Nepomuk[2]-project developed a personal knowledge workbench for the Semantic Desktop [5] providing a comprehensive user interface to create, access, and maintain a PIMO as well as introducing new resources or things into it, e.g., by annotating files, e-mails, and web-pages or writing notes in the semantic wiki. Although the workbench provides valuable means to work with the PIMO for Personal Information Management [6], a comprehensive integration into the user's daily applications was minimal and, hence, the users lacked in-situ support in applications such as email client, web browser or file explorer.

This paper presents an advanced infrastructure for the PIMO ecosystem (section 2) enabling a plug-in architecture for in-situ PIMO access throughout different applications (section 3), thus, providing the knowledge worker one vocabulary for his work, regardless of the application or location. In contrast to the Nepomuk PIMO, which was stored on the user's desktop, this new architecture allows ubiquitous access by storing PIMO data in the cloud, thus, allowing to apply sharing to group members.

2. PIMO infrastructure as semantic middleware on the desktop

As mentioned above, the PIMO is introduced as a knowledge representation layer on the user's desktop. Now, the PIMO is a cloud-based service and provides a service API based on JSON RPC (see Fig. 2). The PIMO Service API uses the PIMO schema with its classes and properties and intended semantics, relies on URIs[3] to identify things and resources, and most importantly, defines a set of methods to access and manipulate the PIMO. In contrast to typical semantic web approaches, the service API does not allow direct access of the core data. So, for instance, the data cannot be read or modified using SPARQL[4] or alike. Instead, a designated set of methods guarantees a consistent and privacy-safe access to the PIMO. It also provides specialized services such as proposals of relevant things for a given text (see information extraction in [8]), a history of used/modified resources and things, as well as, a feed for recently shared concepts, etc. With this approach, we also connected the concept map based knowledge base used in the agile knowledge workflow-tool TaskNavigator [14].

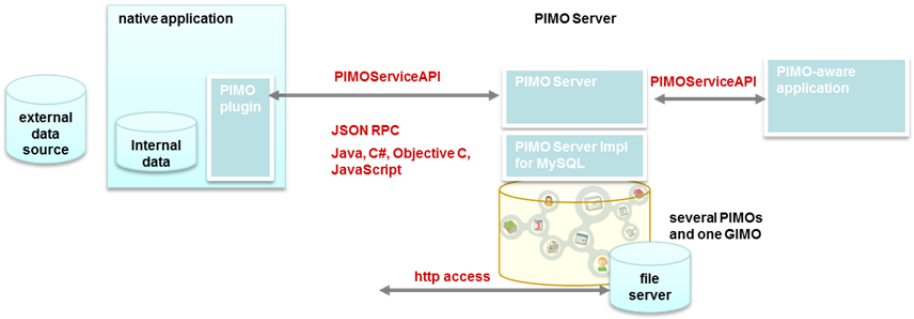

Fig. 2: PIMO architecture

On the client side, we have implementations in different languages such as Java, C#, Objective C, Mozilla's XUL, and JavaScript, allowing to embed in-situ access to the PIMO as a plug-in (such as a sidebar) in different applications (see Fig. 3 to Fig. 5), thus, making it available in standard office applications such as e-mail clients, web browsers, and even in the Windows File Explorer. With the same plug-in mechanism, we also implemented various observation com-

[1] http://www.semanticdesktop.org/ontologies/2007/11/01/pimo/
[2] http://nepomuk.semanticdesktop.org/
[3] Universal Resource Identifier
[4] http://en.wikipedia.org/wiki/SPARQL

ponents in office applications to observe the user's actions and information items used, e.g., with our DragonTalk system for the Mozilla suite[5].

This PIMO infrastructure serves as a *semantic middleware* on the user's desktop realizing the knowledge representation layer and interconnecting various types of applications.

Annotating a new resource will create a new "thing" for that resource in the PIMO. For example, when the PDF file containing the flyer in Fig. 1 (top right) is annotated with the project "ADiWa", a new instance of type *pimo:Document* is created for the file resource. The new thing gets a new, unique URI and is added to the PIMO. This procedure is called "rebirth" as the resource already exists on the file system but is now also represented in the PIMO. The resource's originating location (file path, URL, etc.) is stored as "grounding occurrence" in the PIMO, that way, the provenance is captured and the PIMO GUI can easily "open" the native file when double-clicking on the resource, for example. Metadata such as title/subject, author, recipient, sender, etc. are also present in the PIMO.

Once rebirthed, annotating a thing with some concept in the PIMO is done by adding an association using the PIMO property *pimo:hasTopic*. If the user is willing to disclose a thing, which has an associated grounding occurrence, the user is asked whether the resource should also be shared. Sharing a native file resource then means that the file gets uploaded to the PIMO server and from then on it is available for other users as well as from different devices.

For the research prototype, currently, only a very simple sharing mechanism is implemented, neglecting the task to detect local or remote version changes and updating the versions automatically. Thus, this feature currently only supports *publishing* content like papers, slides, or web pages for a certain topic. However, we understand the importance of an advanced approach for a full-fledged working environment; so, a more sophisticated mechanism is on the agenda.

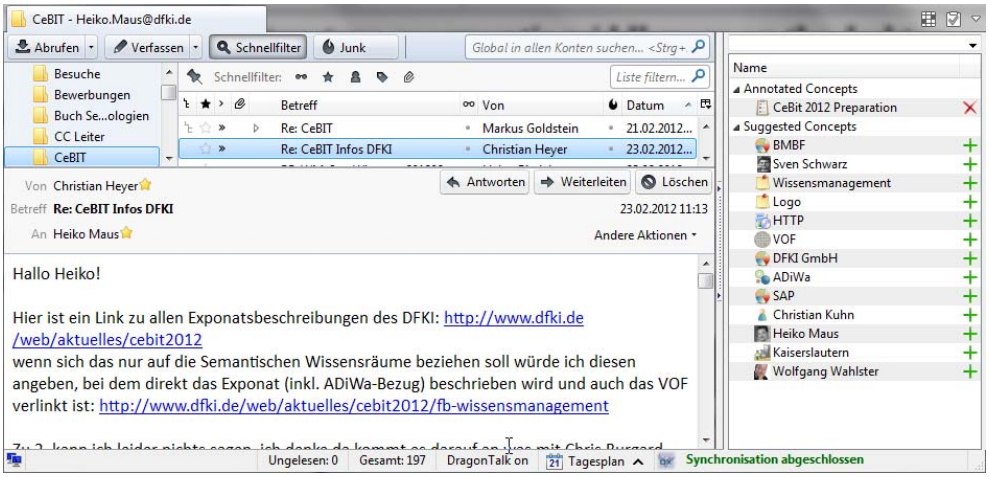

Fig. 3: PIMO sidebar in Mozilla Thunderbird with one annotated thing (a task for a trade fair) and suggesting further things based on analysis of the full text. Annotating a thing allows then to directly open the email from whatever application the thing is accessed.

3. PIMO-enabled applications for knowledge work

From our longstanding research in this area – with questions tackling the PIMO, value-added semantic services, and several implementations of a Semantic Desktop (EPOS, gnowsis, Nepomuk, Refinder[6] from our spin-off gnowsis.com, as well as the PIMO ecosystem presented here) – we see that the main challenge is to face one of the main hindrances in (personal) knowledge management: the individual effort to be invested for getting benefits out of the system. This divides into the individual effort for the ramp-up (how fast can the user start to work beneficially with such a system), into the effort for actually constantly using it, and into effort required for maintaining the knowledge base. And finally, the cognitive effort required for the user to understand the PIMO and to work with it. In the following, we want to address several foci of our research to reduce this effort.

[5] http://dragontalk.opendfki.de/

[6] http://getrefinder.com

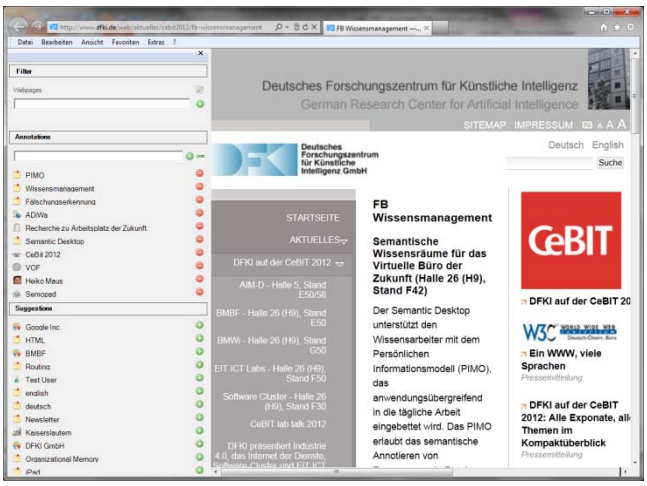

Fig. 4: PIMO sidebar embedded in MS Internet Explorer

Availability in daily activities: First of all, for today's knowledge workers, the PIMO has to be ubiquitous for them and must embrace the information items of their knowledge work. A single Semantic Desktop application alone would require bringing information items to this central place, requiring effort and risking to create yet another knowledge base isolated from the information items in the applications. Now, the previously introduced semantic middleware allows to use the PIMO also in standard applications as long as they can be extended with code using the PIMO Service API. Our current focus is on office applications such as e-mail clients (Mozilla Thunderbird as in Fig. 3 and MS Outlook SmartOffice-Plugin), web-browsers (MS Internet Explorer (see Fig. 4) and Mozilla Firefox), personal information management tools (Nepomuk Personal Semantic Workbench [5]), and task management tools (ConTask [9]). Basic functionality of all plug-ins is to search and access things and to annotate resources with things directly from within the application (see Fig. 3 for an annotated e-mail). That means, users now have the benefit of interconnecting resources from and within different (native) applications using the same vocabulary.

Also of importance for the user is the ease of file handling on his desktop, because the standard office applications rely on files to store documents. There is a lot of structuring done by the user reflecting his mental model from choosing names for files and folders to clustering files into folders up to the folder hierarchy. Handling files on the desktop is still an essential activity in today's knowledge work. Therefore, we also embedded PIMO access into the Windows File Explorer (see Fig. 5, realised as a sidebar plug-in and a namespace extension) with the possibility to easily annotate files and folders with things (and thus, also rebirth files in the PIMO, or to create things out of folders). The sidebar also shows the annotated concepts for selected files and folders and allows a concept-based search/filter for files or folders using PIMO things. That way, the PIMO provides a means for navigation directly in the file explorer itself.

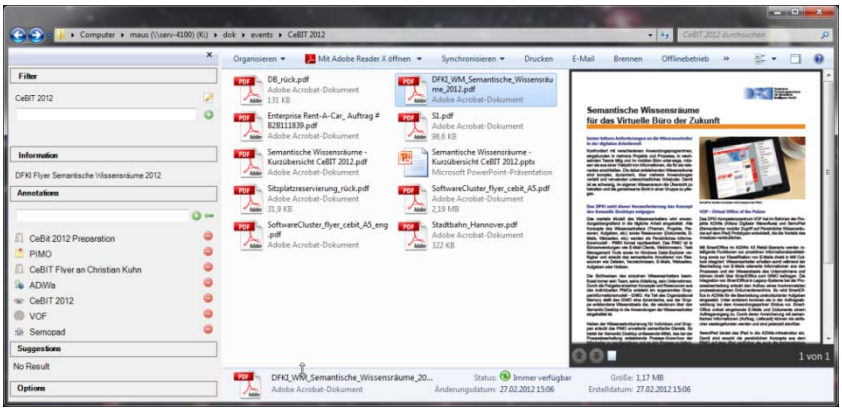

Fig. 5: Semantic File Explorer: PIMO vocabulary embedded in MS File Explorer. Filter Event "CeBIT 2012" is set, a flyer is selected and annotated things are listed

Mental model for knowledge work: To get the most benefit of the PIMO for knowledge work, the relevant part of the user's mental model has to be reflected in the PIMO. We addressed this in [17] by populating the PIMO with the user's native structures (folders-hierarchies in file system, e-mail folders, etc.) found on his desktop as digital footprints of his mental model [17]. A recent approach is bootstrapping a PIMO from a user's e-mail by crawling his e-mail account and identifying PIMO-relevant things [12]. More things, resources and relationships between them are then gathered by the PIMO enhanced daily usage of office applications. After a few days, the PIMO will already provide the network of concepts and resources centered on a user's tasks. Although the formal representation of the PIMO allows for a rich semantic modeling of concepts and their relationships, the study of several PIMOs and their usage over time in [6,16] showed that for most users it is sufficient to see that things are connected and there is no need for a more specific semantic relationship between things for their purposes. Again, to reduce the cognitive effort for users, we apply tagging as an easy-to-use functionality that helps users to quickly weave their personal knowledge space without forcing decisions about the actual semantic relationship. We successfully applied this also in other situations, e.g., tagging tasks in agile knowledge workflows with concepts in the TaskNavigator [14]. However, we still stick to semantic relations, but focus on automatic predictions of relationship types. Dedicated applications can decide automatically on the right (domain specific) semantic relationships. This works in the task management domain (subtasks, executors, used resources), or in a tool for meeting protocols (attendees, agenda). Furthermore, some metadata can be sensed or observed and stored automatically, e.g., in the mobile scenario we can use the sensors of a mobile device (time, location), or, user observation software can observe user behaviour and resource usage automatically (opened files, visited folders or web pages; see [15]).

Direct benefits from the PIMO: Again, derived from the challenge of a satisfactory return on invested effort for the user, direct benefits for the user need to be established. We accomplish this by providing an ecosystem of tools and services using the PIMO and supporting the user's daily requirements for personal knowledge management. These tools and plug-ins are designed to provide direct benefits for the user:

- Tools for personal knowledge management allowing to acquire, create, classify, and organize information items using the PIMO such as the SemanticFileExplorer in Fig. 5.
- Information retrieval components providing fast access to information items via common keyword search, associative search, semantic search, and combinations of them based on the semantic annotations using the PIMO and the full text of the information items.
- Ubiquitous PIMO: One vocabulary throughout all applications, to be accessed via plug-ins to organize and interconnect information items. The PIMO is available everywhere for ease of access and for annotating information items on the fly in order to constantly evolve the PIMO. Further assistance and coverage is achieved by an automatic interpretation and analysis of information items in focus and proactively suggesting annotations and related things (such as the suggesting relevant things for an opened email in Fig. 3).
- Allowing to use information from the group and share information with the group easily (see next section).

Furthermore, the availability of a machine-readable vocabulary of the user's mental model allows easy development of new added-value services supporting the user in his personal knowledge management.. Examples are the semantic search on the user's personal knowledge space [7], ontology-based information extraction using the PIMO as background knowledge [8], personal image collections using the PIMO [10], semantic annotation embedded in office documents [11], and personal trend recognition in the PIMO [18]. From a research perspective, the PIMO allows to get a more precise understanding of the user's current activities and topics. Therefore, in our user observation for context-aware services in [15,17], the PIMO serves as a vocabulary for identifying and formally representing the user's interests, current activities, and the context he is in. With this infrastructure as a basis, we developed context-aware services providing proactive information delivery based on the user's context [15]. Moreover, an agile personal task management embedded in the Semantic Desktop (ConTask, see [9]) uses the PIMO for organizing a user's tasks and their contents, applying task and context oriented proactive information delivery. It observes the user's actions to connect them to the respective tasks and identifies task switches of the user to keep the proactive information delivery always aligned to the user's current work. All these plug-ins and tools are part of the presented PIMO ecosystem providing direct support of activities in the user's knowledge work.

Extending the Personal Knowledge Space to the Group: The views of individual knowledge workers also influence their team, their department, their company. By allowing to share things and resources from individual PIMOs a

Group Knowledge Space is evolving. As this is based on the PIMO, we refer to this as the *GIMO* – the *Group Information Model*. We apply a bottom-up approach, allowing users to easily share things and resources to the group as well as being able to directly use things and resources from the group for the individual work. Again, to infuse the GIMO to the group members and keeping a momentum of sharing and using of information, several value added services are required. In our current ecosystem we use the GIMO also for concept proposals for annotating (personal) information elements, provide RSS feeds for activities on the GIMO (shared concepts and resources, changes, annotations), and being a source for proactive information delivery, e.g., in agile knowledge workflows as in TaskNavigator [14]. Therefore, the GIMO is a dynamic knowledge base evolving from individuals interacting in the group and resembling a part of the Organisational Memory. By means of the Semantic Desktop, it is in turn still available within the knowledge worker's applications, and thus, group knowledge is always available during the individual knowledge work.

With the evolving GIMO from individual PIMOs, we see several research questions arising such as maintenance issues, community building, as well as cooperative domain ontology creation from individual PIMOs (see also earlier work in [13]).

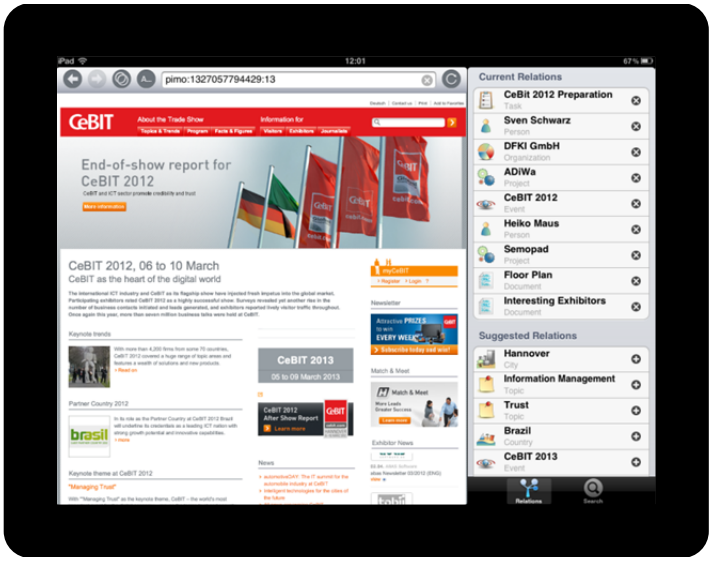

Fig. 6: Semopad: A web page with annotated and suggested things from the PIMO on the iPad

Mobile access to the Personal Knowledge Space: In the Semopad project[7] we bring the power of the Semantic Desktop paradigm to mobile devices. As Apple's iPad is a favourite and often used tablet device, we realized an iOS App for the iPad 2 (see Fig. 6). The software is still being developed and is not yet available via the App Store. Although accesses the PIMO and provides an interface similar to the sidebars we implemented so far such as for Mozilla Firefox. Users can browse the internet and get to see existing and suggested annotations next to the web page at hand.

The related PIMO concepts can be viewed and browsed. For searching things in the PIMO, Semopad realized a facetted search (see **Fig. 7**). That way, the Semopad App can be used in the mobile setting to quickly annotated found web pages, as well as, to efficiently look up information in the PIMO. As the user does not want to enter much text, the project focuses on use cases with minimal required interaction, particularly minimal taps on the (virtual) keyboard. Therefore, the PIMO provides an interactive search for things, requiring only the beginning or part of the concept's name and showing the options while the user is typing. This feature is, of course, also available in the desktop GUIs, but in the App it is crucial.

The App can also be run when being offline and uses a sophisticated caching mechanism which allows high performance in slow network scenarios, e.g., when having to use a bad or slow UMTS connection: The App always shows the cached information first while, in the background, the PIMO is queried for recent changes and respective updates are requested. As soon as an updated version of a displayed thing is available, the display is corrected accordingly. So, when a Semopad user is inspecting a thing and a remote colleague (e.g., in the company building) is adding a new inter-

[7] https://www2.dfki.de/intranet/research/projects/Project_674

esting resource to the thing, the Semopad user sees the changes right away, respectively, as soon as he is online again (after lacking internet connection).

Fig. 7: Facetted search in Semopad and viewing a PIMO concept

4. Conclusions

The paper presented a new ecosystem of applications and plug-ins for the Semantic Desktop approach. The goal is to provide knowledge worker with one vocabulary – the PIMO – across application borders. Furthermore, starting from previous research and systems, the individual PIMOs are now embedded in a GIMO that will give the opportunity to follow a bottom-up approach for knowledge management in groups. We presented several new prototypes and referenced other systems belonging to the ecosystem and also gave outlook to further research that will be addressed in future work. We expect more insights into the question on how the ubiquitous availability will influence the PIMO build up, its usage, and the user benefit in knowledge work as well as the influence to cooperative construction of the GIMO. This insight will be especially complemented by the experience from our spin-off gnowsis.com with their product Refinder which follows a similar approach with individual PIMOs and sharing to friends.

Acknowledgements

Parts of this work were supported by the projects ADiWa (funded by German Federal Ministry of Education and Research - 01IA08006) and Semopad (funded by Stiftung Rheinland-Pfalz für Innovation - 961-386261/1001).

References

1. Sauermann, L. Bernardi, A., Dengel A.: Overview and Outlook on the Semantic Desktop. Proc. of the First Semantic Desktop Workshop at the ISWC Conference 2005, 2005
2. Dengel, A.: Knowledge Technologies for the Social Semantic Desktop. Knowledge Science, Engineering and Management. 2nd Int'l Conference (KSEM-2007), Melbourne, Australia, Springer, LNCS 4798, pp. 2-9, 2007
3. Sauermann, L., van Elst, L., Dengel A.: PIMO - A Framework for Representing Personal Information Models. Proceedings of I-SEMANTICS Conference, 5-7 September 2007, Graz, pp. 270-277, J.UCS, Know-Center, Austria, 2007
4. Dengel, A.: Six Thousand Words about Multi-Perspective Personal Document Management. Proceedings EDM, IEEE Int'l Workshop on the Electronic Document Management in an Enterprise Computing Environment, Hong Kong, 2006
5. Grimnes, G. A., Sauermann, L., Bernardi, A.: The Personal Knowledge Workbench of the NEPOMUK Semantic Desktop. The Semantic Web: Research and Applications. 6th European Semantic Web Conference, ESWC 2009 Heraklion, Crete, Greece, Proceedings, Springer, LNCS 5554, pp. 836-840, 2009

6. Sauermann, L., Heim, D.: Evaluating long-term use of the Gnowsis Semantic Desktop for PIM. The Semantic Web. 7th International Semantic Web Conference (ISWC-2008), 7th Int'l Semantic Web Conference, October 26-30, Karlsruhe, Germany, Springer, LNCS 5318, pp. 467-482, 2008
7. Schumacher, K., Sintek, M., Sauermann, L.: Combining Fact and Document Retrieval with Spreading Activation for Semantic Desktop Search. Proc. of the 5th European Semantic Web Conference, Springer, LNCS 5021, pp. 569-583, 2008
8. Grimnes, G. A., Adrian, B., Schwarz, S., Maus, H., Schumacher, K., Sauermann, L.: Semantic Desktop for the End-User. i-com, Oldenbourg Verlag, 8, pp. 25-32, 2009
9. Maus, H., Schwarz, S., Haas, J., Dengel, A.: CONTASK: Context-Sensitive Task Assistance in the Semantic Desktop. Enterprise Information Systems. 12th Int'l Conference, ICEIS 2010, Funchal-Madeira, Portugal, June 8-12, 2010, Revised Selected Papers, Springer, LNBIP 73, pp. 177-192, 2011
10. Klinkigt, M., Kise, K., Maus, H., Dengel, A.: Semantic Retrieval of Images by Learning from Wikipedia. Knowledge-Based and Intelligent Information and Engineering Systems, Springer, LNCS 6884, pp. 212-221, 2011
11. Rostanin, O., El Agroudy, P.: SWord: Semantic Annotations Revisited. Knowlege-Based and Intelligent Information and Engineering Systems, Springer, LNCS 6882, pp. 410-419
12. Schwarz, S., Marmann, F., Maus, H.: Extracting Personal Concepts from Users' Emails to Initialize Their Personal Information Models. Knowlege-Based and Intelligent Information and Engineering Systems, Springer, LNCS 6882, pp. 430-439
13. Aschoff, F.-A., Schmalhofer, F., van Elst, L.: Knowledge Mediation: A Procedure for the Cooperative Construction of Domain Ontologies. In Engineering Knowledge in the Age of the Semantic Web: Proc. of the 14th Int'l Conference, EKAW 2004, Whittlebury Hall, UK, October 5-8, 2004, Springer, LNAI 3257, 2004
14. Rostanin, O., Maus, H., Suzuki, T., Maeda K.: Using lightweight knowledge modeling to improve proactive information delivery. Proc. of the 2nd Int'l Conference on Agents and Artificial Intelligence (ICAART-2010), January 22-24, Valencia, Spain, INSTICC, Springer, LNCS 6276, pp. 611-614, 2010
15. Schwarz, S.: Context-Awareness and Context-Sensitive Interfaces for Knowledge Work Support. PhD-Thesis. University of Kaiserslautern, Verlag Dr. Hut, 2010
16. Sauermann, L.: The Gnowsis Semantic Desktop approach to Personal Information Management. PhD-Thesis. University of Kaiserslautern, Verlag dissertation.de, 2009
17. Sauermann, L., Dengel, A., van Elst, L., Lauer, A., Maus, H., Schwarz, S.: Personalization in the EPOS project. Proceedings of the Semantic Web Personalization Workshop at the ESWC 2006 Conference, Budva, pp. 42-52, 2006
18. Kubo, S., Tsuji, H., Maus, H., Dengel, A.: Visualizing Personal Trend on Gnowsis Semantic Desktop. IEEE International Conference on Systems, Man, and Cybernetics (SMC 2008), IEEE, pp. 2765-2771, 2008