

Emergent Semantics in Personalized Multimedia Content

Ansgar Scherp
Research Institute OFFIS
Oldenburg, Germany
ansgar.scherp@offis.de

Susanne Boll
University of Oldenburg
Oldenburg, Germany
{susanne.boll|holger.cremer}@informatik.uni-oldenburg.de

Holger Cremer

Abstract: Authoring of personalized multimedia content can be considered as a process consisting of selecting, composing, and assembling media elements into coherent multimedia presentations that meet the user's or user group's preferences, interests, current situation, and environment. In the approaches we find today, media items and semantically rich metadata information are used for the selection and composition task. However, most valuable semantics for the media elements and the resulting multimedia content that emerge with and in the authoring process are not considered any further. This means that the effort for semantically enriching media content comes to a sudden halt in the created multimedia document – which is very unfortunate. In this paper, we propose with the SemanticMM4U framework an integrated approach for deriving and exploiting multimedia semantics that arise with and from the creation of personalized multimedia content and make it available for further use and applications. In this approach, not only the metadata are considered that are semantically evolving from the media elements of the newly created presentation. Also the actual usage of media elements for the authoring can emerge new semantics of the single media elements employed. In the application domain of authoring of digital photo albums we show where and how semantics emerge in the authoring process and how these contribute to an even richer media content pool, both for single media and also for the composition and its later use.

1 Introduction

Authoring of personalized multimedia content is considered as a process of creating and delivering multimedia presentations that meet the user's or user group's preferences and interests as well as current situation and environment. Typically, for creating such personalized multimedia presentations, the following tasks are conducted: first, appropriate media elements are selected from the multimedia storages on the basis of the user profile information and current contextual situation. Here, (semantically rich) metadata information associated to the media elements is exploited to determine the most relevant media elements. The selected media elements are then composed and assembled in time and space into a coherent multimedia presentation, again by exploiting the media element's metadata and user's profile and context information. Finally, the presentation is deployed to the end device of the user for actual rendering, which can be a personal computer but also a mobile device.

Different research groups worldwide have conducted research in the field of creating such personalized multimedia content such as the research institute CWI in the Netherlands (see <http://www.cwi.nl/>) and INRIA Rhône-Alpes in France (see <http://www.inrialpes.fr/>). After almost a decade of research on this issue, we are pleased today to observe that there exist many scientific approaches as well as industrial solutions that provide personalized content to the user.

In recent years, it is becoming increasingly clear that multimedia information does not have unique semantics, but exhibits multiple semantics depending on its context and usage (3rd special Workshop on Multimedia Semantics, 2005). However, these semantics are currently lost in the multimedia content creation process. Although the approaches we find today exploit semantically rich metadata for the multimedia selection and composition task as described above, this highly valuable source of information is currently not considered any further: the actually created multimedia presentations carry none or only a small piece of the metadata exploited for the content creation task before. This means that the effort to enhance and enrich media content with semantically rich descriptions all of a sudden comes to a halt in authoring of personalized multimedia content. What we lose is not only the combined semantics of the multimedia document as a combination of the semantics of the single media items used. Also the semantics that arise from the composition itself are ignored. And finally, the use of media objects for and within a multimedia composition is typically not reflected in the media object's metadata.

In this paper, we propose with the SemanticMM4U framework an integrated approach for deriving and exploiting multimedia semantics from multimedia content authoring. This approach not only employs the semantics and metadata of the media elements for the selection and composition task.

It further aims at exploiting the existing semantics and metadata information for deriving higher-level semantics that emerges by combining media elements into new multimedia presentations and to bring this semantics to the end user.

In Section 2, we present the current state-of-the-art in the fields of multimedia personalization and emergent semantics for multimedia. In Section 3, we present our approach of a SemanticMM4U framework for integrating emergent semantics into the multimedia personalization process. The current state of the implementation of the SemanticMM4U framework is described in Section 4 and the evaluation of our approach in a concrete application scenario is presented in Section 5. Finally, we conclude our paper with Section 6.

2 Related Work

As the proposed SemanticMM4U framework combines the challenging research fields of multimedia content personalization and emergent semantics in multimedia, the related work in both fields is presented in the following. We begin with the state-of-the-art in personalized multimedia content creation and then present the research conducted in the field of emergent semantics for multimedia content.

State of the art in personalized multimedia content creation: On the pathway to a dynamic creation of personalized multimedia content, we find interesting work in the area of adaptive hypermedia systems (De Bra et al., 2002; Wu et al., 2001; Brusilovsky, 1996). The adaptive hypermedia system AHA! (De Bra et al., 2003) is a prominent example. However, though these and further approaches integrate media elements in their adaptive hypermedia presentations, synchronized multimedia presentations are not in their focus. A very interesting research approach towards the dynamic generation of multimedia presentations is the Cuypers system (van Ossenbruggen et al., 2003; Geurts et al., 2001; van Ossenbruggen et al., 2000). It uses constraints for describing the intended multimedia programming and logic programming for the generation of a multimedia document. Within the Opéra project, a generic architecture for the automated construction of multimedia presentations based on transformation sheets and constraints is developed (Bes et al., 2001). This work is continued within the WAM project with the focus on a negotiation and adaptation architecture for mobile multimedia services (Lemlouma & Layaïda, 2004). These and other existing research solutions typically use declarative languages, such as rules, constraints, style sheets and others, to express the dynamic multimedia content creation. However, we observe that the existing approaches find their limits in regard of required effort, applicability and practicability, as well as their power of expressiveness. Many of these systems and approaches for content personalization are targeted at a specific application domain in which they provide a very specific content personalization task. Whenever a complex and application-specific personalization generation task is required, the systems find their limit and need additional programming to solve the problem.

With MM4U, we developed a software framework that provides generic support for the development of personalized multimedia applications (Scherp & Boll, 2005a+c). It relieves application developers from the general tasks in the context of multimedia content personalization and lets them concentrate on the application domain-specific tasks. The framework does not re-invent multimedia content adaptation but is targeted at incorporating and embedding existing research approaches in the field and allows to be extended by domain and application-specific solutions.

State of the art in emergent semantics for multimedia content: Semantics is a discipline developed by Alfred Korzybski (Schuchardt & Kending, 2005) that proposes to improve human behavioral responses through a more critical use of words and symbols (Jain, 2004). In general, this field studies relationship between signs and symbols and what they represent for the purpose and utilization of communication. In regard of emergent semantics for multimedia information, the opinion reinforced in recent years that multimedia information does not have a unique semantics, but exhibits multiple semantics which depend on context and use (3rd special Workshop on Multimedia Semantics, 2005) and which changes over time. Emergent semantics is subject of research conducted from different sides. For example, the work done by (Grosky et al., 2002+2004) considers the semantics of Web pages and the interrelationship of the different documents. Here, Web server logs

are analyzed to derive emergent semantics of the Web pages and the change of their semantics over time and usage. Another approach to emergent semantics considers audiovisual media content. This is pursued, e.g., by (Nack, 2005; Nack & Putz, 2001) for video editing in the domain of news coverage and by (Santini et al., 2001) analyzing the interrelationship and the semantic that emerges through interaction in image databases.

However, also with the composition of media elements into a coherent multimedia presentation new semantics for the created content can be derived from the used media element's metadata. This can take place at several places of the multimedia composition step. However, current approaches either derive emergent semantics prior or post to the composition of the media elements into the presentation. For example, the approach by (Little et al., 2002) first derives emergent semantics from the media element's Dublin Core metadata and uses this semantics to determine a rhetorical structure. Then, the created rhetorical structure is mapped to a temporal course and spatial layout of a multimedia presentation. Thus, the emergent semantics and the created multimedia content are separated in two different representation models. Once, the rhetorical structure is exploited to create the multimedia content, the derived semantics is not further used and thrown away. An example for deriving semantics post to the multimedia composition is presented by (Ding et al., 2005). Here, Flash movies in the binary format are analyzed to derive semantics. However, this analysis process of binary multimedia content is very complicated and time-consuming as each scene of the flash presentation needs to be individually considered and analyzed. The result is a categorization of the Flash movie into a couple of pre-defined categories. The original metadata and semantics of the media elements used for the Flash presentations is not used for this emergent semantics approach.

However, research in regard of emergent semantics for multimedia content during the authoring process is still in its infancies. Although, it is recognized that semantics needs to be handled in every step of the creation process of multimedia content (Hardman, 2005), today's tools tend to concentrate on one step of the multimedia content creation process only and ignore the requirements in regard of semantic information needed by a subsequent processing and usage. What is missing is an approach that derives emergent semantics from the media elements during the composition and assembly task and makes this semantics available for the multimedia presentation consumption as well as for future multimedia authoring tasks. For it, enhanced multimedia document models such as (Westermann et al., 2005) need to be employed, which carry semantically rich metadata. However, these models do not serve as basis for bringing the semantics into the actual multimedia presentation formats created and delivered to the user's or user group's needs. Here, appropriate mechanisms are required for bringing the content represented by such a sophisticated document model into today's multimedia presentation formats such as SMIL (Bulterman et al., 2005) and to exploit the metadata description functionality of these formats, e.g., SMIL's Metainformation Module, to bring the semantics to the user.

3 The SemanticMM4U Approach

Semantics for media and multimedia content is of high potential and benefit when made explicit and machine processable as it allows for a better management, search, distribution, and usage of multimedia information. However, as mentioned above, the semantics that emerge from the creation process of personalized multimedia content is currently left unconsidered. Figure 1 illustrates the creation of multimedia presentations and where we today find an *emergent semantics barrier*. Media content such as image, video and audio is typically associated with technical and semantic metadata. These metadata are used for the selection and assembly of a multimedia presentation. The semantics of the actual creation, however, only reside in the assembly phase and are not made explicit to other steps in the chain from media to multimedia presentation. For example, the parallel presentation of a set of media elements can semantically mean that they belong together. This information should be stored in the media pool for future authoring processes. Also the main keywords of the different media elements used in a new presentation could be summarized in a head of the delivered presentation for search. What we observe today is a multimedia presentation semantics barrier within the authoring process of personalized multimedia content where the (emergent) semantics of the created multimedia presentation is lost.

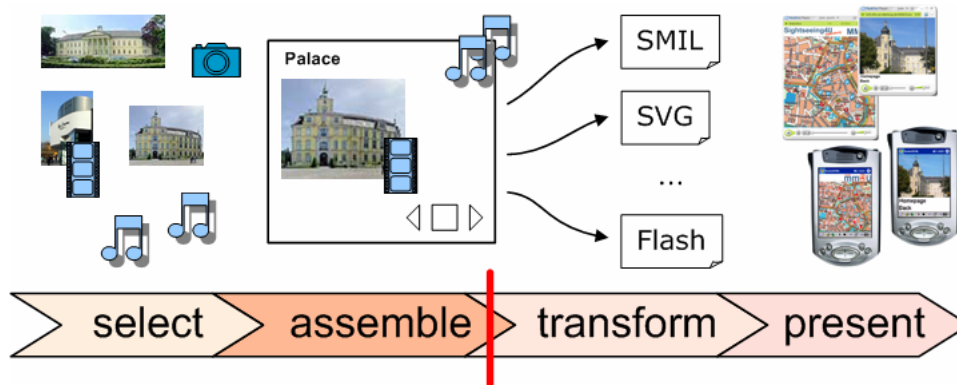


Figure 1: Authoring process of personalized multimedia content and the multimedia presentation semantics barrier

With the SemanticMM4U framework, we propose an integrated approach for deriving and exploiting multimedia semantics in the multimedia authoring process and make these emergent semantics explicit and available for other processing and use. In more detail, the SemanticMM4U approach provides for two central aspects in emergent semantics for personalized multimedia content. Following the process of creating and delivering multimedia presentations, we aim to make the semantics that arise from the authoring process explicit. On the other hand, the use of media elements in the presentation allows us to get a better understanding of the media element itself and also capture these semantics. Figure 2 illustrates the two directions where emergent semantics of the authoring are made explicit for further use. These semantics derivation directions are explained in more detail in the following.




Figure 2: Emergent semantics in the SemanticMM4U approach

Media usage semantics: As depicted in Figure 2, the multimedia presentation’s semantics derived from the single media elements metadata used for the composition and assembly task also affects the single media elements semantics. The fact that a user selects a media element together with other media elements in a multimedia presentation allows us to learn from the usage of the media elements in the presentation about the semantics of the single media elements employed. For example, if a set of media elements are used together in the same section of a presentation this can mean that they all belong to the same semantic concept, such as “Island Norderney”. If media elements come with different semantic descriptions the joint usage may allow deriving or refining the single media element’s metadata. So the media usage brings emergent semantics which arise from the authoring process. Here, new metadata can be derived for the used media elements on basis of their role within the created multimedia presentation as well as the targeted user or user group of the presentation. This emerging metadata for the single media elements can be reflected back into the media management.

Multimedia authoring semantics: With the SemanticMM4U approach, emergent semantics are derived from the media elements’ metadata used for the multimedia presentation. This means that the created presentation at least needs to carry the used media element’s semantics. For example, semantic keywords of the media elements which have been used for selecting the media element for the presentation can also be explicitly stored and made available in a metadata description of the composed presentation. Additionally, the composition process itself creates and derives new metadata. This includes, e.g., the structure and temporal course of the presentation or the layout and media types used.

For the delivery and consumption of the created presentations, the existing metadata standards of the different multimedia presentation formats will be employed to enhance the created content by the derived semantics. This will make the presentation's emergent semantics explicit. Consequently, the created multimedia presentations will be made accessible and processable by future multimedia retrieval systems and by this provide for a better management, search, distribution, and usage of such content. Within the SemanticMM4U framework, new semantics will be derived for the created multimedia presentation from the used media elements metadata. This will be achieved by employing (semi-)automatic multimedia semantics derivation methods and techniques. These include among others semantics derivation by employing rules, knowledge-bases such as taxonomies and thesauri, or by plain programming the semantics derivation functionality. The SemanticMM4U framework aims at abstracting from these different semantics derivation methods and techniques by providing appropriate interfaces and data models. Thus, different approaches for semantics derivation can be employed and combined by the framework. The created multimedia presentations are considered to be new media items with accompanying metadata. Consequently, they can be stored and made accessible via the media stores for use in other presentations, e.g., as MPEG-21 Digital Items (Burnett et al., 2003).

4 Implementation Issues

The SemanticMM4U framework bases on our previous work, the development of the component framework MM4U for the dynamic authoring of personalized multimedia content (Scherp & Boll, 2005a+c). In order to support the different tasks of the general multimedia personalization process in the MM4U framework, we developed a layered architecture. Each layer provides modular support for the different tasks of the multimedia personalization process. The bottom layers provide access to user profile information and media data, followed by two layers for composition of the multimedia presentation in an internal multimedia content representation model and its later transformation into the concrete presentation output formats. A top layer realizes the rendering and display of the multimedia presentation on the end device. To be most flexible in regard of the different requirements of concrete personalized multimedia applications, the framework layers allow for extending the functionality of the MM4U framework. In the work presented in this paper, we aim at extending the MM4U framework as depicted in Figure 3 by considering the aspect of emergent semantics during the (semi-)automatic authoring process of personalized multimedia content. Therefore, the existing MM4U framework will be extended towards a SemanticMM4U framework.

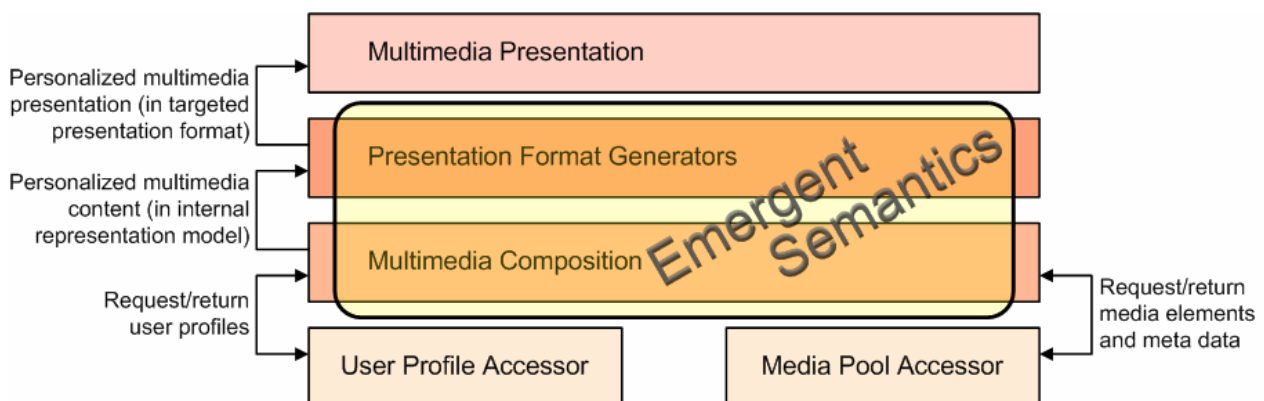


Figure 3: The architecture of the SemanticMM4U framework

Following the multimedia authoring process as illustrated in Figure 1, the new aspects that SemanticMM4U brings to the MM4U framework are touching the layers of multimedia composition and presentation creation. In the following, we elaborate in more detail where and how the emergent semantics influence and alter the different framework components:

- **Impact of emergent semantics for the composition layer:** The current multimedia composition model and metadata concept of the MM4U framework captures the central aspects of multimedia modeling, which are the temporal course, spatial layout, and interaction possibilities of the presentation with the user. With the SemanticMM4U framework, we extend this model by

providing sophisticated support for capturing and processing metadata and semantics in the composition model. Unlike other models (see related work), our internal model provides support for both challenges: composition and multimedia semantics.

- **Impact of emergent semantics for the presentation format generators:** To use the (emergent) semantics post to the multimedia composition step, we need to exploit the semantics of the newly created multimedia content within the transformation step to deploy it with the multimedia presentations in the final presentation formats. We decided to use a RDF-based profile for the presentation's metadata and semantics. We target at using MPEG-7 descriptors as a basis. MPEG-7 defines a comprehensive set of low-level metadata as well as some general high-level metadata. In addition, it provides support for defining domain-specific and application-specific metadata (Manjunath et al., 2002). The semantics, emerged through the assemble process, are saved to the generated presentation.
- **Impact of emergent semantics for user profile and media management:** In the MM4U framework, the user profile accessor and the media accessor are currently used to retrieve the necessary and relevant media information. With the semantics that emerge in the authoring process now new and refined semantics are propagated back to the profile and media store. For example, the media use and presentation structure can be learned and stored in a user profile; the use of a title for a group of images can be stored in the metadata of the single photos.

5 Evaluation and Application Scenario

The SemanticMM4U framework is evaluated by and applied for the development of demonstrator applications for the domain of (semi-)automatic creation of personalized multimedia content. These demonstrator applications will prove the applicability of the framework and provide for an improved retrieval of personalized multimedia content.

The first evaluation of our SemanticMM4U framework is currently conducted in the domain of personalized photo albums. Here, the SemanticMM4U framework is applied within an existing smart authoring tool for the context-driven and content-driven creation of multimedia presentations, called xSMART (Scherp & Boll, 2005b). The xSMART authoring tool integrates the targeted user context into the different authoring steps and exploits this context to guide the author through the content authoring process. The design of xSMART allows that it can be extended and customized to the requirements of a specific domain by domain-specific wizards. These wizards realize the user interface that meets best the domain-specific requirements and effectively supports the domain experts in creating their content targeted at a specific user context.

A concrete wizard we developed for our xSMART authoring tool provides support for the personalized authoring of photo albums. This wizard allows the users to define that only pictures taken during summer time are to be selected and only those, with a good exposure as well as sharpness. For the composition task, the wizard allows for determining the targeted (mobile) end device and allows for selecting different presentation styles for arranging the pictures.

We selected the photo-album scenario for evaluating our SemanticMM4U approach as it is well-known and high expertise available. We expect fast and promising results for emergent semantics in this domain. A screenshot of the current photo-album wizard with support for emergent semantics is depicted in Figure 4. The figure shows a page of a personalized photo album generated by our smart authoring wizard. The composition of a personalized photo album is conducted in several steps: the user starts with the domain-specific photo album wizard, which creates a first presentation on the basis of user selected parameters (e.g., image folders, image options). Then, the user can manipulate the single pages, add text, and change the layout of the images on the pages.

With the SemanticMM4U framework, we aim at deriving semantics during the multimedia composition and assembly step. The example screenshot of a photo album page illustrates where we can derive semantics for the presentation and the images. From the selected style (one big image centered, rest aside), we can derive that the centered image is the most important one of the page. Thus the media item can be "labeled" with a higher relevance tag than the others. From the organization of the five images on the same page, we can deduce that all images have the same topic the user has

entered as “Our Vacation on Island Norderney”. Therefore, the images can be labeled with the same description. Also, available metadata for one image can be used and checked if it might be valid for the other images as well. Here for example, the image on the upper left comes from a friend’s camera and has a GPS position. The organization of the images on the same page and the fact that the images are captured around the same time allows deducing the same GPS position from one photo for the other four images as well. The usage and manipulation of single photos as well as the arrangement of the photos on the pages allows learning and refining the media metadata by exploiting the user interaction during the authoring process.

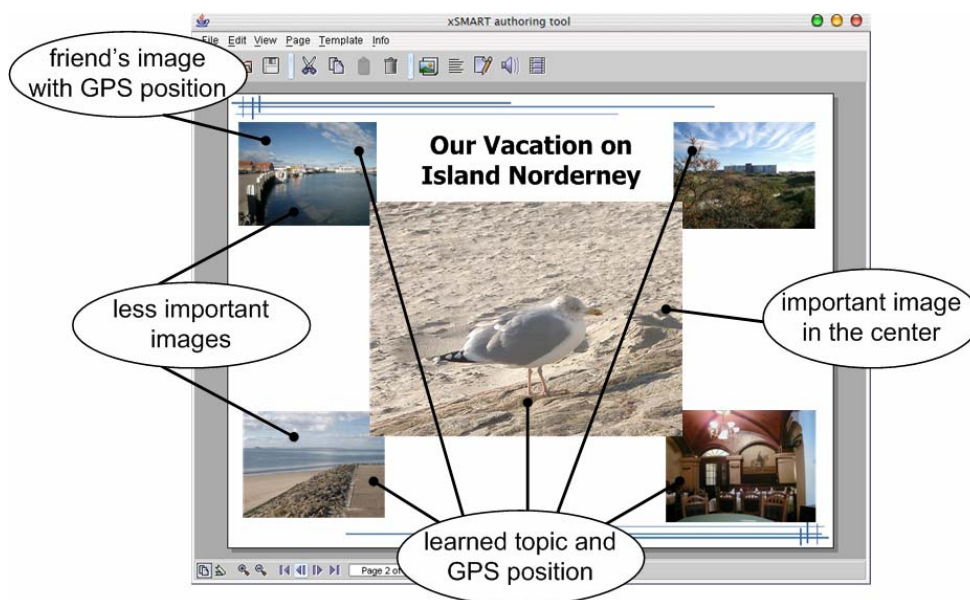


Figure 4: Emergent semantics in our authoring tool xSMART and authoring wizard for personalized photo-albums

In the direction of the presentation, the delivered media presentation, e.g., a SMIL presentation, includes an RDF description that carries the single media element’s metadata such as time, position and also any manual annotation that some of the images have. Concerning the structure of the presentation, the different pages are treated as logically related chapters of the vacation. An available text on the page is used and made explicit as a topic and title for this chapter. The sequence of the pages and the storyline is also made explicit as well as information about spatial or temporal clusters and the derived potential events the photos were taken at. There is almost no limit in the information that can be emerged from the authoring process, though it is highly dependent on the application domain which information is relevant for later search, organization, and delivery of the multimedia content.

6 Conclusion

In this work, we presented a novel approach for emergent semantics in the process of authoring personalized multimedia content. The proposed SemanticMM4U framework not only employs the media element’s metadata for the composition and assembly task, but extends the state-of-the-art by deriving emergent semantics when using media elements for authoring new multimedia content targeted at a specific user or user group and by enhancing the created presentation with this semantics. The SemanticMM4U approach combines the challenging research areas of multimedia content personalization and emergent semantics for multimedia. The presented combination of multimedia content personalization and emergent semantics is a clear consequence from the current state-of-the-art to gain a better understanding of the semantics of multimedia information. As semantics is always dependent on the individual’s role and perception, exploiting the knowledge gained about how to serve the individual user’s or user group’s needs and preferences will help to reach this understanding. On this basis, we will then be able to draw conclusions for the emergent semantics of multimedia presentations created for an individual user or user group.

Making multimedia semantics explicit throughout the multimedia authoring process provides for a better management, search, distribution, and usage of multimedia information. Search engines will be able to gather such presentations and process the provided semantics in order to provide their users in the Internet a better search and retrieval of the content. Consequently, the information needs and requirements of the individual users can be better fulfilled. The proposed approach for exploiting metadata information, deriving higher-level semantics, and making this emergent semantics explicit provides for a better accessibility of the created personalized multimedia content.

Acknowledgements: We thank Prof. Ramesh Jain and Dr. Utz Westermann for their very valuable input and comments for the proposed SemanticMM4U approach.

References

- Bulterman, D., Grassel, G., Jansen, J., Koivisto, A., Layaïda, N., Michel, T., Mullender, S., & Zucker, D. (December, 2005). Synchronized Multimedia Integration Language. Retrieved March 6, 2006, from <http://www.w3.org/TR/SMIL2/>
- Bes, F., Jourdan, M., & Khantache, F. (2001). A Generic Architecture for Automated Construction of Multimedia Presentations. Intl. Conf. on Multimedia Modeling, Amsterdam, Netherlands.
- Brusilovsky, P. (1996). Methods and techniques of adaptive hypermedia. *User Modeling and User Adapted Interaction*, 6(2-3).
- Burnett, I., Van de Walle, R., Hill, K., Bormans, J., & Pereira, F. (2003). MPEG-21: Goals and Achievements. *IEEE MultiMedia*, 10(4).
- De Bra, P., Aerts, A., Berden, B., De Lange, B., Rousseau, B., Santic, T., Smits, D., & Stash, N. (2003). AHA! The Adaptive Hypermedia Architecture. Proc. of the ACM Hypertext Conf., Nottingham, UK.
- De Bra, P., Brusilovsky, P., & Conejo, R. (2002). Proc. of the Second Intl. Conf. for Adaptive Hypermedia and Adaptive Web-Based Systems, Malaga, Spain, Springer LNCS 2347.
- Ding, D., Yang, J., Li, Q., Liu, W., & Wang, L. (2005). What Can Expressive Semantics Tell: Retrieval Model for a Flash-Movie Search Engine. CIVR, LNCS 3568, pp. 123-133.
- Geurts, J., van Ossenbruggen, J., & Hardman, L. (2001). Application-specific constraints for Multimedia Presentation Generation. In Intl. Conf. on Multimedia Modeling, Amsterdam, Netherlands.
- Grosky, W.I., Sreenath, D.V., & Fotouhi, F. (2002). Emergent Semantics and the Multimedia Semantic Web, *SIGMOD Rec.*, 31(4), p. 54-58, ACM Press.
- Grosky, W.I., Sreenath, D.V., & Fotouhi, F. (2004). Deriving Emergent Semantics. Association for Information Systems, SIG on Semantic Web and IS. URL <http://www.sigsemis.org/>
- Hardman, L. (2005). Canonical Processes of Media Production, Proc. of Workshop on Multimedia for Human Communication, November, Singapore, ACM Press.
- Jain, R. (2004). Emergent Semantics and Experiential Computing. White Paper. Retrieved March 6, 2006, from <http://lsdis.cs.uga.edu/SemNSF/Jain-Position.doc>
- Little, S., Geurts, J., & Hunter, J. (2002). Dynamic Generation of Intelligent Multimedia Presentations through Semantic Inferencing, Proc. of the 6th European Conf. on Research and Advanced Technology for Digital Libraries, London, UK, Springer.
- Lemlouma, T., & Layaïda, N. (2004). Context-Aware Adaptation for Mobile Devices. IEEE International Conference on Mobile Data Management. Berkeley, CA, USA.
- Manjunath, B.S., Salembier, P., & Sikora, T. (Ed.). (2002). Introduction to MPEG-7: Multimedia Content Description Interface. Wiley.
- Nack, F. (2005). Capture and Transfer of Metadata During Video Production, Proc. of Workshop on Multimedia for Human Communication, November, Singapore, ACM Press.
- Nack, F., & Putz, W. (2001). Designing Annotation Before It's Needed, Proc. of the 9th ACM Int. Conf. on Multimedia, ACM Press, Sept./Oct. 5, Ottawa, Canada.
- Santini, S., Gupta, A., & Jain, R. (2001). Emergent Semantics through Interaction in Image Databases, *IEEE Transaction of Knowledge and Data Engineering*, 13(3).
- Scherp, A., & Boll, S. (2005a). Paving the Last Mile for Multi-Channel Multimedia Presentation Generation. Proc. of the 11th Intl. Conf. on Multimedia Modeling, Melbourne, Australia, IEEE Computer Society.
- Scherp, A., & Boll, S. (2005b). Context-driven smart authoring of multimedia content with xSMART. Proc. of the 13th annual ACM Intl. Conf. on Multimedia, Singapore, ACM Press.
- Scherp, A., & Boll, S. (2005c). MM4U: A Framework for Creating Personalized Multimedia Content. In: *Managing Multimedia Semantics*, Chapter 11, S. 246-287, U. Srinivasan and S. Nepal (Ed.), IRM Press, p. 246-287.
- Schuchardt, C., & Kending, M. (2005). Alfred Korzybski. Institute of General Semantics. Retrieved March 6, 2006, from <http://time-binding.org/institute/ak.htm>
- van Ossenbruggen, J.R., Cornelissen, F.J., Guerts, J.P.T.M., Rutledge, L.W., & Hardman, H.L. (2000). Cuypers: a semi-automatic hypermedia presentation system, Technical Report, CWI, Netherlands.
- van Ossenbruggen, J., Hardman, L., Geurts, J., & Rutledge, L. (2003). Towards a multimedia formatting vocabulary, WWW '03, Budapest, Hungary, ACM Press.
- Wu, H., de Kort, E., & De Bra, P. (2001). Design Issues for General-Purpose Adaptive Hypermedia Systems. In Proc. of the 12th ACM Conf. on Hypertext and Hypermedia. Århus, Denmark.
- 3rd Special Workshop on Multimedia Semantics, Pisa, Italy, June 22-24, 2005. Retrieved March 4, from <http://www.kettering.edu/~pstanche/wms05homepage/wms05.html>
- Westermann, U., Zillner, S., Schellner, K., & Klas, W. (2005). EMMO: Tradeable Units of Knowledge-Enriched Multimedia Content. In: U. Srinivasan and S. Nepal, *Managing Multimedia Semantics*, Chapter 13, IRM Press.