Semantically-Enhanced Recommendations in Cultural Heritage

Yiwen Wang

Semantically-Enhanced Recommendations in Cultural Heritage

PROEFSCHRIFT

ter verkrijging van de graad van doctor aan de
Technische Universiteit Eindhoven,
op gezag van de Rector Magnificus, prof.dr.ir. C.J. van Duijn,
voor een commissie aangewezen door het College voor Promoties
in het openbaar te verdedigen
op dinsdag 8 februari 2011 om 16.00 uur

door

Yiwen Wang

geboren te Dangtu, China

Dit proefschrift is goedgekeurd door de promotoren:

prof. dr. P.M.E. De Bra
en
prof. dr. A.Th. Schreiber

Copromotor:
dr. L.M. Aroyo

Contents

Preface

Time flies like an arrow. Back in 2006, one month after finishing my master study at the Free University Amsterdam (VU), I started my PhD research within the CHIP project at the Eindhoven University of Technology (TU/e). At that time, CHIP has already being running for around one year, which provided me with a good starting point. From then on, I discovered the topics of personalization, recommender systems, user modeling and semantic web technologies in the domain of cultural heritage. In this period, I encountered pleasure, confusion and inevitable stress. It is now near the end of 2010 and I am finishing my PhD. There are many people involved in the development of this thesis.

First of all, I would like to thank my three supervisors prof. Paul De Bra, prof. Guus Schreiber and dr. Lora Aoryo. I had the luxury that they all actively took part in my supervision. It is from Paul that I learnt doing research is not only a long term commitment, both mentally and physically, but also a self challenge which requires selection and focus. As my daily supervisor, I am grateful to Lora for her endless ideas, sharp remarks and creation of a pleasant working environment. But most importantly, she continuously pushed me forward and at the same time gave me the freedom that allowed me to find my own way. In the background, Guus was pointing out clear directions, providing valuable inspirations and eventually helping me with the details of finishing the thesis.

I also would like to express my gratitude to all the people I collaborated with. Especially my colleague Natalia Stash from the CHIP project with whom I discussed the work, shared the problems, visited the museums and enjoyed the soups. It was Natalia who accompanied me for the full four years and without whom the research would have been very different, and no doubt more lonesome. As another colleague, Lloyd Rutledge gave me enthusiastic support in the first two years and I will always remember the welcome party he organized for me. From the Rijksmuseum Amsterdam, I would like to thank Peter Gorgels and Xenia Henny for coming up with original ideas and inspiring discussions. Many thanks go to the master students Rody Sambeek, Yuri Schuurmans and Ivo Roes from TU/e and

Chapter 1

Introduction

Web 2.0 - the perceived second generation of the World Wide Web is commonly associated with web applications that facilitate information sharing, collaboration and interoperability (O'Reilly 2005). Its focus on "openness" has led to increased interest in open content and in the use of freely available networked applications which may be regarded as open services (Kelly et al. 2008). Visitors are encouraged to actively engage with services and to generate their own content, in contrast to Web sites where they are limited to the passive viewing of information that is provided to them. In this context, institutes, organizations are starting to open up their previously isolated data and services. They aim to provide visitors with maximal access to their resources and services, which will not be limited by constraints such as the device used by the visitor and his/her location.

1.1 General context in Web 2.0

To support the openness in the Web 2.0 enlivenment, the Semantic Web provides a common framework that allows data to be shared and reused from multiple sources. The World Wide Web Consortium (W3C[1]) standardized presentation languages such as Resource Description Framework (RDF[2]) and Web Ontology Language (OWL[3]). These languages are used to describe arbitrary things such as paintings, people or meetings, and record how they relate to the real world in an RDF triple/statement[4], consisting of a subject, a predicate, and an object. It makes the intended meaning of the data, the semantics, explicit in a machine-readable way, which allows for the integration of data. An RDF graph is a set of triples, which express different levels of semantics. By contrast, the semantics in a

---

[1]http://www.w3.org/

[2]http://www.w3.org/RDF/

[3]http://www.w3.org/TR/owl-features/

[4]http://www.w3.org/TR/rdf-concepts/

traditional database or a XML document is usually implicit and needs additional instructions on how to use and integrate them. In recent years, various interesting open data sets have been available on the Web. The most famous example is the W3C Linked Open Data (LOD[5]) project, consisting of over 13.1 billion RDF triples, which are interlinked by around 142 million RDF links as of November 2009. The LOD data sets include DBpedia[6], DBLP bibliography[7], WordNet[8] and FOAF[9]. These data sets are also interlinked. For instance, the DBpedia RDF descriptions of cities includes *owl:sameAs* links to the Geonames data about the city, and FOAF describes persons who *foaf:made* papers in the DBLP bibliography.

At the TED 2009 conference, Tim Berners-Lee described linked data as boxes of data, when connected via open standards, it enables a thousand flowers to bloom[10]. From this, we may ask the question that amounts to: when people access the huge volume of linked data, can we help them to find the flower(s) they like? In other words, the general problem we investigate in this thesis is:

Can we support visitors with personalized access to semantically-enriched collections?

To approach this problem, a lot of work has been done in deploying user modeling and recommendation technologies (Brusilovsky et al. 2007) as a means for personalized information access. As Kobsa distinguished (Kobsa 2001), there are usually three types of data stored in the user model: personal data about user characteristics, usage data about the user's interactive behavior with the system, and environment data that are not related to users themselves. Based on the information collected in the user model, a variety of recommendation algorithms have been proposed (Burke 2002). Amazon[11] and Last.fm[12] are usually thought as good examples of collaborative filtering algorithms (the most popular and widely used algorithms), assessing the similarity between multiple users in order to recommend unseen items to a particular user. By contrast, content-based algorithms (e.g. Pandora[13]) analyze item descriptions to identify items that are of interest to the user. Demographic algorithms (e.g. Pazzani's model (Pazzani 1999)) suggest items based on inferences about user needs and preferences . There are also hybrid systems (e.g. P-Tango (Claypool et al. 1999)) that combine characteristics of multiple recommendation algorithms in order to minimize the disadvantages of each of them and thus to improve the overall performance (Burke 2002). However, most recommender systems, in the last decade, work in a closed or centralized setting,

---

[5]http://esw.w3.org/SweoIG/TaskForces/CommunityProjects/LinkingOpenData
[6]http://en.wikipedia.org/wiki/DBpedia
[7]http://www.informatik.uni-trier.de/ ley/db/
[8]http://wordnet.princeton.edu/online/
[9]http://www.foaf-project.org/
[10]http://www.ted.com/index.php/talks/
[11]http://www.amazon.com/
[12]http://www.last.fm/
[13]http://www.pandora.com/

meaning that access is usually limited by constraints such as different applications, devices, disconnected databases and distributed user data (Ziegler 2004).

Compared with traditional approaches, Semantic Web technologies provide a machine-readable common format to represent heterogenous collections and it also allows users to describe aspects of their social contexts in a standard way. The availability of open structured data adhering to common ontologies enables the integration of data from more diverse sources and it brings new forms of personalized recommendations in a decentralized environment (Peis et al. 2008). For instance, Foafing the music[14] provides music discovery by means of: user profiling (defined in the user's FOAF description), context based information (extracted from music related RSS[15] feeds) and content descriptions (extracted from the audio itself), based on a common ontology that describes the music domain (Celma 2006).

1.2   Project context of CHIP

Within this thesis we proceed from cultural heritage (museums) as an application domain. In recent years, museums are increasingly publishing their digital collections online, experimenting with and implementing interactive and personalized services on their own Web sites (Kelly et al. 2008). All over the world the number of museum Web site visits is growing fast (Chan 2008). The expectation is that more and more people will spend time preparing their visit before actually visiting the museum and look for related information reflecting on what they have seen or missed after visiting the museum. It can also be expected that museum curators want to enhance visitors' experiences in the more personalized, intensive and engaging way promised by an improved Web (Wang et al. 2009a).

In this context, the Dutch Science Foundation (NWO) funded the Cultural Heritage Information Personalization (CHIP[16]) project in early 2005, as part of the Continuous Access to Cultural Heritage (CATCH[17]) program in the Netherlands. CHIP is a collaborative project between the Rijksmuseum Amsterdam[18], the Technische Universiteit Eindhoven[19] and the Telematica Instituut[20]. As mediators between the technical and the art worlds, working inside the museum allowed the whole CHIP team to realize a real application-driven approach by performing frequent interviews with curators and collection managers as well as having close contact with real museum visitors to extract realistic use cases and requirements.

As a PhD student, I joined the CHIP project in July, 2006 when it had already been running for a year. At that stage, the team had been cooperated with the

---

[14]http://foafing-the-music.iua.upf.edu/
[15]http://web.resource.org/rss/1.0/spec
[16]http://www.chip-project.org/
[17]http://www.nwo.nl/catch
[18]http://www.rijksmuseum.nl/
[19]http://w3.tue.nl/
[20]http://www.telin.nl/index.cfm?language=en

MultimediaN E-Culture[21] project and the STITCH[22] project for the semantic en-
richment of the Rijksmuseum digital collections. Based on it, the first version of
the CHIP demonstrator, called the Art Recommender, was developed, which pro-
vides content-based recommendations for artworks and art concepts. User ratings
from the Art Recommender were stored in a traditional database. In order to get a
direct insight, I started my work with the first evaluation of the Art Recommender,
which tests the effectiveness of recommendations for the real Rijksmuseum visi-
tors. Besides, I designed the minimal user model ontology to store user ratings
to replace the original database schema. All this work is more fully reported in
Chapter 2.

1.3   Research questions and approach

The general problem we investigated in this thesis is: *can we support visitors
with personalized access to semantically-enriched collections?* In order to solve
this problem, we formulate four research questions with respect to user modeling
(RQ 1 and 2) and personalized recommendations (RQ 3 and 4) in cultural heritage.

RQ 1. *Can we acquire user information in a non-intrusive way?*

For recommender systems, it is important to collect user information for
providing personalized recommendations. In order to minimize the intrusiveness
in that users must provide information in advance, we build an interactive rating
dialog with representative samples of artworks for a quick instantiation of the user
model. We address typical issues for user modeling, such as the cold-start problem
for first-time users and the sparsity problem and discuss the solutions. We
perform two evaluations to test the effectiveness of personalized recommendations
for users and to compare different ways for building an optimal user model for
efficient recommendations.

RQ 2. *What is a minimal user model to store user information?*

The first research question serves a input to the second question. Besides
the user's ratings, there are many different types of user information such as
the demographic data and information about the users museum tours. To store
all information, we design a minimal user model ontology as a specialization of
FOAF and use the event ontology SEM[23] to model the user's behaviors during
the tour, e.g. the sequence of artworks in the tour, the user's current position
and the time spent. By using standard existing user model ontologies, we aim to

---

[21]http://e-culture.multimedian.nl/
[22]http://www.cs.vu.nl/STITCH/
[23]Simple Event Model (SEM) http://semanticweb.cs.vu.nl/2009/11/sem/

provide a shared understanding of user information.

RQ 3. *Can we use the semantic structure of collections to improve recommendation algorithms?*

To study this question, we take three steps. Firstly, we develop a content-based recommendation algorithm based on the domain ontology. It recommends related artworks and concepts via artwork features. Secondly, we identify different types of semantic relations within one vocabulary and across multiple vocabularies. The various relations are used to recommend more explicitly related items. Thirdly, we adopt an existing method of instance-based ontology matching to build implicit relations between concepts and combine both explicit and implicit relations for recommendations. On top of it, we define four inference steps and try to generalize our approach as a framework for such semantically-enhanced recommender systems. We perform evaluations for each step respectively. We test the effectiveness of recommendations in step 1, and the number of recommended items and precision in step 2. We measure the recommendation accuracy and discuss the added values of providing serendipitious recommendations and explanations for recommended items in step 3.

RQ 4. *How can we present semantically-enhanced recommendations?*

We develop three tools for particular functions: (i) a Web-based Art Recommender; (ii) a Web-based Tour Wizard, and (iii) a Mobile Guide on PDA and iPod that can be used in the physical museum space. To facilitate navigation and browsing, we adopt existing techniques like Spectacle[24] and Simile[25] in the Art Recommender in order to cluster multiple recommendations based on relations. In the Tour Wizard, we present artworks in the museum tours with different views such as the historical time-line and the museum map. In addition, the system automatically derives the relations which are applied to retrieve explicitly or implicitly related concepts and artworks in order to explain the underlying recommendation inference to users. We evaluate the performance of the Art Recommender in terms of the recommendation effectiveness and usability issues. Due to several constraints from the museum side, we augment the evaluation with a qualitative analysis of personalized museum tours provided by the Tour Wizard and the Mobile Guide on PDA. Besides, we test whether the sequence of recommended artworks in the tour follows an efficient route through the museum with the mobile Guide on iPod.

---

[24]http://www.aduna-software.com/products/spectacle/
[25]http://simile.mit.edu/

1.4   Thesis outline

In Chapter 2, we give an overview about the semantic enrichment of the Rijksmuseum collections, the minimal user model ontology which only stores the user's ratings, and the first implementation of the content-based recommendation algorithm in our first tool, the Art Recommender.

In Chapter 3, we describe how to create personalized online museum tours using our second tool, the Museum Tour Wizard. Besides, we explain the conversion of online museum tours to handhelds using the third tool, the Mobile Guide.

In Chapter 4, we update the Mobile Guide tool with a real time routing system. It can adapt museum tours based on the user's location in the physical museum and his/her ratings of artworks and concepts.

In Chapter 5, we identify a number of semantic relations within one vocabulary and across multiple vocabularies. We apply all these relations in recommendations and test the results in terms of usefulness.

In Chapter 6, we define reusable inference steps for such semantically-enhanced recommender systems. As a follow-up work of Chapter 5, we propose a hybrid approach combining explicit and implicit recommendations based on the semantic structure in the collections.

In Chapter 7, we give an example of reusing user interaction data (tags) to enrich the user model for generating recommendations and we investigate problems that arise in mapping user tags to domain ontologies.

In Chapter 8, we provide the conclusion of what we have done. We also discuss what we have not done but which may follow from our work in CHIP, related projects in CATCH, and other cultural heritage projects.

1.5   A topic-based reading guide

The thesis is organized according to papers that resulted from our work in CHIP. These papers cover results on three main topics of our research (Fig. 1.1): *metadata vocabularies, user model* and *recommendation algorithms*. Metadata vocabularies focuses on the semantic enrichment of museum collections, providing a foundation to our work. User model addresses research questions about acquiring user information (RQ 1) and the storage of user information (RQ 2). Based on metadata vocabularies and user model, we study different recommendation algorithms in order to provide personalized recommendations (RQ 3).

Besides, there are also two additional topics: *tools* and *evaluations*. We present the results for end-users in tools (RQ 4) and test our approach in evaluations, which plays an essential role in our user-centered design method. In Table 1.1, we give an overview of each topic in different development stages.
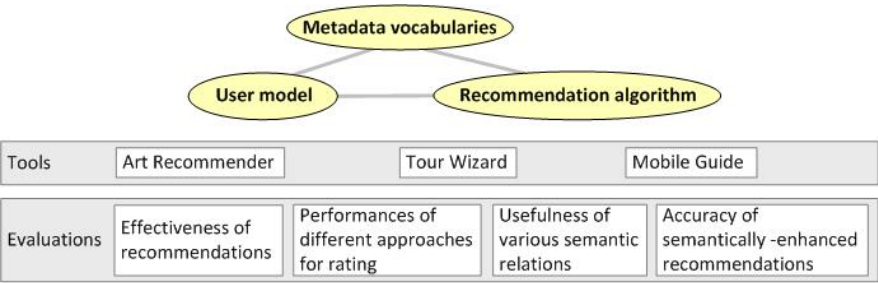
Figure 1.1: A topic-based reading guide

Table 1.1: A topic-based reading guide

Topic	Development in stages	Section
Metadata vocabularies	Mapped to standard vocabularies Getty (ULAN, TGN, AAT) and Iconclass	2.3
User Model	i. Store user ratings, as a specialization of FOAF	2.5
	ii. Store user viewing, tours (artworks and sequence) and mapped to the Simple event model (SEM)	4.4
	iii. Integration of distributed user model	7.4
Recommendation algorithms	i. Apply content-based recommendation (CBR) using the Lapalace method	2.4
	ii. Identify various semantic relations to enhance CBR	5.3
	iii. Use instanced-based ontology matching to build implicit relations and combine explicit and implicit relations for CBR	6.3
Tools	i. Art Recommender	2.4
	ii. Tour Wizard and Mobile Guide	3.4
	iii. Mobile Guide extended with a routing system	4.3
Evaluation	i. Test the effectiveness of recommendations	2.8
	ii. Compare different approaches for rating	2.8
	iii. Compare the usefulness of semantic relations	5.4
	iv. Test the accuracy of semantically-enhanced recommendations	6.4

1.6   Collaborations

The research in this thesis is a collaboration with many people, in particular, with my colleagues from the CHIP project. Lloyd Rutledge, as a post-doc, clarified the various data issues in mappings and contributed to the development of the first prototype of the Art Recommender. As a scientific programmer, Natalia Stash developed the CHIP tools, proposed the first content-based recommendation algorithm for the Art Recommender and helped me with the settings and analysis of evaluations. From the Rijksmuseum, Peter Gorgels and Xenia Henny came up with the original idea of the Art Recommender and they were actively involved in our discussions about the various interfaces of the CHIP tools.

Rody Sambeek, Yuri Schuurmans and Ivo Roes from TU/e joined CHIP for their master graduation projects. Rody and Yuri developed the first prototype of the Mobile Tour Guide (RFID + PDA-based). Based on their work, Ivo developed the second version of the Mobile Tour Guide (iPod-based).

As senior colleagues from VU, Shenghui Wang introduced the instanced-based ontology matching for building implicit relations between concepts; Laura Hollink helped me identify various semantic relations in the domain ontology; Annette ten Teije provided me with inspiration and materials for the work on reusable knowledge elements; and Willem Robert van Hage introduced the original Simple Event Model to enrich the CHIP user model. Also, together with Natalia Stash, he contributed to the extension of the Tour Wizard with a real-time routing system.

My main contribution to the CHIP project, as reported in this thesis covers five main topics: user model, recommendation algorithm, user interface design, reusable knowledge elements and evaluation. For the user model, I designed a minimal user model ontology to store user ratings (Section 2.5), extended it with other existing ontologies (Section 4.4) and explored the interoperability of distributed user models across applications (Section 7.4 and 7.5). For the recommendation algorithm, I identified both explicit and implicit semantic relations in the domain ontology (Section 5.3) and applied them in the recommendation algorithm in order to improve the accuracy and allow for serendipity and explanations (Section 6.3). For the three CHIP tools (Art Recommender, Tour Wizard and Mobile Guide), I contributed to the user interface design in collaboration with Fabrique[26] (Section 2.5, 3.4 and 4.3). For the reusable knowledge elements, I defined the task of semantically-enhanced recommendations and decomposed the task into four inference steps (Section 6.2). Following user-centered design method, I performed a number of evaluations, to test the effectiveness of the original recommendation algorithms (Section 2.8) and the accuracy of the semantically-enhanced recommendation algorithm (Section 6.4), to explore alternatives for quickly building a user model representing his/her interests in the collection (Section 2.8), and to compare the usefulness of different semantic relations in the domain ontology (Section 5.4).

---

[26]http://www.fabrique.nl/

Chapter 2

Generating Ontology-based Art Recommendations

The semantically rich background knowledge about the art domain provides a basis to our research. On top of it, we deploy user modeling and recommendation technologies in order to provide personalized services for museum visitors. Firstly, we develop an interactive rating dialog of artworks and art concepts for a quick instantiation of the user model, which is built as a specialization of FOAF. Secondly, we implement a content-based recommendation (CBR) algorithm, which recommends related artworks and concepts based on the user's ratings. Following a user-centered design cycle, we performed two evaluations with visitors to test the effectiveness of recommendations and to compare different ways for building an optimal user model for efficient recommendations.

As a starting point, this chapter gives an overview about the semantic enrichment of the Rijksmuseum collections, the minimal user model ontology which only stores the user's ratings, and the first implementation of the content-based recommendation algorithm. It serves as input to Chapter 3 and 4.

This chapter was published as a final version as Recommendations Based on Semantically-enriched Museum Collections in the International Journal of Web Semantics (Wang et al. 2008b) and was co-authored by Natalia Stash, Lora Aroyo, Peter Gorgels, Lloyd Rutledge, and Guus Schreiber; and an initial version at Interactive user Modeling for Personalized Access to Museum Collections: The Rijksmuseum Case Study in the proceedings of the User Modeling (UM) Conference (Wang et al. 2007) and was co-authored by Lora Aroyo, Natalia Stash and Lloyd Rutledge.

2.1 Introduction

Museum collections contain large amounts of data and semantically rich, mutually interrelated metadata in heterogeneous distributed databases (Hyvonen et al. 2005). Semantic Web technologies act as instrumental (van Gendt et al. 2006)

in integrating these rich collections of metadata by defining ontologies which accommodate different representation schemata and inconsistent naming conventions over the various vocabularies. Facing the large amount of metadata with complex semantic structures, it is becoming more and more important to support users with a proper selection of information or giving serendipitous reference to related information. For that reason, as observed in (Adomavicius and Tuzhilin 2005; Brusilovsky et al. 2007), recommender systems are becoming increasingly popular for suggesting information to individual users and moreover, for helping users to retrieve items of interest that they ordinarily would not find by using query-based search techniques. From a museum perspective (Bowen and Filippini-Fantoni 2004), personalized recommendations do not only help visitors in coping with the threatening "information overload" by presenting information attuned to their interests and background, but is also considered to increase users' interest and thus stimulate them to visit the physical museum as well.

The Web 2.0 phenomena enables an increasing access to various online collections. The users range from first-time visitors to art-lovers, from students to elderly. Museum visitors have different goals, interests and background knowledge. With the help of Web 2.0 technologies they can actively participate on the Web by adding their comments, preferences and even their own art content. Meanwhile, Web languages, standards, and ontologies make it possible to make heterogeneous museum collections mutually interoperable (Hyvonen et al. 2005) on a large scale. All this transforms the personalization landscape and makes the task of achieving personalized recommender systems even more challenging.

The rest of the chapter is structured as follows. In section 2, we discuss the research challenges, in particular, for recommendations in the open Web context. Then, in section 3 we explain how the museum collection is enriched by using common vocabularies and in section 4 we elaborate on the content-based recommendations for artworks and topics. Further, in section 5, we describe the user model specification and explain the technical architecture (section 6) with an illustrative use case (section 7). Results of two user evaluations are given in section 8. Finally, we discuss our approach and outline directions for future work.

2.2   Research challenges

While the open world brings heterogeneous data collections and distributed user data together, it also poses problems for recommender systems. For example, how to deal with the semantic complexity; how to enable first-time users to immediately profit from recommendations; and how to provide efficient navigation and search in semantically enriched collections. To address the issues, we identify three main research challenges for recommender systems on the Semantic Web:

(i) Enhancing recommendation strategies

In (Hyvonen et al. 2005; Schreiber et al. 2008), we see examples of how ontology engineering and ontology mapping enable content interoperability through rich semantic links between different vocabularies in heterogenous museum collections. This, however, raises new problems for recommender systems applied in such a context, for example, how to deal with the semantic complexity of different types of relationships for recommendation inferencing and how to increase the accuracy and define the relevance of recommendations based on the semantically-enriched collection. Currently, there are many recommendation strategies (Hook et al. 1996; Berkovsky et al. 2007; Brusilovsky et al. 2007) to address these issues: *collaborative filtering* compares users in terms of their item ratings (e.g. Amazon.com[1] and last.fm[2]); *content-based recommendation* selects items based on the correlation between the content of the items (e.g. Pandora[3] and MovieLens[4]). Ruotsalo and Hyvönen proposed an *event-based* (Ruotsalo and Hyvonen 2007) recommendation strategy that utilizes topics from multiple domain ontologies to enhance the relevance precision. In CHIP we have deployed a *content-based* (Wang et al. 2007) strategy, which uses users' ratings on both artworks and art topics in a semantically-enriched museum collection.

(ii) Coping with cold-start and sparsity problems

The heterogeneous population of museum visitors increasingly grows. However, most users are still "first-time" or called "one-time" users to both virtual and physical museums (Bowen and Filippini-Fantoni 2004). Thus, coping with the *cold-start* problem becomes even more crucial for recommender systems applied in the museum domain. In other words, how do we allow first-time users to immediately profit from the recommender system, without requiring much user input beforehand? In addition, in the process of enriching the museum collections, there is an increase in the number and size of semantic structures used. This far exceeds what the user can rate and thus creates the problem of rather *sparse* distribution of user ratings over the collection items. It becomes difficult to recommend effectively when there are not sufficiently many ratings in a large collection. To solve these two closely-related problems, a *hybrid user modeling* approach is widely used (Zakaria et al. 2002; Brusilovsky et al. 2007), combining both user and content centered attributes for generating recommendations. In CHIP, we follow a two-fold approach. On the one hand, we build a non-obtrusive and interactive rating dialog (Denaux et al. 2005) to allow for a quick instantiation of the user model, and, on the other hand, we realize this dialog over the most representative samples for the collection of artworks in order to enable a fast

---

[1]http://www.amazon.com/
[2]http://www.last.fm/
[3]http://www.pandora.com/
[4]http://www.movielens.org/login

population of ratings on artworks and topics (Wang et al. 2007).

(iii) Supporting recommendation presentation and explanation

Due to the heterogeneous character of the data, it is becoming more and more important to facilitate navigation and search in multi-dimensional collections (Albertoni et al. 2004). How to let users explore a large amount of heterogeneous information and still allow for a comprehendable overview? Among the different techniques for visualization clustering (Albertoni et al. 2004), faceted browsers provide a convenient and user-friendly way for hierarchical navigation, as exemplified in MUSEUMFINLAND[5] and E-culture projects[6]. In CHIP, we focus on using and exploring the effectiveness of existing techniques like Spectacle[7] and Simile[8] to cluster multiple recommendations based on properties and present them with different views (e.g. timeline and museum map). Additionally, there is also the problem of explanation, i.e. how to provide users a logic insight in recommendations based on the semantic structure of the collection. Traditional ways to cope with this is using histograms of other users' ratings or likeness to previously rated items (Brusilovsky et al. 2007). In CHIP, explanations are given based on semantic relationships of artworks and topics, which has shown to improve the transparency for recommendations (Cramer et al. 2008).

2.3   Metadata vocabularies

The Rijksmuseum digital collection is stored in two databases: ARIA[9] (educational Website-oriented database) and ADLIB[10] (professional curator database). The current CHIP demonstrator works with the ARIA database, which consists of 729 of the museum's most popular artworks, 486 themes, 690 encyclopedia keywords and 43 catalogue terms. The ARIA database has two main problems: (i) *inconsistent descriptions*: artworks are annotated with different descriptions without using any standard vocabularies; and (ii) *flat structure*: no semantic relationships are described except for general hierarchical relationships between topics (e.g. top, broader and narrower topics) and themes, which brings a severe obstacle for content-based recommendation inference. To address this problem we have focussed on enriching the ARIA database with shared vocabularies. For this, the E-culture project provided the RDF/OWL representation using three Getty vocabularies[11] (ULAN, AAT, TGN) (van Assem et al. 2004) and the CATCH STITCH

---

[5]http://www.seco.tkk.fi/applications/museumfinland/
[6]http://e-culture.multimedian.nl/
[7]http://www.aduna-software.com/products/spectacle/
[8]http://simile.mit.edu/
[9]http://www.rijksmuseum.nl/collectie/ontdekdecollectie
[10]http://www.rijksmuseum.nl/wetenschap/zoeken
[11]http://www.getty.edu/research/conducting research/vocabularies/

project produced mappings to Iconclass thesaurus[12] (van Gendt et al. 2006). We also use SKOS Core[13], created for the purpose of linking thesauri to each other. It specifies the *skos:narrower, skos:broader* and *skos:related* relationships between ARIA topics. Mapping to common vocabularies introduces a semantic structure to the ARIA collection. Table 2.1 gives an overview of all mappings.

Table 2.1: Mappings between ARIA data and other vocabularies

Source data	Vocabulary	Mapped topics	Total topics
Metadata techniques, materials and artists styles	AAT	283	2825
Metadata artists names	ULAN	263	485
Metadata creation sites	TGN	69	507
Metadata subject themes	Iconclass	178	503

The metadata of artworks in CHIP is defined by VRA Core[14] interpreted here to be a specialization of Dublin Core[15] for describing works of art and images of works of art. Fig. 2.1 gives a top-level overview of the RDF Schema used in CHIP, where concepts for places (creation places, birth and death places) in ARIA refer to the geographic location concepts in TGN; artist names in ARIA refer to artist names in ULAN; art styles in AAT are linked to artists in ULAN, and via the link to artists in ARIA the concept of 'style' is introduced in the Rijksmuseum collection; and, finally, subject themes in ARIA refer to concepts in Iconclass. For example, in Fig. 2.1, the artwork "The Jewish Bride" is created by "Rembrandt" (ULAN concept) in "1642" (ARIA concept) in "Amsterdam" (TGN concept). It uses material "Oil paint" (AAT concept) and has a subject "Cloth" (Iconclass concept). Artist "Rembrandt" is born in "Amsterdam" (TGN concept) and has a style of "Baroque" (AAT concept).

To enlarge the scope of the recommendations and to address the scalability aspects of our approach, we plan to include also the ADLIB database (70,000 objects) in the current demonstrator. The enrichment of this collection has already been provided by the E-culture project.

2.4 Recommendations for artworks and topics

In CHIP, a user can start the exploration of the Rijksmuseum collection by first building a user profile, which is driven by an interactive rating dialog (Aroyo et al.

---

[12]http://www.Iconclass.nl/libertas/ic?style=index.xsl
[13]http://www.w3.org/2004/02/skos/
[14]http://www.vraweb.org/resources/datastandards/vracore3/categories.html
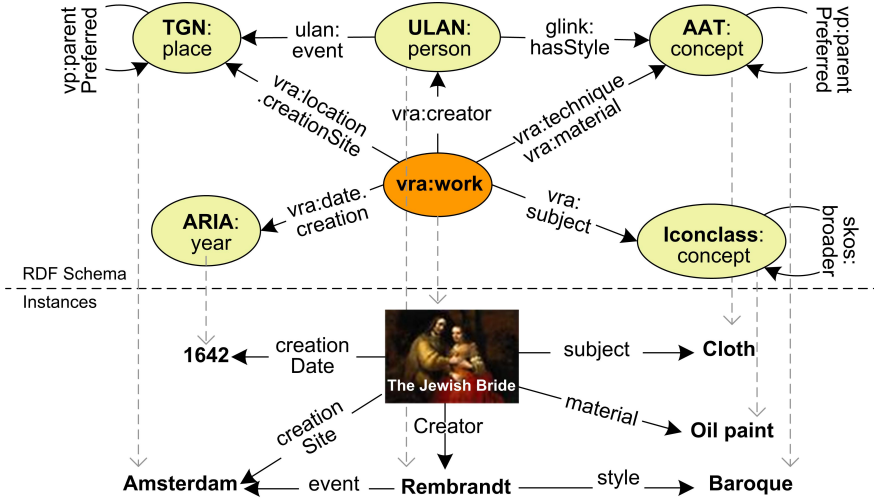[15]http://dublincore.org/

Figure 2.1: Metadata vocabularies in RDF Schema

2007) over the museum collection. In this rating dialog, we distinguish three steps:

Step 1. The user gives ratings to both artworks and associated topics on a 5-degree scale of preference.

Step 2. Based on the semantic relationships, the Art Recommender calculates a *Belief value* to predict the user's interest in other artworks and topics.

In this calculation of belief values for directly linked topics, a smoothing method, (called *Laplace smoothing*), is used: $\theta_j = \frac{N_j + \lambda}{N_{presented} + N_{states} \times \lambda}$

where: θ_j is the probability that the user likes a topic with j stars, N_j is the number of times the topic appears in a set of rated artworks (e.g., artworks the user rated as "I like it"), $N_{presented}$ is the number of times the topic is presented among rated artworks, λ is the smoothing parameter (often set to 1), and N_{states} is the number of rating states (5 in our case).

Using this formula, we then calculate the belief value for topics and artworks:

$$Belief_{topic} = \sum_{j=1}^{5} \theta_j \times W_j \qquad Belief_{artwork} = \frac{\sum_{t=1}^{T} Belief_{topic}}{N_{topics}}$$

where: W_j is the rating of the artwork and N_{topics} is the number of topics.

In other words, the rating of an artwork propagates a belief value to all topics that are directly linked to this artwork and likely to some semantically related topics. The belief value of each topic is used, in turn, to determine the belief value

for artworks.

Step 3. The user may give a rating to either recommended artworks or topics and this is collected as user feedback on the recommendations in the same scale to refine the recommendations presented.

The use of common vocabularies makes it possible to infer additional artworks and topics via properties such as *vra:creator, vra:creationSite* and *vra:materialMedium* (Wang et al. 2008c). Following the content-based recommendation strategy, we allow for the enlargement of the recommendation scope through meaningful links. Also, it is partially helpful for solving the cold-start and sparsity problems. Even with a limited number of ratings, the demonstrator still may produce recommendations through the semantic relationships and order them based on the belief value. For example, if the user rates the artwork "The Nightwatch" with 5 stars, the artwork "The Sampling Officials" and the topics "Rembrandt van Rijn" and "Lastman, Pieter" will be recommended. The underlying inference is that "The Nightwatch" has a creator "Rembrandt van Rijn", who also painted "The Sampling Officials", and he has the *student-of* relationship with "Lastman, Pieter". The rich semantic relationships offer explanations for users to understand why a recommendation is produced. By allowing users to rate recommended artworks and topics, it enables a fast rate-recommend loop for refining the user's preferences and increasing the accuracy of recommendations.

Besides the semantic-driven recommendation based on content, we have explored various approaches to address the cold-start and sparsity problems. By consulting museum domain experts, we present users a subset of artworks containing representative topics to rate first in the rating dialog. In such a way, the user profile collects user ratings with well-balanced distributed topics in a short time and makes it possible to quickly generate recommendations through the entire collection.

As an example of distributed user data integration, we have mapped a small set of iCITY[16] user tags to CHIP art topics. The result of this experiment (Wang et al. 2008a) suggests that the user tags may be used to populate the user model in CHIP and enable instant generation of recommendations. However, as we discussed in (Carmagnola et al. 2008), this approach depends heavily on the correctness of the mappings. Another constraint is that the user tags are mostly seen as a stream of concepts that can be interpreted in various of ways, where the museum vocabularies are static.

2.5   A user model specification

Our goal of building a user model in CHIP is to provide a shared and common understanding of user information and behaviors for enhancing the personalized

---

[16]http://icity.di.unito.it/

access to museum collections. Ideally, the user model needs to store (i) user's personal information; (ii) objects that the user has interacted with; (iii) user's activities over the objects (e.g. the user *rates* an object with a value); and (iv) the corresponding contextual information such as time, place and device. All these data allow us to get information of the user in context.

Currently, we have built a minimal user model as a specialization of FOAF[17]. Main classes and properties from FOAF used in CHIP are *foaf:Person* and *foaf:holdsAccount*.

- Class: *foaf:Person* is used to represent the information about a person who holds an account *chip:User* on a Web site. Account specific information is described by *chip:User*, a subclass of *foaf:OnlineAccount*.

- Property: *foaf:holdsAccount* is used to link a *foaf:Person* to a *chip:User*.

The core class in the user model is the *RatedRelation*. It uses the definition of semantic N-ary relations[18] to represent additional attributes describing a relation. For example, Saskia rates artwork "Nightwatch" with a value of 5. This rate relation contains information in the original three arguments: who has rated (Saskia), what is rated (Nightwatch), and what value the rating gives. Each of the three arguments in the original N-ary relation gives rise to a true binary relationship. In this case, there are three properties: *hasRated, ratedObject* and *ratedValue*, as shown in Fig. 2.2. The additional labels on the links indicate the OWL restrictions on the properties. We define both *ratedObject* and *ratedValue* as functional properties, thus requiring that each instance of *RatedRelation* has exactly one value for Object and one value for Value.

There are in total 5 classes in the range of ratedObject property: *vra:Work, ulan:Person, tgn:Place, aat:Concept* and *ic:Concept*. These objects are well-defined with properties in Fig. 2.1 Metadata vocabularies in CHIP RDF Schema. In the definition of the *User* class (of which the individual Saskia is an instance), we specify a property *hasRated* with the range restriction going to the *RatedRelation* class (of which *RatedRelation_1* is an instance). In addition, we have defined the *Tour* class and two related properties: *hasTour* and *tourWork*. The range of *tourWork* is the class *vra:Work*.

Further extension of this specification would require more indepth treatment of contextual information (e.g. device, time, location) and how this is linked to user activities, such as rating an artwork or creating a tour. In addition, also observational data, e.g. artworks visited, time spent with artworks, could be useful to collect, and may possibly be used to increase recommendation efficiency, effectiveness and relevance. For example, does recording the time spent with an artwork, allow us to infer an actual preference for that artwork, even when it is not included

---

[17]http://www.foaf-project.org/
[18]http://www.w3.org/TR/swbp-n-aryRelations/

*External classes and properties
from FOAF*

foaf:Person

foaf:holdsAccount

foaf:OnlineAccount

*Internal classes and properties
in CHIP*

subClass of

chip:User

chip:hasRated   chip:hasTour

chip:
ratedValue

chip:Value   chip:RatedRelation   chip:Tour

chip:ratedObject   chip:tourWork

*Classes from the
museum vocabularies,
refers to Fig. 1
Metadata vocabularies*

tgn:
Place   ulan:
Person   aat:
Concept   iconclass:
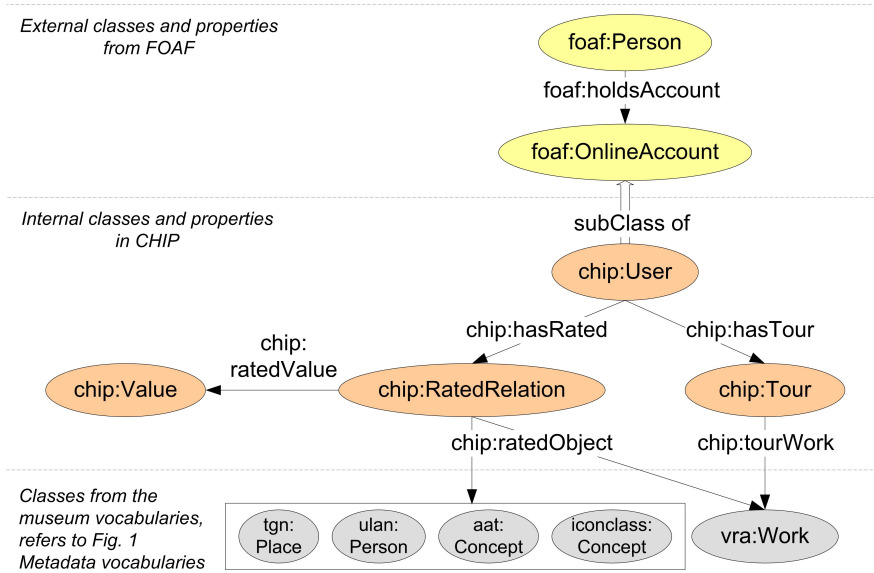Concept   vra:Work

Figure 2.2: Main classes and properties in the CHIP User Model

in the tour or not rated? If we know where a user has been, when visiting a city, does this allow us to infer a consistent interest in particular topics?

2.6 Architecture and implementation

Fig. 2.3 shows the core CHIP components, third-party open APIs, which deliver semantic search results in CHIP (E-Culture API) or additional user data (iCity API) and tools that CHIP uses for data visualization.

The server-side CHIP core components are described below:

- **Collection data** refers to the enriched artwork collection, currently the Rijksmuseum ARIA database, maintained in a Sesame Open RDF memory store and queried with SeRQL.

- **User data** contains user models stored in OWL and tour data stored in XML. To be used by the Mobile Tour Guide, the user models currently have to be transformed to XML.

- **Web-based components** are an Art Recommender and a Museum Tour Wizard realized as Java Servlets and JSP pages with CSS and JavaScript.
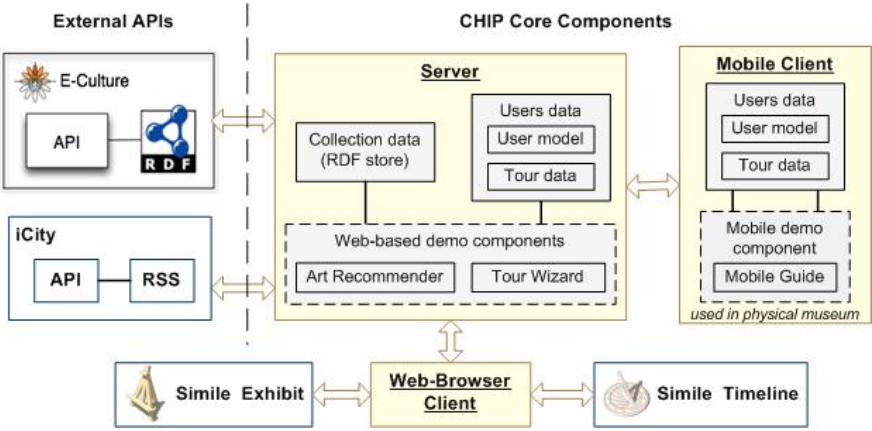
Figure 2.3: CHIP Overall Architecture

Another CHIP client, implemented on a PDA (MS Windows Mobile OS) contains a standalone application Mobile Guide. It is an RFID-reader-enabled device and could also work offline inside the museum and subsequently be synchronized with the server-side on demand. The user profile and the tour data (both in XML) can be downloaded from the CHIP server to the mobile device to be used during the tour in the museum. When the museum tour is finished, the user data can be synchronized with the user profile on the server.

Fig. 2.4 presents the details with respect to the usage of the E-Culture API for semantic search in CHIP. Each user query in CHIP is sent to the E-Culture server, which sends a JSON file back with a list of artworks related to the search query. For every artwork we get a score (relevance of the search result) and a path (search path in the graph). We then further process the JSON file and add more CHIP-specific information to each artwork, like concepts that are associated with this artwork (from the collection data) and the artwork rating (from the users data). The resulting CHIP JSON file is sent to the Simile Exhibit tool to be presented in a faceted view.

In order to experiment with user tag interoperability between the CHIP demonstrator and third party applications, we have adopted an open API to request and link user data from iCity using an RSS feed. Once the user's personal (login) information is authenticated in a dialog between iCity and CHIP, we map the iCity user tags to the CHIP vocabulary set (ARIA shared with Getty and Iconclass) by using the SKOS Core Mapping Vocabulary specification.
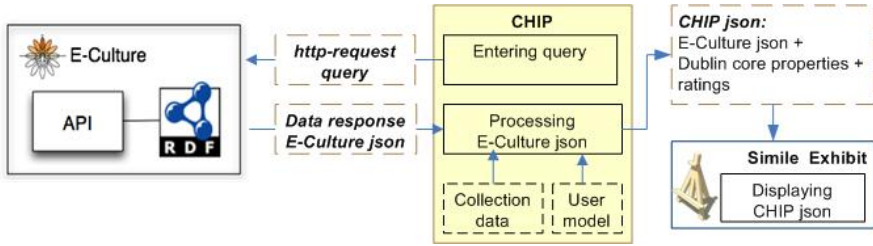
Figure 2.4: Application of E-Culture API in CHIP

2.7 Usage scenario

In this section we describe a typical usage scenario of the CHIP demonstrator in order to illustrate the main user-system interactions.

Saskia is planning her first-time visit to the Rijksmuseum Amsterdam. She does not know a lot about the collection and she would not be able to spend much time there either. Here is how the CHIP demonstrator could help her:

- finding out what she likes in the Rijksmuseum collection

- preparing a personalized museum tour (in terms of time to spend and number of artworks to see)

- storing the data of her visit so that she can later on use it

To login on the CHIP online demonstrator Saskia needs to create a user account. Once logged in, she can choose either the *Art Recommender* tab, to quickly get acquainted with the Rijksmuseum collection and find out her art interests, or she can choose the *Tour Wizard* tab to create different personalized tours and see their layout on the Rijksmuseum map or on a historical timeline. A general *Semantic Search* option supported with an autocompletion function is available, if she wants to search for artworks or topics.

Everywhere in the CHIP demonstrator Saskia can give a rating (in a *5-degree rating scale*) from 1 star (I hate it) to 5 stars (I like it very much) on an artwork or a topic presented on the screen. Each rating of an artwork results in: (i) directly including the artwork with the rating in her user profile, (ii) using the updated user profile to generate a *list of recommended artworks* and a *list of recommended topics*. For each recommended artwork or topic, Saskia can click on the *"why"* (see Fig. 2.5[19]) for an explanation. For recommended topics, "why" explains which artworks with this topic have been rated positively, and for recommended artworks,

---

[19]The screenshots are based on the design by Fabrique (http://www.fabrique.nl/).

it explains which topics from these artworks have been rated positively. As shown in Fig. 2.6, the artworks "Dead peacocks" is recommended because it contains art concepts (or called properties) which are also included in the positively rated artwork "Night watch". Also, Saskia can rate recommended artworks or topics and update her user profile for a further refinement of recommendations.
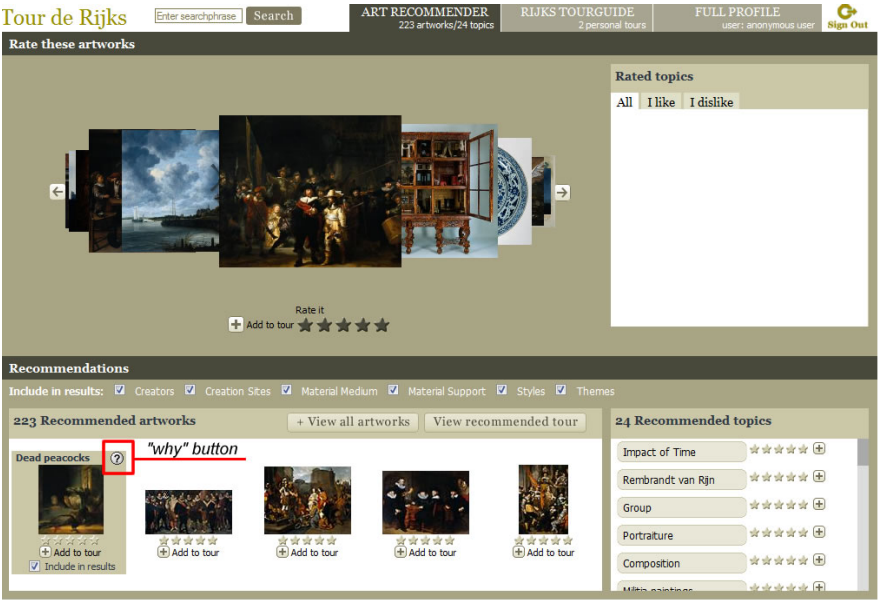

Figure 2.5: "why" button in the Art Recommender

Based on the collected ratings from Saskia, the *Museum Tour Wizard* automatically generates two tours: "Tour of favorites" containing all her positively rated artworks and "Tour of recommended artworks" containing the top 20 recommended artworks. Saskia can explore the tours by viewing the artworks on a museum map (see Fig. 2.7) or on a historical timeline. She can also create new tours by using the search option for finding topics or artworks to add to the tour.

When Saskia is in the museum she can upload her tours on a PDA and use it for guidance. Artworks currently unavailable in the exhibition are filtered out, but are still to be seen on the PDA as background information. For example, Saskia's tour of favorites consists of 15 artworks and is estimated to last for 75 minutes. But she wants to spend at the maximum one hour, so the Mobile Guide reduces her tour to 12 artworks. When she is ready to start, the Mobile Guide recommends
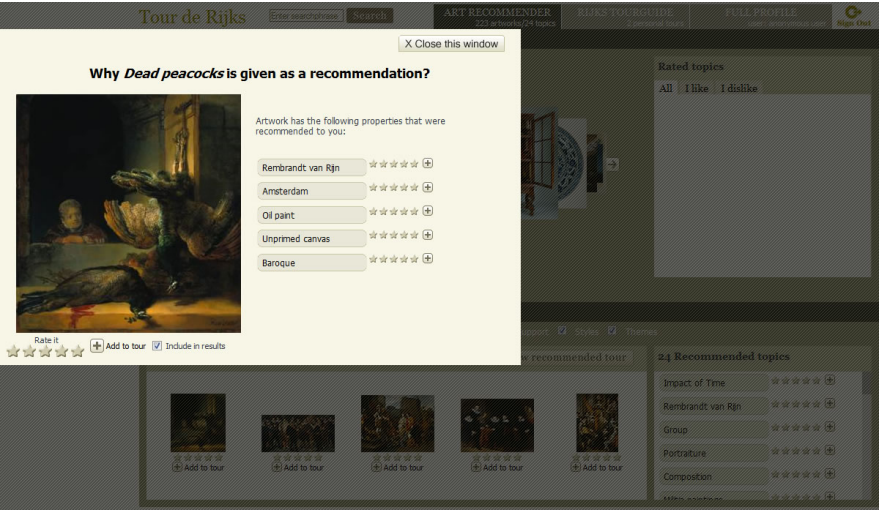
Figure 2.6: Explanation for recommended artworks in the Art Recommender
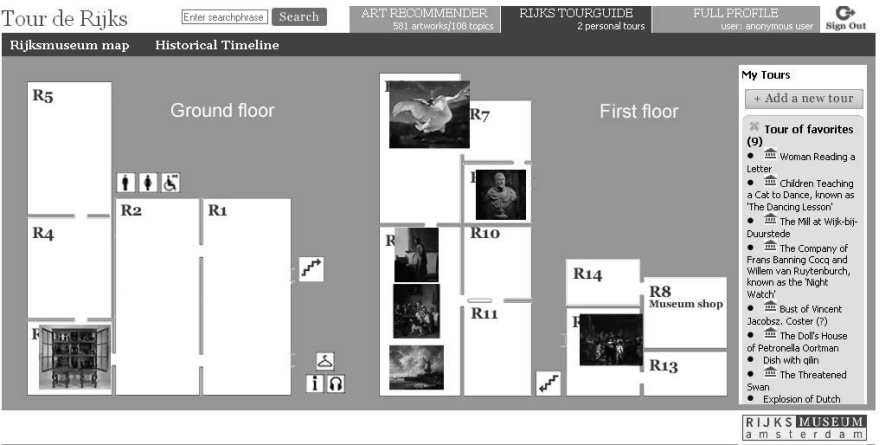

Figure 2.7: Visualizing museum tours on the museum map in the Tour Wizard

her a sequence of artworks and a route to follow.

The usage scenario assumes that all artworks in the museum are tagged with RFID tags. During the tour, Saskia can request information about new artworks by using the RFID tag reader attached to the PDA, which plays an audio file and provides an option to rate this artwork. After listening to the audio and rating the artwork, she follows the initial tour. When the tour is finished, Saskia may synchronize her updated user profile on the PDA with the user profile that was created earlier online. In this way, she has saved all her interactions in the museum and maintained an updated user profile online.

2.8   Evaluation

The overall rationale of the evaluation is to follow a user-centered design cycle in the construction of each part of the CHIP demonstrator. We have performed two initial evaluations at Rijksmuseum Amsterdam with real users to test particular aspects of the demonstrator and derive requirements for further development.

Evaluation I: *effectiveness of recommendations, novices vs. experts*

The goal of the first evaluation (Wang et al. 2007) is to test the effectiveness of the content-based recommendations with the CHIP Art Recommender. 39 users participated in this study. They used the CHIP Artwork Recommender in an average of 20 minutes. The knowledge of the users of the Rijksmuseum collection was tested with questionnaires before and after the test session with the CHIP demonstrator. Our hypothesis was:

The Art Recommender helps novices to elicit or clarify their art preferences from their implicit or unclear knowledge about the museum collection.

To test the hypothesis, we have compared the *precision* of users' topics of interest before and after using the Art Recommender (rating and getting recommendations) (Wang et al. 2007). Looking at the large variety of users, we defined an *expert-value* as a weighted sum of user's personal factors (e.g. prior knowledge of the museum collection, frequency of visiting the museum, interest in art) collected from the questionnaire to distinguish between novice and expert users. As reported in (Wang et al. 2007), the results confirmed our hypothesis, a significant increase of precision was found for novices, while there is a slight increase for experts. However, the distinction between novices and experts is not clear-cut. Plotting the precision on a continuous range of the expert value, we observed, ignoring extreme values, a convergence as expert level increases.

In addition, we have derived four dominant factors about the museum visitors target group. Most of the users appear to be:

- Small group with 2-4 persons and a male took the leading role (67%)

- Middle aged people in the range of 30-60 years old (62%)

- No prior knowledge about the Rijksmuseum collections (62%)

- Strong interest in art (92%)

From this, we get a clear image what are the characteristics of the main target users. The main questions in this context are: (i) what kind of interaction and personalization topics do we need for providing personalized access to the museum collection? (ii) How to structure, store and use the user characteristics to refine the current user model?

Evaluation II: *Representative samples for rating, sparsity and cold-start*

The second evaluation was performed online with 63 participants, most of them are first-time users of the CHIP demonstrator. Based on a functionally-enhanced CHIP Art Recommender, which allows to search for artworks and topics, we explored different alternatives for getting recommendations through the entire collection, to solve the *sparsity* and partially the *cold-start* problem. The evaluation consists of two parts: Part 1 is to let users assess 45 well-distributed topics and Part 2 is to randomly split users into six different groups to rate artworks and topics in a short time (limited to 5 minutes). These six groups follow different alternatives to build their user profiles according to two independent variables: (i) sequence of artworks, which are presented in the Art Recommender for users to rate; and (ii) target of ratings. These two variables ranged over the following values: *Sequence of artworks* (random, expert-sorted, expert-sorted + self-selected); and *Target of ratings* (rate artworks, rate artworks and topics). Here "expert-sorted" means that domain experts selected the first 20 artworks, which overall cover a well-balanced distribution of topics through the entire collection. After that, artworks appear in the order of the number of topics each contains. The "expert-sorted + self-selected" condition allows to search for artworks and topics based on "expert-sorted". Table 2.2 gives an overview of the results according to the six groups using different approaches, where: R(Random), E(Export-sorted), S(Self-selected), Ra(Rate artworks) and Rt(Rate topics).

Table 2.2: Evaluation II: Results in six groups

Group	1	2	3	4	5	6
Sequence of artworks	R	R	E	E	E+S	E+S
Target of ratings	Ra	Ra+Rt	Ra	Ra+Rt	Ra	Ra+Rt
Number of user ratings	96	151	170	224	157	203
Match of preferences	24%	30%	45%	48%	49%	44%

The results show that: first, the "expert-sorted" sequence of artworks works very well for first-time users to quickly build their user profiles with well-distributed

topics through the entire collection; and second, "rating both artwork and topics synchronously" increases the total number of the user's contributions (ratings) and it seems to improve the precision of recommendations; however, at some moment, it might lead to information overload.

All in all, the two evaluations gave us some critical insights in: (i) how to further specify the target group and adapt the user interaction and interfaces for the main groups of users, (ii) how the sequence of artworks affects the recommendation relevance and ranking. Further we learned about the context in which the users are visiting the museum, e.g. in small groups of 2-4 persons, and usability issues of the mobile device.

2.9   Discussion and future work

In this chapter, we demonstrated how Semantic Web technologies are deployed in a realistic use case to provide personalized recommendations in the semantically enriched museum collection. The semantic enrichment provides relational and hierarchical structure which we further exploit in combined artwork and topic based recommendations. The evaluation suggests that this approach helps especially novices to elicit their art preferences about the collection.

However, it also brings a problem with respect to calculating the recommendation relevance. For example, if the user rates an artwork, we currently treat all its properties, such as "creator", "creationSite" and "material" with equal strength in the recommendation strategy, where they could carry different importance for each user. In other words, the "creator" could be more interesting to the user than the "material". Moreover, material is likely to be a less discriminative factor for recommendations, as most of the artworks in this collection are of the same material. Thus, each artwork property should be assigned with a different weight in the recommendation strategy. Even more, the relevance of each property for a given user should be dynamically adjusted according to the user's ratings, or used with a default value when not enough user ratings are available. If a user mostly rates values of the property of "creationSite", these should have a priority in recommendations. As follow-up work, we look for solutions to solve this problem in later chapters. We perform an evaluation in Chapter 5 to compare different properties (or artwork features) as well as various semantic relations in order to find which relations are useful for users. Based on these findings, we define weights for specific relations and propose a hybrid recommendation algorithm in Chapter 6. We compare this new algorithm with the old one which was introduced in this chapter in terms of recommendation accuracy.

Web 2.0 enjoys increasing popularity and offers a rich network with a large number of user communities and a staggering amount of user generated content. For recommender systems this suggests, as a main opportunity, the integration of distributed user data for recommendations. Such integration would amount to a

unified user model that can be used across multiple applications, enriching the potential for recommendations by using the distributed user data. However, to realize such a user model, issues of storage, linking, representation and inference must be solved. As a first step of defining such a user model specification, we proposed to extend the existing FOAF specification with possibilities to express user activities and interests in objects. As a second step, we mapped the CHIP user model to an existing event model ontology SEM in Chapter 4. The goal of the mapping is to store users' information during the museum tour, e.g. visited artworks, users' position in the museum, and locations of artworks in the exhibition.

As observed in (Greaves and Mika 2008), Web 2.0 is a user centered community, whereas the Semantic Web must be regarded as primarily a network connecting professional data through semantic relations. When we extrapolate this observation to our approach in CHIP, the major challenge is not to linking data from social networks and other Web 2.0 applications, but to bridge the gap between the semantic structure of museum collection data, which is professional semantics, and the variety of meanings found in open social networks, which rely on what is commonly called emergent semantics. The direction of bridging this semantic gap, as suggested by (Gruber 2008), is to add structure to user data, as a function of how this data links to repositories of information. One way of creating such a structure, as proposed for SIOC in (Bojars et al. 2008), is to characterize social networks not as relations between people, but rather as object centered sociality. Objects could simultaneously be characterized by semantically linked meta data, obtained from professionals. Admittedly, this is still a long way from collective intelligence (Gruber 2008), but it is likely a significant step towards providing better recommendations, that take the users social context into account. In Chapter 7, we provide an example of collecting distributed user models for interoperability. In this example, we extract user tags about cultural events gathered by another application iCITY and map these tags to the museum domain ontology. These mappings are used to enrich the user model for generating recommendations in the CHIP Art Recommender.

Chapter 3

Creating Personalized Museum Tours

We introduced the Art Recommender in Chapter 2. In this chapter, we present two other tools: Tour Wizard and Mobile Guide. Based on the user's ratings, the Web-based Tour Wizard recommends museum tours consisting of recommended artworks that are currently available for museum exhibitions. The Mobile Guide converts the recommended tours to the mobile devices PDA that can be used in the physical museum space. Due to several constraints, we augment the evaluation with a qualitative analysis of personalized museum tours provided by the Tour Wizard and the Mobile Guide.

This chapter was published as: Cultivating Personalized Museum Tours Online and On-Site in the International Journal of Interdisciplinary Science Reviews 2009 (Wang et al. 2009a) and was co-authored by Lora Aroyo, Natalia Stash, Rody Sambeek, Yuri Schuurmans, Guus Schreiber and Peter Gorgels.

3.1 Introduction

In recent years, the purpose of museums has shifted from merely providing static information of collections to providing personalized services to various visitors worldwide, in a way suiting visitors' personal characteristics, goals, tasks and behaviors. Personalization enables changing "the museum monologue" into "a user-centered information dialog" between the museum and its visitors (Bowen and Filippini-Fantoni 2004). This interactive dialog occurs not only in the real museum, but also in the "virtual museum" (Schweibenz 1998) on the museum Web site. Museums are increasingly experimenting with and implementing more personalized and interactive services on their own Web sites. All over the world the number of museum Web site visits is growing fast (Chan 2008). Visitors spend more and more time on the museum Web sites to do things, e.g. to discover interesting artworks, prepare a museum tour, or learn related knowledge about artworks, usually in relation to a (possible) physical museum visit. This brings a great challenge for museums to provide a personalized and extended museum experience for visitors in an immersive museum environment, which includes both the virtual museum

(online) and the real museum (on-site).

In this context, the CHIP (Cultural Heritage Information Presentation) project has been working at the Rijksmuseum Amsterdam[1] since early 2005, as part of the NWO-CATCH[2] (Continuous Access to Cultural Heritage) program. CHIP is a cross-disciplinary research project, combining aspects from cultural heritage (museum) and computer science. From the museum perspective, it poses three issues: (i) how to acquire visitors' interests in the museum collection; (ii) what kinds of personalized services can be provided on the museum Web site and in the real museum space; and (iii) how to link visitors' museum experiences online and on-site and what approaches can be deployed to increase visitors' motivation to return to the immersive museum environment (online and on-site). From the computer science perspective, our main research challenges are: (i) to enrich the museum digital collection with semantic structures; (ii) to recommend artworks and related concepts in a way suiting different users' art interests; (iii) to build an interactive and dynamic user model that stores users' various information; and (iv) to create personalized online museum tours and to convert these online tours to on-site tours on the mobile device.

To address these issues from both disciplines, we have so far taken the following steps: i) used technologies associated with what has been called "the Semantic Web"[3] to enrich the museum digital collections by mapping them to existing common vocabularies; (ii) created an interactive user model as an extended domain-overlay to acquire and store users' art interests and other information; (iii) developed three different tools within the CHIP demonstrator, namely, *the Art Recommender, the Tour Wizard* and *the Mobile Guide*. The Art Recommender applies content-based recommendation techniques to recommend artworks and concepts based on the user model. The Tour Wizard generates personalized online museum tours containing recommended artworks and allows users to create new tours by adding/removing artworks. The Mobile Guide converts online tours to on-site tours on the mobile device and guides users' visits in the real museum environment. Following a user-centered design method, we have performed a series of empirical user studies (Wang et al. 2008b) with real users to derive the requirements for building these tools and to access the quality of personalization provided by the tools.

In this chapter, we focus on describing the creation and conversion of online and on-site museum tours implemented in the Tour Wizard and the Mobile Guide tools. The descriptions of the semantic enrichment of museum digital collections, the user model and the Art Recommender tool are explained in (Wang et al. 2008c). The rest of chapter is structured as follows: In Section 2, we discuss related work about existing museum tours and in Section 3, we give a use case of such tours. Then, in

---

[1]http://www.rijksmuseum.nl (05/03/09)
[2]http://www.nwo.nl/catch (05/03/09)
[3]http://www.sciam.com/article.cfm?id=the-semantic-web&print=true (05/03/09)

Table 3.1: Exploring existing museum tours

Museum and Tour type	Tour description
Rijksmuseum Amsterdam (Human-guided tour)	The visitor follows a human guide, which selects artworks and gives corresponding information to visitors using speech, gestures or extra material.
Rijksmuseum Amsterdam (Audio tour)	Most artworks are labeled with a number, which are coupled to an audio track on the visitor's audio device.
Tate Britain (Online tour)	The visitor sees a virtual museum representation on a museum map. Rooms can be selected and each room contains a set of artworks.
Metropolitan Museum of Art (Online tour)	The visitor can select six different virtual reality rooms and then navigate the virtual rooms and the artworks inside the rooms.
Van Gogh Museum Amsterdam (Multimedia tour)	The visitor walks through the museum following a timeline of Van Gogh's life. Artwork information can be seen on a PDA from the artwork list.
Netherlands Architecture Institute Rotterdam (Multimedia tour)	Artworks have sensors that can be scanned using a PDA. If a sensor is scanned, the corresponding information will be presented on the PDA.

Section 4, we describe how to create online museum tours using the Tour Wizard and how to export the tours using the Mobile Guide and give users guidance during their tours in the physical museum. Further, a qualitative analysis of these tools is given in Section 5. Finally, in Section 6, we discuss our approach and outline directions for future work.

3.2 Related work

Museum tours offer visitors a unique experience in the museum and special insights about the museum collection. There are mainly four types of museum tours: *human-guided tours, audio tours, online/virtual tours*, and *multimedia tours* (Wang et al. 2008c).

The traditional human-guided and audio tours are usually available in most mu-

seums. In recent years, enhanced Web technologies have enabled increasing access to museum Web sites. As a trend, more and more museums create online/virtual tours on the museum Web sites for online visitors across the world. Besides the online tours, with the support of mobile computer technology, multimedia tours are becoming increasingly important to visitors by enhancing their museum experience (Anderson and Blackwood 2004). Many museums offer multimedia tours, which are implemented on different mobile devices. These tours strengthen the exhibits by allowing visitors more informed enjoyment and knowledge, hence greater engagement with the artworks (Sakamura 2003). However, the study conducted at the Van Abbemuseum with electronic handheld guides indicated that the different handheld devices did not perform better than the traditional paper guide, although the handheld devices could not be used to their full potential as audio and video data were not present (Bartneck et al. 2006).

Our mandate in CHIP to enhance personalized museum experiences both on the Web site and in the real museum space dictates a focus on the online and multimedia museum tours. From the exploration stage, we found that most online and multimedia tours suffer from two main problems. The first is lack of content personalization and dynamic adaptation according to the visitors' interests and the contextual information. Most tours contain a fixed list of artworks, which is the same for everyone or for visitors from the same pre-defined user groups (e.g. groups of tourists, students, experts). The second problem is lack of connection between online tours and on-site/multimedia tours, which are usually separated tours without any connections. These two problems became our main challenges in building the personalized online and on-site museum tours.

3.2.1 Providing personalized content

For most online and multimedia museum tours, in order to deliver personalized content, the visitor's interests and contextual information are usually required. The user information can be inferred implicitly by observing users' behavior in the museum or during their interactions with the multimedia device; it can also be provided explicitly by the users (Bowen and Filippini-Fantoni 2004). The data are stored in the user model and are exploited in the process of content generation to describe or recommend objects potentially relevant for users.

These types of solutions are quite complex and therefore have been developed mostly in the context of academic research. For example, the wearable computer (Fig. 2.a), developed at MIT Media Lab, delivers audio and visual narration adapting to the user's interest from her physical path in the museum and length of stops (Sparacino 2002). The PEACH project (Rocchi et al. 2004) develops a PDA-based museum tour application (Fig. 2.b), whose content is adapted to the visitor, location-aware and only available in certain locations in the museum. The INTRIGUE project (Fig. 2.c), which relies on user-modelling, recommends sight-

seeing destinations by taking into account the preferences of heterogeneous tourist groups (Ardissono et al. 2003). Another application is the iPod Multimedia Tour (Fig. 2.d) designed for the Saint Louis Art Museum[4] by Schwartz and Associates

Figure 3.1: Other multimedia museum tours

Creative[5], which won the 2007 Muse Award[6]. (The St Louis Museum is one of the first in the world to offer a tour on the iPod.)

For content personalization in CHIP, we built a user model to collect the user's interests automatically from his or her interactions. Based on the user model, we adopted a content-based recommendation strategy to recommend both artworks and art concepts, that might of interest. In this way, our system enables the delivery of personalized content.

3.2.2   Supporting the virtuous circle of museum visits

The term "virtuous circle" was coined by Ailsa Barry from London Natural History Museum (Barry 2006). It means creating a connection between the online (virtual) and the on-site (real) information through functions such as bookmarks allow people to save information of interest from the museum interactions (e.g. from Kiosks, PDAs) and access it after the visit via e-mail or on a personalized page available on the museum Web site (see Fig. 3.2.a). The essence of the virtuous circle is that, the visitor can start the museum tour either from the Web or in the museum, and can extend the tour from the Web to the museum and back to the Web, or vice versa.

There are two main reasons to link the visitor's experiences online and on-site into a virtuous circle. First, such linking supports a continuous learning experience. By activating previous knowledge, it helps retain memories over time, enables the person to pursue individual interests, and allows him or her to focus more on experimentation, discovery and the aesthetic experience during the visit. Second,

---

[4]http://www.slam.org/ (05/03/09)
[5]http://www.sacreative.com/ (05/03/09)
[6]http://www.mediaandtechnology.org/muse/2007muselist.html (05/03/09)
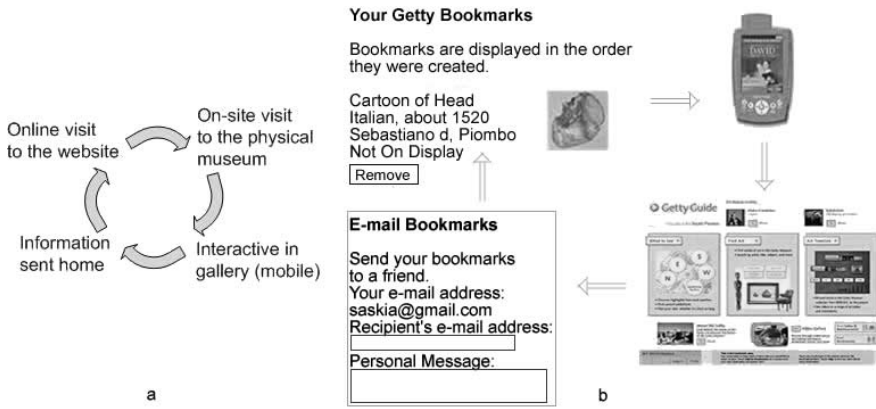
Figure 3.2: (a) Virtuous circle of the museum visit (b) and the GettyGuide example

it can strengthen the visitor-museum relationship by driving traffic to the museum Web site and stimulating further interest in the digital collection.

In practice, there have been several museum projects (e.g. Tate Modern, Science Museum Boston[7] and the GettyGuide[8]) to encourage the exploration of the virtuous circle. As illustrated in Fig. 3.2.b, GettyGuide has kiosks that allow users to bookmark objects within their collections, and these are then e-mailed back. However, commonly these e-mailed bookmarks contain distinct information and are not directly linked back to the museum Web site (Barry 2006). They therefore do not really encourage visitors to expand or continue their experiences further within the virtual space.

To maintain the virtuous circle, we implemented the distributed user model, which stores a user's various information during the online and on-site museum tour. Once the tour is finished, the user model is synchronized on these two different clients (the Web site and the mobile device) for the user's next time visit. In such a way, we aim to extend the personalized museum experience in a more long-lasting and engaging way.

3.3   Usage scenario

To explain a possible scenario, let's suppose a new visitor Saskia plans to visit Rijksmuseum Amsterdam for the first time. She does not know much about the museum collection and she has only limited time for the visit. Here is an illustrative

---

[7] http://www.mos.org (05/03/09)
[8] http://www.triplecode.com/projects/getty.html (05/03/09)

scenario before, during and after the museum visit.

Before the museum visit.   Considering her limited visiting time, Saskia wants to make her visit efficient so that she may be sure to see some artworks which are really interesting to her. Thus, she decides to prepare herself a bit before she goes to the museum. She checks the museum Web site and looks for some artworks that she would like to see. However, because of the online collection of artworks on the museum Web site, Saskia is confused by too much information, so she needs recommendations of artworks that (i) match her art interests; (ii) are currently available in the museum exhibition; (iii) fit her time-constraint.

During the museum visit.   After the preparation, Saskia visits the Rijksmuseum Amsterdam. At the reception, she rents a mobile museum guide, with audio and with a detailed text description of artworks. In addition, Saskia wants to load the data she prepared beforehand. She expects the resulting combination of her data and the museum's to be presented on the mobile device, indicating the actual locations of artworks from the tour and the route linking these artworks. During the visit, Saskia sees some new artworks and wants to receive more information about them.

After the museum visit.   Afterwards, Saskia becomes more interested and excited about the museum collection. She wants: (i) to know more about what she has seen in the museum; (ii) to learn new aspects about artworks, which are related to her art interests; and (iii) to keep up-to-date with new artworks coming in the museum which might be interesting for her.

3.4   Personalized museum tours

The goal of museum tours within the CHIP demonstrator is to enhance the visitor's museum experience in a more intensive, long-lasting and engaging way, by linking the museum experiences both online and on-site. Following a user-centered design method, we have so far developed three tools within the CHIP demonstrator in a coherent way, namely, *Art Recommender, Tour Wizard* and *Mobile Guide.*

- The *Art Recommender* helps users to discover their art interests in the museum collection and to store them in a corresponding user model.

- The *Tour Wizard* generates online museum tours containing interesting artworks recommended from the first tool, *Art Recommender*). The online tours can be presented both on a geographical museum map and in a historical timeline.

- The *Mobile Guide* converts online museum tours (generated from the *Tour Wizard*) to the on-site tours on the mobile device, and assists the user to find his or her way during the visit. When the tour is finished, it sends the user's real behaviors to update the user model on the Web server.

To further understand the relations among these three tools and how they work together, we give an architectural diagram of core components in Figure 3.3.
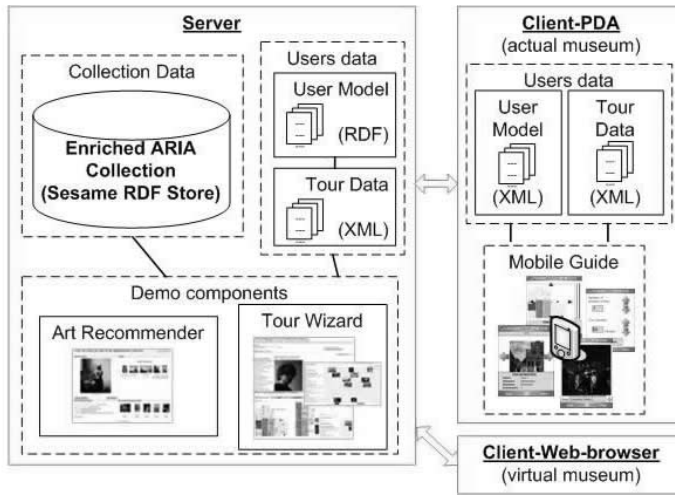
Figure 3.3: CHIP Architecture: Core components

The CHIP demonstrator is based on a client-server architecture. There are three core components on the server-side (Aroyo et al. 2007): (i) *Collection data* refers to the enriched museum collections, currently the Rijksmuseum ARIA database, maintained in a Sesame Open RDF memory store and queried with SeRQL. (ii) *User data* contains user models stored in RDF and tour data stored in XML. To be used by the Mobile Guide Guide, the user models currently have to be transformed to XML. (iii) *Web-based demo components* are the Art Recommender and the Tour Wizard realized as Java Servlets and JSP pages with CSS and JavaScript.

Another CHIP client, implemented on a PDA (MS Windows Mobile OS) contains a standalone application Mobile Guide. It is an RFID (Radio Frequency Identification) reader enabled device and can work offline inside the museum and subsequently be synchronized with the server-side tools on demand. The user model and the tour data (both in XML) can be downloaded from the CHIP server

to the mobile device to be used during the tour in the museum. When the museum tour is finished, the user data can be synchronized with the user model on the server. The second version of the Mobile Guide is now being prepared and will be implemented on an iPod (Roes et al. 2009).

In this chapter, we focus on describing the creation and conversion of online and on-site museum tours using the Tour Wizard and Mobile Guide. For detailed descriptions about the Art Recommender, the semantic enrichment of the collection (metadata vocabularies) and the specification of the user model, see (Wang et al. 2008b).

3.4.1 Web-based Tour Wizard

Based on the ratings stored in the user model, the Tour Wizard automatically generates personalized museum tours of artworks. It contains recommended museum tours and user-created tours. The recommended tours contain artworks, that might be of interest to the user according to his or her ratings of presented artworks and concepts. The user could also create tours by adding or removing artworks. The tours can be presented both on the Rijksmuseum map (Fig. 2.7) and on a historical timeline (Fig. 3.4).
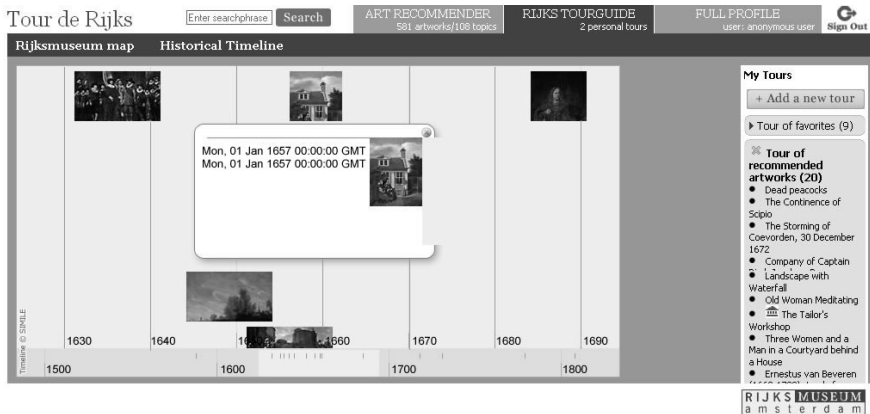
Figure 3.4: Tour Wizard: museum tours in the timeline bar

Tour Wizard allows users to semantically search for artworks or related concepts to add them to the tours. This function is supported by the search API of the MultimediaN E-Culture project (Schreiber, 2006). For example, a user Saskia wants to make a tour about artworks created by the Dutch painter Rembrandt van Rijn. If she searches "Rembrandt", the system will return 4 types of results

(see Fig. 3.8.a): (i) Creator "Rembrandt van Rijn"; (ii) Artworks which contain "Rembrandt" in the title, e.g. "The Prophetess Anna (known as 'Rembrandt's Mother')", "Self portrait of Rembrandt van Rijn" and "Study for a statue of Rembrandt"; (iii) Theme "Rembrandt's cycle"; and (iv) Other creators/painters who are related to "Rembrandt van Rijn", e.g. his teacher "Peter Lastman", his student "Dou Gerrit".

To return to Saskia: she wants to see all of Rembrandt's works and add them to a tour, so she can click on the first search result, which is the creator "Rembrandt van Rijn". The system then will present the description about Rembrandt van Rijn to her and give an overview of all 22 artworks (see Fig. 3.8.b). By viewing these artworks, Saskia could add all of them to her Rembrandt tour or select some of them to add to the tour.

3.4.2   PDA-based Mobile Guide

To export the online museum tours on the mobile device (PDA) and give guidance to users during their visit in the real museum space, we implemented a stand-alone PDA-based Mobile Guide on the HP Ipaq device with RFID reader for user-positioning.

Figure 3.5 illustrates the main functions of the Mobile Guide: (a) select and download online tours; (b) set up the constraints of the tour (e.g. time spent and number of artworks to see); (c) request and receive detailed information (text, image and audio) about an artwork in the tour; (d) receive detailed description about the artwork and rate this artwork in a 5-star scale; (e) indicate the user's current positioning and show the tour route; and (f) retrieve information about the room such as the number of tour artworks that are available in the room.

To download the online tours on the device, the Mobile Guide needs to invoke a Mobile data application on the server (see Fig. 3.6), which is created for exporting and importing information in XML. Then, a Servlet called GetTours will be invoked to fetch the tour data from the data store using SeRQL and returns the information to the PDA as an XML file using a DOM approach as a separate component called the XML Writer. The generated XML file retrieves all data from the tours and returns to the PDA.

Different from online tours, on-site tours in the real museum space encounter a number of constraints, e.g. the availability of artworks, time duration and the route . In the Mobile Guide, we proposed a mapping mechanism: (i) to filter out unavailable artworks from the total set of artworks in the selected tour; (ii) to allow users to limit the number of tour artworks to see and set up the total time duration; and (iii) to link all available artworks and indicate the route for the visit.

In the Rijksmuseum scenario, each artwork is tagged with a passive RFID tag, which is connected with the PDA. We track the user's position by scanning the location of the corresponding artwork. Once the visit is finished, the Mobile Guide
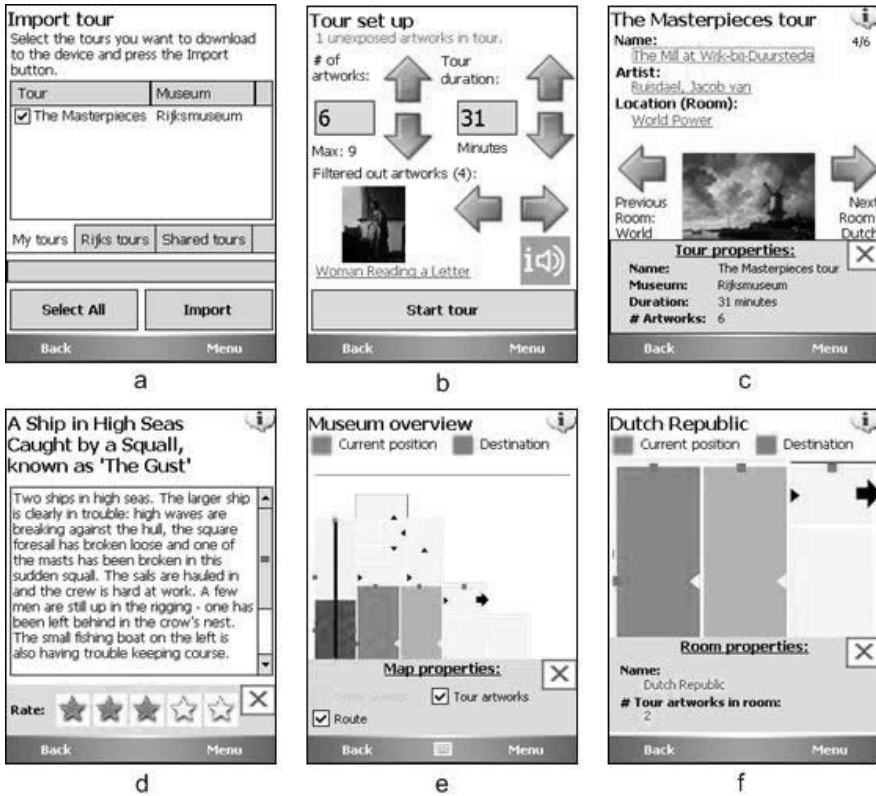
Figure 3.5: Screenshots of the Mobile Guide on the PDA

interactions are synchronized with the user model maintained on the CHIP Web site. As indicated in Figure 3.6, the synchronization is performed by the PostUM Servlet, which receives the user model from the Mobile Guide as a Post variable.

3.5 Qualitative analysis

Following a user-centered design method, we have performed a series of user studies to test the effectiveness of personalized recommendations generated by the Art Recommender (Wang et al. 2007); to explore various alternatives to build a user model representing the user's interests in a short time (Wang et al. 2008b); and to derive requirements for building museum tours (see Section 2).
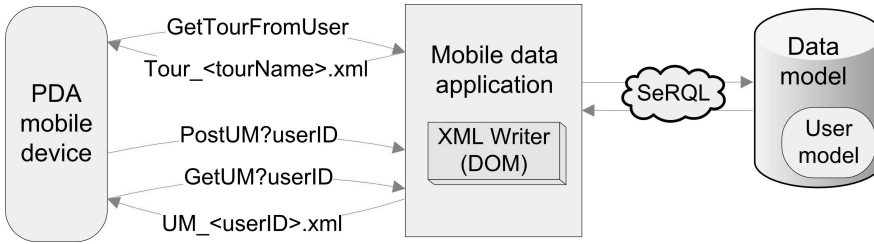
Figure 3.6: Download the online tour and update the user model

However, it is difficult to perform an empirical evaluation on an application mainly used for scientific research in order to access the quality of personalized online and on-site museum tours provided by the Tour Wizard and the Mobile Guide. The problem is the constraints from the museum side, such as permission to use the real museum environment, the attachment of RFID tags to artworks in the current exhibition, and the availability of mobile devices and related hardware. So we have to augment the user studies with a qualitative analysis of personalized
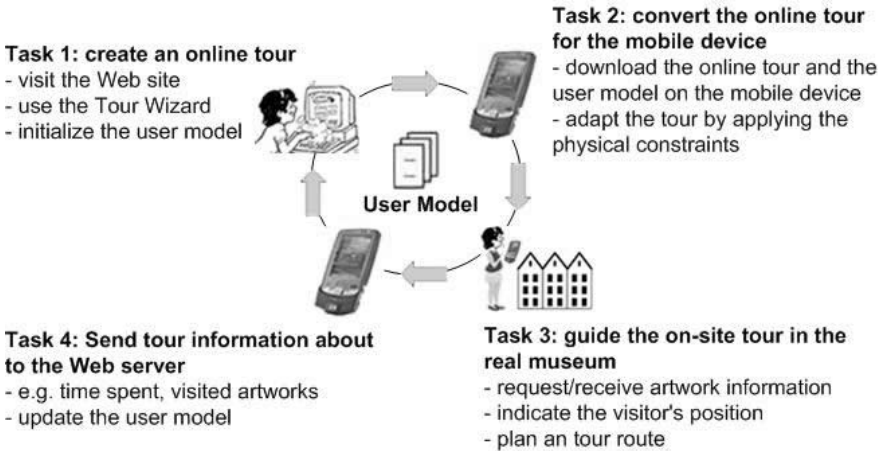
Figure 3.7: Four tasks supporting the virtuous circle of museum visits

museum tours provided by the Tour Wizard and the Mobile Guide, to identify possible issues in usability and topics for future research. To support the "virtuous circle" of museum visits (Fig. 3.2.a), we define four tasks in a pre-defined sequence and discuss related issues/problems in each task. As depicted in Fig. 3.7, the

distributed user model plays a central role, as it is automatically initialized and updated based on ongoing user interactions on the Web server or on the mobile device; and enables the personalization of content.

Task 1: Create an online museum tour on the Web site. The user can visit the museum Web site at home and use the online Art Recommender to rate presented artworks and art concepts. While he or she rates artworks and concepts, the user model is automatically updated to store the declared art interests. Based on the dynamic user model, the Art Recommender will recommend artworks and concepts that fit these interests, and the Tour Wizard will generate online museum tours containing recommended works and allow for adding/removing artworks. From the previous user study for the Art Recommender (Wang et al. 2007), we found that the system efficiently helps users, especially novice users to elicit their art interests in the museum collections and recommend artworks in a way suiting their interests. However, as a sequence of recommended artworks, does the recommended museum tour fit the user's interests? Is the selection of artworks representative for the whole museum collection?

Task 2: Convert the online tour for the mobile device. Once the user gets the mobile device (PDA), the Mobile Guide will download user-selected online tours and the user model. For the adaptation from the online tour to the on-site tour, the Mobile Guide needs to: (i) filter out unavailable artworks for the current exhibition; (ii) sort available artworks based on the location; and, (iii) apply physical constraints (number of artworks and time spent) to adjust the tour. As a preliminary estimate, we presume that each artwork takes 5 minutes. However, it might be quite different for different individuals. Another issue is the user interface on the mobile device, e.g. how to present the artworks with different types of information (image, text and audio) on the relatively small screen of the mobile device?

Task 3: Guide the on-site tour in the real museum. During the tour, the user can request and receive information about new artworks by reading the passive RFID tag attached to the artworks, which also indicates the visitor's current location. With the support of various wireless communication and localization technologies (e.g. RFID, GPS, infrared, bluetooth), it is possible to provide functions that allow social activities for users. Based on the contextual information during the tour (e.g. time, the sequence of artworks, user's activities), how to dynamically adapt the tour? For example, the new artworks the user adds are located in the rooms which have already been visited, and the user does not have much time left, in this case, how to dynamically adapt the tour according to the changes, e.g. plan the new route, arrange the rest of artworks.

Task 4: Send the tour information to the Web. When the user finishes the museum tour, the Mobile Guide will send tour information to update the user model on the Web server. Currently, we only store ratings of visited artworks and

related art concepts. There are still some issues remained, e.g. what are the other contextual information items we need to store from the tour, and how to use these for content personalization in a next visit.

From the analysis, we see that the user model plays an essential role. It stores the user's interactions on two clients (Web and mobile device) and enables the personalization of content. In order to enhance the personalized museum tours, we need to take into account also different aspects of the user model, like the user groups, the context, device, etc. How to store the user data in a standard way that can be shared with and understood by other applications is an important topic for further research that we have partially addressed in (Wang et al. 2008b).

3.6   Discussion and future work

In this chapter we have proposed an approach to exploit personalized museum tours suiting different users art interests and to link the online and on-site tours in an intensive and long-lasting way based on an interactive and dynamic user model. We proposed a method to import online tours from the Web server (Tour Wizard) to the mobile device and to synchronize user data on the mobile device with the Web server. While moving from the online to the on-site tours, physical aspects of the museum are considered, e.g. time spent and number of artworks in the tour. We presented a mapping mechanism for this conversion. Furthermore, we tried to capture innovative new functionality for mobile museum tours like user guidance and user positioning. User guidance and user positioning are used to offer museum visitors a dynamic tour experience. However, the current Mobile Guide implements a basic use case. In the future, we plan to extend the Mobile Guide with the following features or possibilities.

Dynamic adaptation.   When wireless communication is provided, it brings an opportunity for providing dynamic adaptation during the Mobile Guide. For example, the user can receive new recommendations when he or she includes a new artwork during the Mobile Guide. Correspondingly, the whole plan of the tour (e.g. the route, total time spent, rest of artworks in the tour) could be dynamically adjusted according to the changes. Or if an artwork is heavily crowded on view, the system might recommend an alternative tour. One step further, the limitations imposed by the fact that the user model is based on explicit ratings, we might think of observing and storing implicit user information, as for example how much time the user stands in front of a painting, and use this as an additional source for the dynamic adaption. In Chapter 4, we extend the current Mobile Guide with a real-time routing system. It sorts recommended artworks in a sequence and provides users with an efficient tour route through the museum based on their locations and the positions of recommended artworks in the museum.

A variety of Web-applications and devices. The current Mobile Guide runs on Windows Mobile. To support a larger spectrum of devices from museums and users, clients for other operating systems can be implemented. For instance to support more smart phones: a Symbian client can be developed or to support iPhones: an implementation for iOS can be created. As an exploration, we implemented the new version (with the real-time routing system) of the Mobile Guide on an iPod in Chapter 4.

Wireless communication and orientation. Wireless communication technologies such as Bluetooth or Wi-Fi can be used to share data between devices. This allows for providing social functionalities like sharing tours with friends or sharing notes about artworks (Graziola et al. 2005) in the hotspot area. Additionally, interactive maps and Location-based technologies (e.g. Infra Red, RFID, GPS, Bluetooth) can be applied to facilitate visitors orientation. In the new version of Mobile Guide, we use a wireless connection for the exchange of information.

User interaction. The user interaction of the current Mobile Guide on the PDA has been set-up primarily to be functional and usable. Special attention is dedicated to support a small touch screen controlled by a human finger instead of a stylus. In Chapter 4, we implemented the new version of Mobile Guide on an touch-screen based iPod. However, we still need to create a nice look-and-feel for users, which obviously will be a target for future work to be carried out later.

Figure 3.8: Tour Wizard: Semantic search function in the Tour Wizard

Chapter 4

Adapting Museum Tours on Handhelds

We introduced the first version of the Mobile Guide (on PDA) in Chapter 3. As a follow up work, we present in this chapter the second generation of Mobile Guide (on iPod) with a real-time routing system. This routing system is implemented based on the SWI-Prolog Space package, using : (i) the user profile containing users' preferences and current location; (ii) the semantically enriched Rijksmuseum collection and (iii) the coordinates of the artworks and rooms in the museum. In addition, we extend the user model ontology by storing the user's behaviors during the tour. In such a way, we maintain a dynamic user model which connects the user's interactions with tools online and on-site. In the evaluation, we test whether the sequence of recommended artworks in the tour follows an efficient route through the museum.

This chapter was published as "Finding Your Way through the Rijksmuseum with an Adaptive Mobile Museum Guide" in the Proceedings of the Extended Semantic Web Conference (ESWC) 2010 (van Hage et al. 2010) and was co-authored by Willem Robert van Hage, Natalia Stash, and Lora Aroyo.

4.1 Introduction

Cultural heritage and museum collections provide a wide variety of objects, which could be of interest to different visitors. To meet the diversity of preferences and backgrounds of visitors museum curators offer tours on different topics. However, these topics usually are selected based on the highlights of the collection and the resulting tours include a fixed and predefined sequence of artworks to view. An audio tour provides more freedom in determining your own sequence of artworks while visiting a museum. However, the set of artworks to choose from is still a predefined one and is the same for all visitors. Currently, museums turn to multimedia guides in order to bridge the gap between the visitor's interests and the static museum tours. Personalization is one way to provide dynamics related to

visitor's interests, which subsequently could enhance visitor's experiences (Roes et al. 2009). An adaptive mobile museum guide acts as a museum expert and provides the user with information adapted to the current situation (Kruger et al. 2007). For example, the *MIT Media Lab*[1] audio and visual narration adapts to the user's interest acquired from the physical path in the museum and length of the user stops. The mobile museum guides developed within *Hippie* (Oppermann and Specht 1999) and *PEACH* (Rocchi et al. 2004) projects provide content adaptation based on technical restrictions of specific presentation devices as well as visitor's preferences and knowledge. The difference between two projects is that Hippie museum guide uses stationary and mobile devices in a sequential way (e.g., a user prepares his museum visit on the personal computer at home and then uses the mobile device while actually visiting the museum), the PEACH museum guide combines both mobile and stationary devices in parallel. The mobile museum guide built within the *Sotto Voce* (Aoki 2002) project takes into account the special needs of groups visiting a museum and facilitates social interaction between group members. *AgentSalon* (Sumi 2004) system users are provided with mobile devices and are monitored while exploring the museum. The system can infer an overlap between users' interests and experiences and fosters communication between the users with stationary devices. *ARCHIE* (Luyten et al. 2006) provides a socially-aware handheld guide that stimulates interaction between group members. They can communicate with each other either directly (by voice) or indirectly (by collaborative games) by means of their mobile guides. By using a personal profile it allows to adapt the interface and tailor the information to the needs and interests of each individual user. The user profile evolves slowly by observing how the user interacts with the digital content, e.g. asking for more, or bookmarking it, may indicate interest while stopping an explanation prematurely may indicate a lack of it. The *Kubadji* mobile tour guide[2] aims at deriving users' interests from implicit behavior (e.g. artworks viewing times), recommendation of items of interest and personalization of the content delivered for these items via the handheld device. Besides it uses a collaborative filtering approach for predicting users' viewing times of unseen exhibits from his viewing times at visited exhibits. The context-aware museum tour guide presented in (Chou 2005) is used to give directions to the visitor and is adjusted as the tour progresses dropping one or more exhibits if the visitor falls behind the tour or suggesting additional exhibits or taking a break at a nearby restaurant if the visitor has extra time. The environment also supports peer-to-peer interactions between visitors, allowing them to find each other, share ratings and comments about exhibits. A number of museums, e.g. Tate Modern, Science Museum Boston, are already exploring the potential of personalized museum guides, currently available on their websites.

    A major bottleneck in the realization of this personalization is how to collect

---

[1]http://www.media.mit.edu/
[2]http://www.kubadji.org/

the necessary information about the user's (constantly evolving) interests (Roes et al. 2009) without intruding on the visitor too much. Typically, for large scale online access personalization can be achieved through usage of stereotypes (e.g. students, novices, art experts, children) or through deducing a user profile from observation of their online browsing and searching (or in museum viewing) behavior. In this way, personalized virtual tours are ways for visitors to construct their own narratives[3]. In addition, the indoor localization of people and objects plays a critical role in order to implement and successfully deploy such a system. Two tasks are considered in this context (Kruger et al. 2007):

- **Detecting user's location** inside the museum requires a positioning system that considers the boundaries and constraints (i.e. the walls, doors, stairs) of the physical indoor space. Methods using different hardware solutions have been proposed to increase the accuracy of the indoor user position.

- **Assessing user's context** in terms of artworks in her neighborhood, which artworks have been already seen by the users, how much time has the user already spent in the museum and additional temporal constraints (e.g. how much time is available), what are visitor's general interests in art, and potentially also their physiological and the emotional state (Kruger et al. 2007).

Having the limited resources of mobile guides in mind, most of representation and processing of relevant knowledge needs to be carried out remotely in the infrastructure. To reduce complexity and to ensure reusability of the knowledge representations and inference mechanisms a flexible web-based approach is required that allows different types of systems to exchange and augment information on users and particular situations (Kruger et al. 2007). In the following sections we discuss briefly the CHIP project, the routing mechanism of SWI-Prolog Space package and pay a special attention to the SPACE-CHIP demonstrator.

4.2 Finding routes through the Rijksmuseum

The Art Recommender supplies a list of recommended artworks that are ordered by the estimated likelihood that the user will find them appealing based on manual ratings. Even though the rooms in the Rijksmuseum have themes, such as works about the Dutch republic or works by Rembrandt and his pupils, these themes do not necessarily coincide with the preferences of the user. This means that even a small set of recommended artworks can be distributed over the entire museum. In order to improve the user experience of the museum visit, we reorder the results of the Art Recommender so that there is an efficient way to walk from one to the other. This route minimizes the walking effort, while maximizing the number of top recommendations. Also, it takes into account an optional maximal walking

---

[3]Virtual Museum (of Canada), `http://www.museevirtuel-virtualmuseum.ca/`
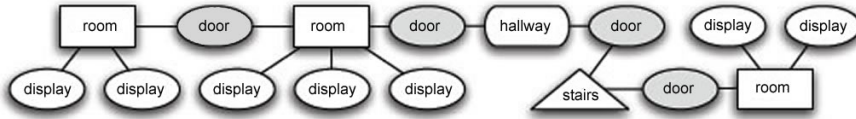
Figure 4.1: Example connectivity graph

distance and the number of artworks. This helps the user to decide where to go given limited time. (The feeling of missing something important can cause people to linger too long in the "wrong" rooms and therefore to miss their favorite works.)

An Easy Traveling Salesman Problem   Computing an efficient route through a museum is very similar to the traveling salesman problem. However, for a few reasons, theoretically at least, a significantly easier problem than the general traveling salesman problem, for which the greedy nearest-neighbor search algorithm is considered a sub-optimal solution. First, if you consider the artwork displays, rooms, doors, hallways, and stairs to be nodes in a connectivity graph, then this graph is not fully connected, because there are walls and floors in the way. This is illustrated in Fig. 4.1. Second, the rooms are considered units when the exhibits are created, which means it makes sense to view all works from a single room together. This means it is nearly always a good idea to delay transitions across doors until all displays in a room have been visited. And third, floor transitions take a lot of effort, especially up by stairs, or either way by elevator, because you have to wait for the elevator. For these reasons there are only a few sensible paths through the museum. Locations are grouped per room and then per floor. If you set the transition weight of the edges in the connectivity graph to the experienced distance instead of the actual distance then nearest neighbor search will always send the visitor to works within the same room first before making the transition to another room (or even floor), which is good in the case of the Rijksmuseum, but which is bad in the case the general traveling salesman problem, because it causes local optima.

Implementation of the Nearest Neighbor Router   The SWI-Prolog space package (van Hage et al. 2009) provides nearest neighbor search. However, this nearest neighbor search is unaware of the restrictions posed by the walls and floors. Therefore, we base our routing on a connectivity graph search algorithm that uses intersection queries as opposed to nearest neighbor queries. First we compute a connectivity graph between all the artwork displays, rooms, stairwells, etc. that takes into account where the doors are. Then we compute the weighted shortest path between all the displays. The weight is based on graph distance, the type of

the transition (*e.g.* moving to another floor is more expensive), and on the distance between locations inside a room (*e.g.* how far displays are from each other or from a door). This shortest distance matrix is used to compute an efficient path along all the recommended artworks. The exact method we use to calculate the routes is as follows:

- Pre-compute artwork distance matrix once
    1. define that stairs, hallways, toilets, are rooms
    2. define that works are on display in the museum
        (a) give the display a $\langle x, y, z \rangle$ coordinate
    3. define what it means to be connected
        (a) places (displays, doors) space_intersect with same room
        (b) places are stated to be connected (stairs to stairs on other floor) by A chip:connectsTo B
    4. assert A chip:connectsTo B for each connected pair $\langle A, B \rangle$
    5. make connectivity graph of chip:connectsTo
    6. compute weights for each transition
        (a) graph distance plus distance within room
        (b) door transitions get a higher graph distance than display-display transitions
        (c) stairs transitions get an even higher graph distance
    7. compute and cache upper triangle matrix of weighted graph shortest path distances between all places

- Apply routing algorithm for each request
    1. fetch set of recommended works (given by Art Recommender)
    2. fetch current position (given by user interface)
    3. fetch remaining time in museum (given by user interface)
    4. fetch maximum number of artworks to route (given by user interface)
    5. greedy nearest neighbor search in weighted distance graph until list of recommended works is empty:
        (a) look up nearest recommended work
        (b) remove work from list of candidates
        (c) add path from current position to work to recommended route
        (d) set current position to location of work

     (e) add length of path to total length of recommended route

  6. while total path length of recommended route takes longer than remaining time in museum

     (a) remove furthest artwork from current position

     (b) apply greedy nearest neighbor search again (step 5)

4.3   SPACE-CHIP demonstrator

Imagine the following usage scenario: Our user prepares a visit to the Rijksmuseum. He provides his opinion about a number of Rijksmuseum artworks and topics through the Art Recommender e.g. rates the painting "Woman Reading a Letter" and the artist that made the painting Johannes Vermeer with 4 stars meaning he likes them. These ratings result in the list of recommended artworks that form a Tour of Recommended Artworks that the user can view in the Tour Wizard[4].

    The user is going to follow this tour inside the museum with the help of the CHIP Mobile Guide. The routed tour is shown in Fig. 4.3. We use icons in a different color to indicate artworks that are in the tour and connect them with the tour line. The user location is indicated with an icon at the entrance door on the ground floor. During the visit the user views artworks that are in the tour but is also attracted by other artworks outside his tour. In order to give a notification to the system that the user has viewed an artwork he has to click on a corresponding icon on the museum map and in the popup window showing artwork description (see Fig. 4.2) he has to click on "Viewed" icon. If the user clicks on a "Viewed" icon for an artwork that is in his tour then the tour route remains the same. Otherwise the tour may be re-routed taking into account the user's interest in that artwork. He can also give ratings to any artwork he sees. These actions result in the tour being dynamically adapted taking into account the history of his visit (seen artworks), changing interests and current location. (However if the user wants to follow the initial sequence of recommended artworks and does not want the tour to be adapted he can select a corresponding option in the tour configuration). If the user, for example, likes the works by Frans Hals and Ferdinand Bol he comes across on his way to the recommended Johannes Vermeer works, he can add a rating by selecting the work on the map and submitting a new rating (see Fig. 4.2). This automatically updates the tour. The updated tour is shown in Fig. 4.4. For the sake of clarity we have highlighted the works from

---

[4]For the demonstration purposes we simulate the user's experience with the mobile device by showing the tour map in the Tour Wizard tool. In difference with the original version of the tool in this Tour Wizard we indicate with icons the (imaginary) artworks locations. Semantically enriched data about Rijksmuseum collection only provides information about the room number where a particular artwork is located but does not provide information about the exact artwork location.
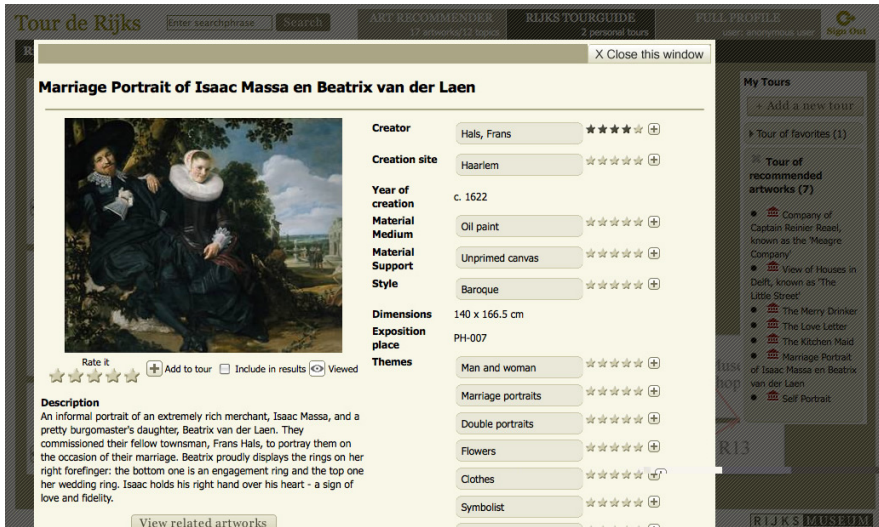
Figure 4.2: Artwork description page

the original tour with red, the new Frans Hals recommendations with yellow and the new Ferdinand Bol work with blue.

4.4 Mapping the CHIP user model to SEM

In order to provide data exchange between CHIP and the SWI-Prolog Space package we mapped the CHIP user model specified using RDF/XML to the Simple Event Model (SEM)[5] which is proposed by van Hage et al. (van Hage et al. 2009) and is just a formalization in RDF using SEM.

As shown in Fig. 4.5, we defined `chip:User` as a sub class of the `sem:Actor`, who participates in the `sem:Event`. In our case, there are three different types of events: (i) rating, (ii) viewing, and (iii) tour. In a rating event, the user rates a `sem:Object` with a `chip:ratedValue` from "-1" to "1". The viewing events are usually part of the tour events, since the user views a `sem:Object` during the tour. In a tour event, the user adds a `sem:Object` into a particular tour. All of the objects added in the tour will be ordered in a sequence based on their locations in the museum, which are described using the `rdf:_n`[6] as a sub property of the

---

[5]For this work we use this version: `http://semanticweb.cs.vu.nl/2009/04/event/`. A newer version is available at `http://semanticweb.cs.vu.nl/2009/11/sem/`.

[6]`http://www.w3.org/TR/rdf-schema/`
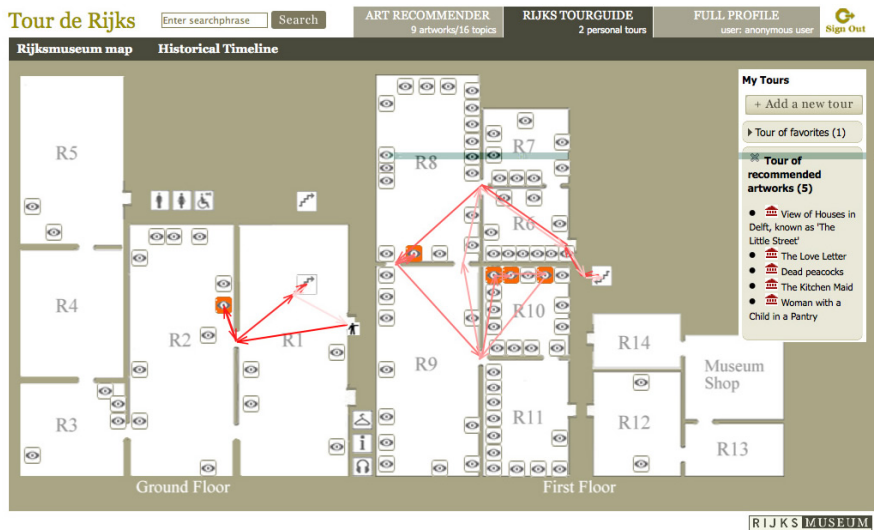
Figure 4.3: Initial route of the tour of recommended recommended artworks
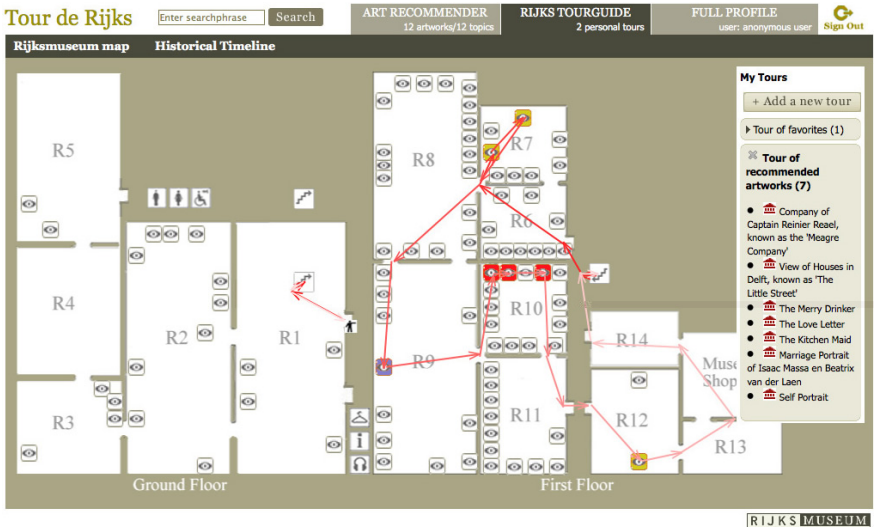
Figure 4.4: Re-routed tour of recommended artworks
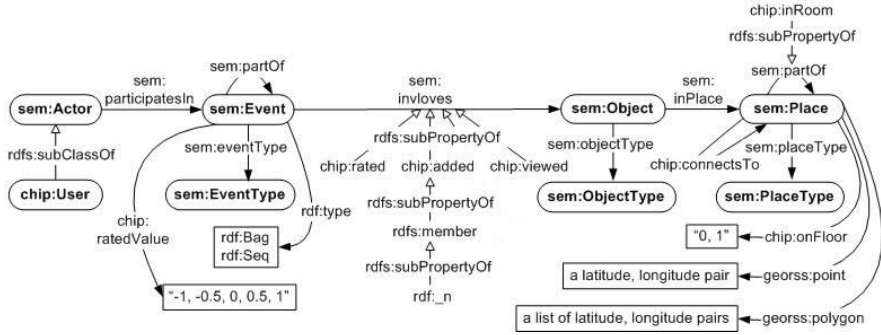
Figure 4.5: Mapping CHIP user model (UM) to the simple event model (SEM)

`rdfs:member`.

Suppose Saskia is a CHIP user who participated in three events: (i) she rated artwork_9 with a value of "0.5"; (ii) she added three artworks (artwork_5, artwork_9 and artwork_16) in the tour; and (iii) she viewed artwork_16 in this tour. Using the routing algorithm, artworks are ordered in a sequence: artwork_16, artwork_5 and then artwork_9. In Fig. 4.6 we give the corresponding code that describes these information in the user profile and tour profile.

To indicate the locations of objects in the museum, we use various properties: `chip:inRoom`, `chip:onFloor`, `georss:point`[7] and `georss:polygon`. There are also many different types of places in our case, such as display (the place type for artworks), room, door, hallway, stair, elevator, restroom, etc. Two places are connected by using the `chip:connectsTo` property.

4.5   Evaluation

We consider two issues for evaluation: (1) that recommendations are useful for the users and (2) that sequence of recommendations follows an efficient route through the museum in a reasonable time that allows real-time interaction with the system.

With respect to the first issue we performed a study *Effectiveness of recommendations, novices versus experts*. Our conclusion was that the **Art Recommender** helps novices to elicit or clarify their art preferences from their implicit or unclear knowledge about the museum collection (Wang et al. 2008b). Compared to the novices, the experts (mainly museum domian experts) do not seem to benefit from it a lot, although there is a slight increase of 0.23 (the increase for the novices is 1.18) which indicates that the system also helps the experts to elicit their art

---

[7] http://georss.org/simple

```
:Saskia a sem:Actor ;
      sem:participatesIn :event_1, :event_2, :event_3 .

:event_1 a sem:Event ;
      sem:eventType :rating ;
      chip:rated :artwork_9 ;
      chip:ratedValue "0.5" .

:event_2 a sem:Event ;
      sem:eventType :viewing ;
      sem:partOf :event_3 ;
      chip:viewed :artwork_16 .

:event_3 a sem:Event, rdf:Seq ;
      sem:eventType :tour ;
      rdf:_1 :artwork_16 ;
      rdf:_2 :artwork_5 ;
      rdf:_3 :artwork_9 .
```

Figure 4.6: Code example of the CHIP User Model based on SEM.

preferences.

With respect to the second issue, we measured the speed of the router. To determine the speed of the router we measured the CPU time taken on a 2.66 GHz Intel processor, given enough memory to store the cached distance matrix between the artworks. The result is shown in Fig. 4.7. Even though the performance curve shows exponential growth in terms of the number of artworks that need to be routed, the total time needed for the routing stays within reasonable bounds for the number of artworks in a realistic tour through the Rijksmuseum. This performance could be significantly improved by further optimizing the data structure that stores the distance matrix. At the moment this is a binary tree. An array matrix would provide faster access. Furthermore, we guarantee that the router always favors within-room transitions over between-room transitions, which in turn are always favored over floor transitions. Given the limited connectivity between the rooms and floors this guarantees an efficient path.

4.6   Discussion and future work

Existing adaptive mobile museum guides differ in the ways they construct the user model, the ways they provide personalized experience inside the museums, and the devices that they use. Many projects focus on social communication between the
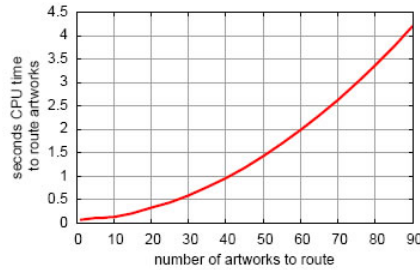
Figure 4.7: Time performance of the router for n artworks. The router needs exponential time to derive an efficient route, but stays within a reasonable time for real-time response for the number of artworks that a realistic tour comprises.

visitors (e.g. friends, group members) while following a tour. Currently, CHIP does not take the social aspect into account - neither for generating recommendations, nor for communication inside the museum. This could be one of the improvements to the CHIP demonstrator. The strong points of CHIP however are the distributed user model and the ability to view the CHIP Mobile Guide in any browser. No additional software installation is required while typically museums would provide the visitors with the PDA's running pre-installed software and ask them to provide some personal information to start creating their user models. With the presence of Wifi inside the Rijksmuseum the visitor could use his own device (iPod touch or iPhone) to follow the CHIP Mobile Guide that uses the user model and the tours information stored on the CHIP server.

We consider several directions for future work.

First of all, implementing the demonstrator for use in a realistic situation (inside the museum) with real time data. In the current prototype we simulated the user's experience with the mobile device by showing the tour map in the Tour Wizard tool. To indicate the fact that the user has seen an artwork he has to first click on a corresponding icon on the map and then on a "Viewed" icon in the popup window that opens. Time issues while following the tour in this way are not taken into account. The next step in developing the demonstrator would be the implementation of the real-time user localization and re-routing the user by taking into account the time that he spends viewing artworks, moving between artworks, taking stairs, etc.

Secondly, designing and evaluating the user interface for guiding the user in a realistic situation. At the moment the user is only provided with the museum map that indicates the tour route and the current user location. Based on the map the user has to figure out where to go next. It would be interesting to consider the possibility of guiding the user "locally" by an indication about where to go

next from the current point, like turn left/right, etc. It would be also interesting to consider using technologies like Google Goggles[8] to show information about an artwork when the user points with his device to it.

And thirdly, experimenting with various re-routing algorithms. The current algorithm can provide re-routing of a tour or sequencing of a given set of artworks from the tour (generated by CHIP Art Recommender and Tour Wizard tools) based on the user's position or closeness to a certain artwork from the tour. In addition, the routing mechanism uses the museum coordinates. It does not take into account the information from the user model. It would be interesting, however, to consider more complex algorithms that would also take user preferences into account, and possibly decide to add more artworks to the tour that might be interesting for the user based on the user's closeness to them.

---

[8]http://www.google.com/mobile/goggles/

Chapter 5

Enhancing Recommendations using Semantic Relations

In Chapter 2 we presented the first implementation of the content-based recommendation (CBR) algorithm in the Art Recommender. To deal with the semantic complexity in the collections, we identify in this chapter a number of semantic relations within one vocabulary (e.g. *borader/narrower*) and across multiple vocabularies (e.g. *style, birthPlace*). We evaluate which relations are useful for CBR in terms of the number of recommended items and precision. We explore navigation patterns of users. The results give a first insight and help us derive preliminary weights for each semantic relation that we will discuss in Chapter 6.

This chapter was published as "Semantic Relations for Content-based Recommendations" in the Proceedings of the International Conference on Knowledge Capture (K-CAP) (Wang et al. 2009b) and was co-authored by Natalia Stash, Lora Aroyo, Laura Hollink and Guus Schreiber.

5.1 Introduction

The main objective of the CHIP (Cultural Heritage Information Personalization) project is to demonstrate how Semantic Web and personalization technologies can be deployed to enhance access to digital collections of museums. In collaboration with the Rijksmuseum Amsterdam[1], we have developed the CHIP Art Recommender: a content-based recommender system that recommends art-related concepts based on user ratings of artworks. For example, if a user gives the famous painting "Night watch" a high rating, the user will get its *creator* "Rembrandt" recommended.

The semantic enrichment of the Rijksmuseum InterActief (ARIA)[2] database enables the opportunity to recommend a wide range of concepts via different se-

---

[1]http://www.rijksmuseum.nl
[2]http://www.rijksmuseum.nl/collectie/ontdekdecollectie

mantic relations. These relations link concepts not only within one vocabulary (e.g. *teacher/studentOf, broader/narrower*), but also across two different vocabularies (e.g. *hasStyle, birth/deathPlace*). For example, if a user likes the artist "Rembrandt", the system could recommend his *teacher* "Pieter Lastman" and his art *style* "Baroque", or even its *narrower* concept "Renaissance-Baroque styles and periods" and its *broader* concept "European styles and periods".

However, for recommender systems, the use of semantic relations also poses a problem. Not all related items are useful or interesting for end users. If the user likes the artist "Rembrandt", besides his teacher and art style, the system could also recommend his *death place* "Amsterdam" or even the broader geographic location "Noord-Holland", which might not be of interest for users. Thus, our main challenge is to find which semantic relations are generally useful for content-based recommendations. Furthermore, we aim to derive the navigation patterns in order to improve the accuracy of recommendations. Our hypothesis is that by choosing specific semantic relations, the recommender system could retrieve more related items without decreasing the accuracy and interestingness. In the experiment, we tested the Art Recommender with end users by applying both artwork features and semantic relations to recommend related concepts. Using artwork features as a baseline, we compared the recommendations via different semantic relations in terms of accuracy and interestingness.

The chapter is organized as follows: Section 2 presents related work about the use of semantic relations for recommender systems. Section 3 briefly introduces the metadata vocabularies and identifies a number of semantic relations as well as artwork features. In Section 4 we describe our demonstrator, a content-based art recommender system and explains the design of the experiment. Section 5 discusses the results. We conclude and discuss the future work in Section 6.

5.2   Related work

In recent years, many recommender systems have appeared that use Semantic Web technologies, where information is well-defined in an open standard format that can be read, shared and exchanged by machines across the Web (Berners-Lee et al. 2001). Peis et al (Peis et al. 2008) classified semantic recommender systems into three different types: (i) vocabulary or ontology based systems; (ii) trust network based systems constructed with FOAF[3]; and (iii) context-adaptable systems that use additional ontologies about e.g. the current time, place of the user. In this chapter, we focus on the first type (vocabulary-based recommender systems) and discuss how various semantic relations to enhance recommendations.

Metadata vocabularies or domain ontologies are so far mainly used for content-based recommender systems. the CULTURESAMPO portal (Ruotsalo and Hyvonen 2007) recommends images based on semantic relations between selected images

---

[3]Friend of A Friend: http://www.foaf-project.org/

and other images in the repository. In particular, they used the *has-part/part-of* relations with a fixed weight to determine the ontological relevance of recommendations. A similar approach is adopted in the ConTag project (Adrian et al. 2007), which extracts similar topics using the *broader/narrower* relations for recommendations. By knowing user preferences, Blanco-Fernández (Blanco-Fernández et al. 2008) inferred semantic associations between user preferences and relevant instances from the domain ontology in order to provide personalized recommendations of TV programs.

In CHIP we have developed a content-based recommender system, the Art Recommender. Compared with the content-based recommender systems mentioned above, the Art Recommender works with four different semantic metadata vocabularies (see Section 3), which provide richer semantic relations: not only hierarchical relations such as *broader/narrower* within one vocabulary, but also more sophisticated relations across two different vocabularies, e.g. *hasStyle* and *birth/deathPlace*. These semantic relations might be helpful to partially solve the cold-start and over-specialization problems for content-based recommender systems. For example, (i) when there are few ratings, the system could use semantic relations to provide additional concepts; (ii) the use of semantic relations within one vocabulary or across multiple vocabularies might retrieve new concepts, which might be surprising or interesting for users.

5.3   Identifying semantic relations

The CHIP Art Recommender currently works with the Rijksmuseum ARIA database, containing images and metadata descriptions of artworks. The mapping of metadata from ARIA to Iconclass[4] and to the three Getty thesauri[5] (the Art and Architecture thesaurus (AAT), the Union List of Artists Names (ULAN) and the thesaurus of geographic Names (TGN)) brings rich semantic structure to the Rijksmuseum collection and creates the opportunity to recommend a wide range of art concepts via various semantic relations (Wang et al. 2008b). As shown in Fig. 5.1, we listed 4 basic artwork features (Relations 1-4) which link an artwork and its associated concepts, as well as 11 semantic relations (Relations 5-15), which link concepts within one vocabulary and across two different vocabularies.

Relations 1-4 are artwork features, which have already been implemented in the original Art Recommender for the inference of recommended concepts. As an example, if a user likes the artwork "Night watch", we could recommend the *creator* "Rembrandt" from ULAN, the *creation site* "Amsterdam" from TGN, the *material* "Oil painting" from AAT, the *subjects* "Cloth" from Iconclass and "Wealth in the Republic" from ARIA.

---

[4]http://www.iconclass.nl/libertas/ic?style=index.xsl
[5]http://www.getty.edu/research/conductingresearch/
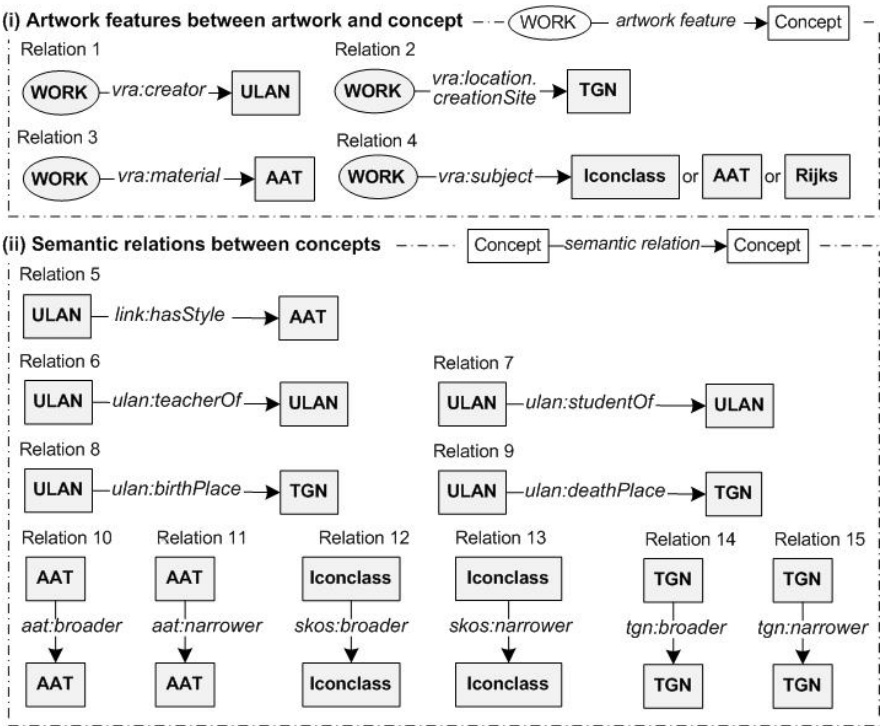
Figure 5.1: Overview of artwork features and semantic relations based on the metadata vocabularies

Relations 5-15 are semantic relations linking concepts within one vocabulary and across two different vocabularies. We applied these semantic relations in the experiment in order to get insights into which relations are useful for content-based recommendations. In more detail, Relation 5 (*link:hasStyle*) links an artist to his/her art style(s), across the ULAN and AAT vocabularies, e.g. "Rembrandt" in ULAN has an art style "Baroque" in AAT. Relations 6 and 7 are the *ulan:teacher/studentOf* relations linking two concepts within the ULAN vocabulary. For example, "Rembrandt" is the teacher of "Gerrit Dou" and the student of "Pieter Lastman". Relations 8 and 9 are the *birth/deathPlace* relations between artists and geographical locations where she was born or died, across the ULAN and TGN vocabularies, e.g. "Rembrandt" in ULAN was born in "Leiden" in TGN, and died in "Amsterdam" in TGN. Relations 10-15 are the general *broader/narrower* relations within the AAT, Iconclass and TGN vocabularies. Although the relations are the same, the types of concepts mapped to the three vocabularies are different: (i) concepts mapped to AAT are mainly art styles, e.g. "Rococo revival" has a broader concept "Modern European revival styles", (ii) concepts mapped to Iconclass are general subjects, e.g. "Musical" has a narrower concept "Music instruments" and, (iii) concepts mapped to TGN are geographic locations, e.g. "Amsterdam" has a broader concept "Noord-Holland".

5.4   Evaluation

Our goal is (i) to investigate which semantic relations are useful for content-based recommendations in comparison with standard artwork features, and (ii) to look at the combined use of semantic relations and artwork features in sequence, which might derive some navigation patterns from users in order to enhance the accuracy of recommendations and to be reused for other recommender systems.

5.4.1   Target system: Art Recommender

To address these goals, we applied both artwork features and semantic relations for content-based recommendations of art concepts in the Art Recommender[6]. Considering artworks are recommended based on related/recommended art concepts, in order to get a clear insight, we only looked at how semantic relations and artwork features influence related/recommended art concepts in this experiment. We leave the exploration of how they affect related artworks for recommendations as a next step in future work.

The user interface of the Art Recommender (see Fig. 5.2) was split in two parts: the upper part is the rating dialog with a slide show of artworks, which allows the user to browse artworks from the collection and give ratings to them with 1-5 stars (i.e. I hate it, I dislike it, neutral, I like it, and I like it very much).

---

[6]http://www.chip-project.org/demoUserStudy3/
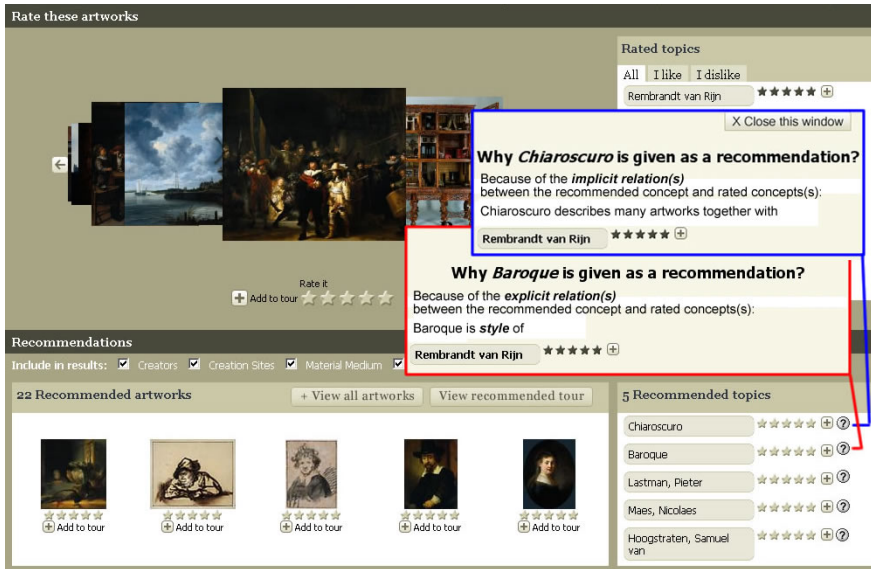
Figure 5.2: Explanation for recommended concepts in the Art Recommender

In the bottom part recommended concepts are shown, based on the ratings given by users to the artworks in the upper part. Then the user rates (with 1-5 stars) the recommended concepts shown in the bottom part to express how much she likes each recommendation. The list of recommended concepts will be dynamically updated based on the last rating given for an artwork or concept. In addition, in the "Why recommended" option (see Fig. 5.2), an explanation is provided about which feature or relation was used for each recommendation. The user is then asked to give 1-5 stars indicating how interesting she finds the concept recommended via this feature or relation (*interestingness*). This dimension of interestingness puts the recommended concept back in context, which helps user to understand the inference of recommendations by using particular artwork feature(s) or semantic relation(s).

5.4.2   Method

At the beginning of each session, participants were asked to fill out a questionnaire about: (i) their age, (ii) whether they are familiar with the Rijksmuseum collection, (iii) experience with recommender systems in general, (iv) expectation from art recommendations, and, (v) for what purpose they will use art recommendations.

After completing the questionnaire, we briefly introduced the Art Recommender and explained the recommendation process. Using the Art Recommender, users were asked to follow two steps:

Step 1 (Pre-task): to find an artwork that she likes from the artwork slide show (to start the process the user needs to give a rating of either 4 or 5 stars; the recommender does not start-up with negative ratings). As a baseline, it will recommend the first set of related art concepts by applying artwork features based on the rated artwork.

Step 2 (Main task): to rate the first set of recommended concepts. Based on the ratings of concepts, the system will produce a second/new set of recommended concepts by applying semantic relations, which also allows users to rate. At any point for each recommendation the user can click on "Why recommended" and give her feedback on whether she finds this recommendation via the particular artworks feature or semantic relation interesting or not on a 5-degree scale. Step 2 gave us an insight in the performance of the concepts recommended via semantic relations in comparison with the concepts recommended directly via artwork features.

Users were asked to repeat this process for at least 5 times in order to rate enough recommended concepts via either artwork features or semantic relations. At any point, the user could stop rating recommended concepts and go to select another artwork from the slide show. Then the same process is repeated for each rated artwork.

5.4.3 Dimensions and metrics

Using artwork features as a baseline, we tested the results of recommended concepts via semantic relations in terms of two dimensions: *accuracy* and *interestingness*.

- *Accuracy*: by directly asking the user whether she likes this recommended concept, which is shown as "Ratings" in the Art Recommender in Fig. 5.2.

- *Interestingness*: by giving the explanations of "Why recommended", it asks the user whether she finds the concept recommended via the particular artwork feature or semantic relation interesting.

Precision, Recall and Mean Absolute Error (MAE) are most popular metrics to evaluate recommender systems (Burke 2002; Herlocker et al. 2004) and to measure the usefulness of semantic relations in query expansion for information retrieval systems (Hollink et al. 2007; Navigli and Velardi 2003; Tudhope et al. 2006). *Precision* represents the probability that a recommended item is relevant, *Recall* represents the probability that a relevant item will be recommended, and *MAE* measures the average absolute deviation between a predicted rating and the users' true rating (Herlocker et al. 2004).

However, in our case, we could only apply *precision*, but not *recall* and *MAE*. Because it is difficult to determine the total number of relevant items. As Burke

discussed in (Burke 2002), relevance is subjective from an end user's standpoint and it often changes when the user gets explanations for recommendations. As Herlocker discussed in (Herlocker et al. 2004), it is also not appropriate in our case to use *MAE*, where a list of recommended concepts is returned but users often only view concepts that she is interested and cares about errors in concepts that are recommended. Thus in the experiment we only use precision to measure accuracy and interestingness for recommended art concepts. To divide the concepts into relevant or irrelevant concepts, we defined a threshold value on the used 5-star scale, which converts 4 and 5 stars to "relevant" and 1-3 stars to "not relevant". In terms of accuracy, relevant concepts refer to the recommended concepts that the user likes with 4 and 5 stars, and in terms of interestingness, relevant concepts refer to the recommended concepts that the user finds interesting with 4 and 5 stars. Below we explain how we calculate it:

$$Precision = \frac{Correct\ Hits}{Total\ Rec.Rated}$$

Correct Hits is the total number of relevant concepts that are recommended by the system and have been rated by the user with 4 and 5 stars in terms of accuracy and interestingness respectively.

Total Rec.Rated is the total number of concepts that are recommended by the system and have been rated by the user with 1 to 5 stars in terms of accuracy and interestingness respectively. *Total Rec.* is the number of all recommended concepts with or without user ratings. To avoid the data sparsity problem (Burke 2002) (i.e. the number of recommended items far exceeds what a user can rate), we only use the number of "Total Rec.Rated" to compute the precision and we do not include the number of "Total Rec.", because we do not have user feedback on concepts without ratings (Herlocker et al. 2004). However, we will provide the number of "Total Rec." (in Table 1) to get an idea of how many concepts could be recommended via an artwork feature or a semantic relation.

5.4.4   Results

In a period of three weeks, in total 48 users participated. The experiment took about 20-35 minutes per person. Each user gave on average 53 ratings for artworks and concepts. Below we describe the participants characteristics collected with the questionnaire.

- *Age*: in the categories of 20-30 years old (65%) and 30-40 years old (21%)

- *Familiar with the Rijksmuseum collection*: not familar with the collection (27%) and a bit familiar with the collection (46%)

- *Experience with recommender systems in general*: every few months using recommender systems, such as Amazon.com (68%)

- *Expectation from art recommendations*: want to get accurate art recommendations that match their art preferences (79%) and interests (83%)

- *For what purpose will use art recommendations*: want to keep up-to-date with new information about artworks/concepts (93%), to reflect on what has been seen in the museum (75%), and to understand her art interests better (79%)

Table 5.1: Experiment results for artworks features and semantic relations

Nr.	Artwork features/ Semantic relations	Total Rec.	Accuracy			Interestingness		
			Total Rec.Rated	Correct Hits	Precision	Total Rec.Rated	Correct Hits	Precision
			Artwork features					
1	vra:creator	332	111	74	0.67	97	80	0.82
2	vra:location.creation Site	182	83	33	0.40	61	34	0.56
3	vra:material	159	92	39	0.43	47	21	0.45
4	vra:subject	3245	1054	528	0.50	768	453	0.59
1-4	all artwork features	3918	1340	674	0.50	973	588	0.60
			Semantic relations					
5	link:hasStyle	82	38	24	0.63	46	34	0.73
6	ulan:teacherOf	291	135	57	0.43	127	90	0.71
7	ulan:studentOf	92	55	24	0.44	67	46	0.68
8	ulan:birthPlace	184	44	14	0.32	48	21	0.43
9	ulan:deathPlace	130	42	11	0.26	55	14	0.25
10	aat:broader	69	23	12	0.53	19	11	0.60
11	aat:narrower	125	31	17	0.55	26	16	0.62
12	skos:broader	404	224	112	0.50	131	67	0.51
13	skos:narrower	1198	506	263	0.52	425	213	0.50
14	tgn:broader	82	22	5	0.22	15	2	0.15
15	tgn:narrower	1204	35	6	0.16	23	3	0.13
5-15	all semantic relations	3861	1155	524	0.45	1007	533	0.53

Table 1 gives an overview for artwork features and semantic relations. We calculated: (i) Total number of recommended concepts, (ii) total number of recommended and rated concepts, (iii) correct Hits (recommended and rated with 4 or 5 stars); and, (iv) precision for *accuracy* and *interestingness* respectively.

As a baseline, artwork features provide in total 3918 recommended concepts and reach an average precision of 0.50 for accuracy and 0.60 for interestingness. In comparison, semantic relations bring 3861 new recommended concepts and reach an average precision of 0.46 for accuracy and 0.53 for interestingness, which are only slightly lower than artwork features. For the individual artwork features and semantic relations, we found that:

(i) Artwork feature *vra:creator* and semantic relations *link:hasStyle* and *aat:broader/narrower* produce the most accurate recommendations and they are also the most interesting relations from the users' point of view. This could be explained by observing that artist and art style (concepts in ULAN and AAT) are intrinsically related to the artworks and an important reason why people might like an artwork or related artworks.

(ii) Semantic relations *ulan:birth/deathPlace* and *tgn: broader/narrower* that recommend geographic locations perform very badly. In particular, the *tgn:broader/ narrower* relations have the least values for accuracy and interestingness. To understand why *tgn:broader/narrower* and *ulan:birth/deathPlace* relations perform "so badly", we looked at the experiment data in detail. For example, many users like the artist "Rembrandt", however, in most cases they found his birth place "Leiden" and his death place "Amsterdam" not relevant. In comparison, users like recommended concepts such as his art styles, his teacher(s) and students(s). Another example, "Utrecht" is also a popular concept often rated with high scores, but its narrower location "Vianen" is always rated as a not-relevant concept, since it is unfamiliar to most users. This suggests that, for art recommendations, semantic relations *tng:broader/narrower* and *ulan:birth/deathPlace* might not be useful or interesting for users because they are not intrinsically related to artworks but only to locations or artists. This might also explain why users rarely rated locations recommended via these relations (with a low number of *Total Rec.Rated*). In comparison, artwork feature *vra:creationSite* gives much better results, probably it is more related to artworks.

(iii) Artwork feature *vra:subject* and semantic relations about subjects *skos:broader/narrower* produce the largest number of recommended concepts and correspondingly resulted in most user ratings. With respect to accuracy and interestingness, they score on the average.

To explore potential correlations between accuracy and interestingness, in Fig. 5.3, we plotted these two dimensions for artworks features and semantic relations. Interestingly, there is a strong positive correlation between accuracy and interestingness (Peason's $R=0.89$, p value<0.01). This means that for an artwork feature or semantic relation, the more accurate recommendations it produces, the more interesting users find the recommendations, and vice-versa. An exception here is the semantic relation *ulan:teacher/studentOf*. As shown in Table 1, although the accuracy precisions for these two relations are slightly lower (0.43, 0.44) and the interesting precisions for them are very high (0.71, 0.68). This explains why semantic relations could partially solve the over-specialization problem (see Section 2) by recommending surprising or interesting items, even though the recommendations are not always quite accurate.

The setup of the experiment gives us an opportunity to look at the combined use of artwork features and semantic relations in sequence. As explained in Section 5, every positively rated artwork/concept resulted in a new set of recommended
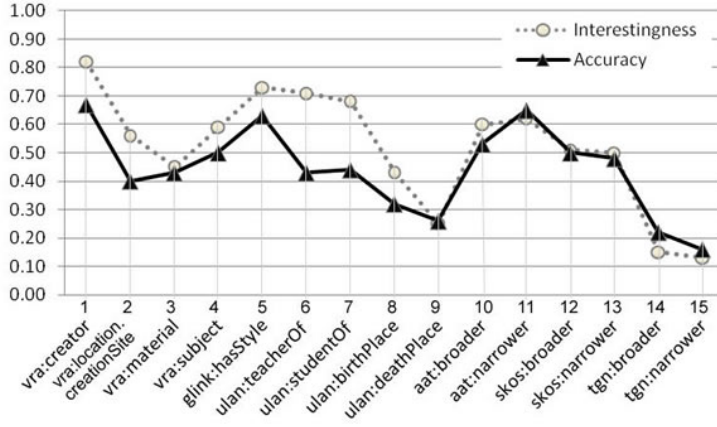
Figure 5.3: Correlation between accuracy and interestingness

concepts that the user could rate. In theory this process can go on until no new recommendations are found, but in practice most users stopped after three or four steps (Hollink et al. 2007). These sequences of ratings allow us to examine the quality of recommendations based on sequences of semantic relations and artwork features.

We first removed all sequences for which we have less than 10 user ratings. From our previous user studies (Wang et al. 2007), 10 ratings seems to be a minimum to get a reliable estimate of the quality of recommendations. We then calculated the mean of accuracy precision and interestingness precision (P_{mean}) for the remaining features and relations. Fig. 5.4 shows the sequences of recommended concepts that received more than 10 ratings, and their P_{mean} values at each step. From Table 1, we can calculate that the P_{mean} is 0.55 for all artwork features and 0.49 for all semantic relations. Using these two values as references, in Fig. 5.4 we highlighted artwork features (used in Step 2) that have a P_{mean} greater than 0.55 in black and semantic relations (used in Step 3 and 4) that have a P_{mean} greater than 0.49 in grey. Interestingly, we found three potentially useful navigation patterns of combined artwork feature and semantic relations:

- artwork -> creator -> style -> broader/narrower styles

- artwork -> creator -> teacher/student -> styles

- artwork -> subject -> broader/narrower subjects

We observe that all three patterns show a decrease of P_{mean} in each step, which might be due to the fact that the concepts are gradually more remote from the
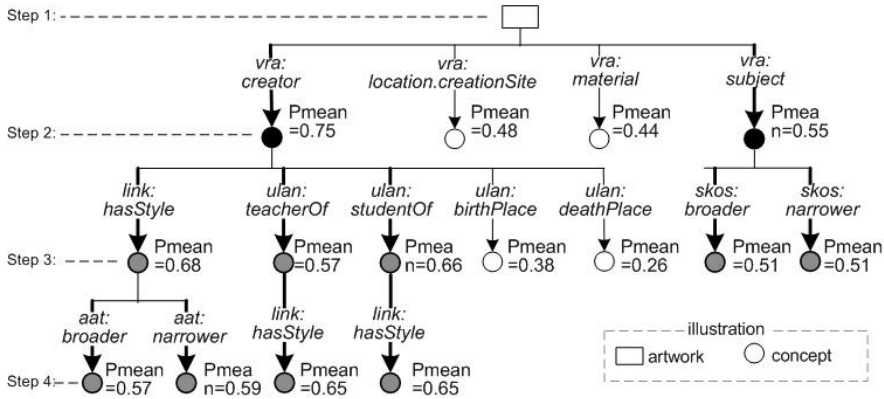
Figure 5.4: Combining artwork features and semantic relations in sequence

artwork. The only exception is Step 4 in Pattern 2 (from teachers and students to art styles). Still, at each step in the patterns, the P_{mean} value remains relatively high above the average. The three patterns could potentially be used to recommend remotely linked concepts without asking users' explicit feedback/ratings on each step. For example, if a user likes the artwork "Night watch", following the second pattern, it could recommend concepts "Rembrandt" (*creator*), "Pieter Lastman" (*teacher*), "Renaissance" (*the teacher's art style*), "Gerrit Dou" (*student*), and "Baroque" (*the student's art style*), without explicitly asking the user to rate "Rembrandt", "Pieter Lastman" and "Gerrit Dou".

5.5   Discussion and future work

Metadata vocabularies provide rich semantic relations that can be used for recommendation purposes. We examined the performance of both semantic relations and artwork features with the content-based CHIP Art Recommender in terms of accuracy and interestingness. Our results demonstrate that artwork features (*vra:creator*) and semantic relations (*ulan:teacher/studentOf, link:hasStyle*) that recommend concepts in the ULAN and AAT vocabularies produce the most accurate recommendations and also give the most interesting recommendations from the users' point of view. This might be due to the fact that these artwork features and semantic relations which recommend concepts in domain-specific vocabularies are closely related to the domain of art. In comparison, semantic relations considering geographic locations in TGN (e.g. *tgn:broader/narrower, ulan:birth/ deathPlace*) score very low on both accuracy and interestingness. A similar ob-

servation applies to the TGN vocabulary, which is a relatively much more general vocabulary and not related to the art domain, in comparison with the ULAN and AAT vocabularies.

Based on the performance of individual semantic relations and artwork features, we derived optimal navigation patterns of combined features and relations with multiple intermediate concepts. These patterns can potentially be used to effectively recommend indirectly linked concepts without asking the user's explicit feedback for the intermediate concepts.

In summary, we found that vocabularies which are relatively close to the domain are usually more useful for content-based recommendations than vocabularies, which are more general. In particular, for recommender systems in the domain of art, ULAN and AAT vocabularies which contain concepts about artists and art styles proved to be more useful for art recommendations than the TGN vocabulary which contains concepts about geographic locations. We may conclude that the use of specific semantic relations can enhance content-based recommendations by (i) retrieving more related concepts, which partially solves the cold-start problem; (ii) providing more interesting or surprising recommended concepts by using combinations of artwork feature and semantic relations, which partially solves the over-specialization problem.

As a preliminary result, the three navigation patterns we derived from the experiment might be very interesting for both users and recommender systems in similar domain of art. For future work, we are primarily interested in association rule mining and decision trees that may produce optimized results. For example, some internal nodes of the presented patterns may be pruned.

In addition, we plan to investigate the weights for different semantic relations based on the user ratings collected from the experiment. These weights can be used in computing predicted values for recommended concepts. For example, if a user likes "Rembrandt", recommendations of his *student* "Gerrit Dou", his *art style* "Baroque" or his *death place* "Amsterdam" would receive different predicted values based on the different weights of the semantic relations. The predicted values of recommended concepts can then be used to determine the predicted values for recommended artworks. In this way, we will gain insights about how the various semantic relations influence both recommended concepts and artworks. Inspired by the work from Mobasher (Mobasher et al. 2004), Ruotsalo and Hyvönen (Ruotsalo and Hyvonen 2007), the weight for each relation should not be a fixed value but a dynamic value which is calculated according to several factors, e.g. the relevance of a concept with respect to an artwork *TD-IDF* (Baeza-Yates and Ribeiro-Neto 1999), the number of user ratings of a particular artwork or concept, and the semantic distance or similarity between two concepts by using latent semantic index (LSI) (Berry et al. 1995). We further study these issues in Chapter 6. Based on the evaluation, we give a preliminary weight for each specific relation and apply these weights in a hybrid recommendation algorithm.

Our findings about which semantic relations are most beneficial to recommendations and our future work about applying weights for various relations could also be used for collaborative filtering recommender systems. For example, Mobasher's work (Mobasher et al. 2004) shows that well-selected semantic relations can be used to populate related items in order to compute the similarity between users for collaborative filtering recommender systems. This might be helpful to partially solve the cold-start and sparsity problems for recommender systems in general. Following this direction, we could apply the method of calculating the weights for various semantic relations in the recommender system and try different recommendation strategies (e.g. content-based, collaborative filtering and the hybrid approach) in order to compare the quality of recommendations in a large scale quantitative experiment. As an example, we collaborated with the Kubadji[7] project to explore a hybrid algorithm combining both content-based and collaborative filtering algorithms. Our idea is to propagate the user's ratings to related objects in the user profile, and then based on the extension of overlaps between users' profiles, to compute the similarities between users for recommendations. However, the results are not sufficient due to several reasons, which we will further explain in Chapter 8.

---

[7]http://hum.csse.unimelb.edu.au/kubadji/

Chapter 6

Defining Inference Steps for Semantically-Enhanced Recommendations

As a follow-up work of Chapter 5, we focus in this chapter on how we can use the semantic structure to enhance the content-based recommendations (CBR). On top of it, we aim to define reusable inference steps for such semantically-enhanced recommender systems from a perspective of knowledge engineering. For the semantic enhancement of CBR, we propose a hybrid approach combing both explicit and implicit recommendations based on the semantic structure in the collections. We evaluate our approach in terms of recommendation accuracy and discuss the added values of providing serendipitous recommendations and supporting explanations for recommended items.

This chapter was published as "Enhancing Content-based Recommendations with the Task Model of Classification" in the proceedings of the Knowledge Engineering and Knowledge Management by the Masses (EKAW) Conference 2010 (Wang et al. 2010) and was co-authored by Shenghui Wang, Natalia Stash, Lora Aroyo and Guus Schreiber.

6.1 Introduction

In recent years, the Semantic Web has put great effort on the reusability of knowledge. However, most work deals with reusable ontology and ontology patterns [1], there is hardly any work on reusable reasoning patterns, except the work from van Harmelen and ten Teije (van Harmelen et al. 2009). They made a first attempt at finding reusable task types and decomposing these tasks into a number of primitive reasoning patterns for Semantic Web applications. In CHIP, we collaborate with the Rijksmuseum Amsterdam[2] and built an art recommender system based

---

[1] http://ontologydesignpatterns.org
[2] http://www.rijksmuseum.nl

on the semantically-enriched collection with the mappings to standard vocabularies (Wang et al. 2008b). Inspired by the work from van Harmelen and ten Teije, we pose the question: can we identify reusable knowledge elements that can help designers of such recommender systems on the semantic web? In this chapter, we address the following research challenges:

(i) *Finding reusable inference steps for recommender systems based on rich semantic vocabularies*

As a first attempt, we analyze our demonstrator[3] (called the "CHIP Art Recommender") and identify several tasks, e.g. browse, search and content-based recommendation. In this chapter, we focus on the task of content-based recommendation and decompose this task into a number of inference steps: realization, classification by concepts, classification by instances, and retrieval.

(ii) *Bridging the vocabulary gap*

For the semantic enrichment of museum collections, most concepts of artworks have been mapped to common vocabularies for semantic-based knowledge representations (Hyvonen et al. 2005)(Schreiber et al. 2008). However, because of the complexity of the museum collections, it still contains many concepts/terms that can not be mapped to common vocabularies. These unmapped concepts are often described in non-standard schemas or in different languages. In this context, how can we bridge the discrepancy between the semantically-structured data and the remaining unstructured/unmapped data? How can we combine data from these two parts for recommendations?

To address this issue, Isaac et al. (Isaac et al. 2007) proposed a method of instance-based ontology matching. The basic idea is that the more significant the overlap of instances/artworks of two concepts is, the closer these two concepts are, and the level of significance is calculated by the corrected Jaccard measure (Isaac et al. 2007). We adopted their method in our system to build an implicit relation between two concepts even though there are no explicit semantic relations annotated between them. In such a way, most of unmapped concepts are linked with mapped concepts via implicit relations and this allows for further inference.

(iii) *Improving accuracy, serendipity and explanation for recommendations*

It hardly needs arguing that the semantic enrichment of collections could retrieve more related items (Wang et al. 2009b). However, we still face the issue of how to maintain a relatively high accuracy for recommended items. This problem becomes even more complicated when there are multiple explicit and/or implicit relations involved for a recommended item, how can we still compute an accurate prediction for this item in a way suiting the user's art preference? Besides the accuracy, there are some other issues that also affect the user's satisfaction to

---

[3]http://www.chip-project.org/demo/

a recommender system, e.g. recommending unexpected and new/unknown items
(serendipity) and providing users an insight in the logic underlying the recommen-
dations (explanation) (Herlocker et al. 2004).

To compute the prediction for related items, Ruotsalo and Hyvönen calculate
the relevance of a concept with respect to an artwork using the TF-IDF metric
and give default weights for the general *braoder/narrower* relations (Ruotsalo and
Hyvonen 2007). Mobasher et al. propose the combination of user ratings of a
particular artwork/concept and the semantic distance or similarity between two
concepts by using latent semantic index (LSI) (Mobasher et al. 2004). Our ap-
proach is to calculate the values from explicit and implicit relations separately,
and then combine the values from these two parts by setting a parameter α. By
tuning the value of α (between 0 and 1), we could change the strength of explicit
(obvious) and implicit (serendipitous) recommendations. In addition, we develop
a "Why recommend" function for each recommended item, explaining the various
relations between the user's rated items and the recommended item.

The chapter is organized as follows: In Section 2, we identify task types and
corresponding inference steps. In Section 3, we explain the semantic-enhanced
recommendation strategy. Further, in Section 4, we test our strategy with the
CHIP Art Recommender in terms of accuracy, serendipity and explanation. We
conclude and discuss the future work in Section 5.

6.2   Task types and inference steps

The CHIP Art Recommender contains three different tasks: (i) browse, (ii) search,
and (iii) content-based recommendation. For the first two tasks (browse and
search), we adopted the definitions from van Harmelen and ten Teije (van Harme-
len et al. 2009). In this chapter, we focus on defining the third task (content-based
recommendation) and analyzing the corresponding inference steps.

6.2.1   Defining the task of content-based recommendation

The standard content-based recommendation (CBR) usually takes the user profile
plus the domain ontology and returns a set of instances, which might be of interest
to the user (van Harmelen et al. 2009). In the case of CHIP, the system stores
the user profile in the form of both a set of concepts and a set of instances. Based
on the user profile and the domain ontology, it recommends both related concepts
and instances via various relations from the collection.

As described in Table 1, we use formal preliminaries: a terminology T is a set
of concepts c organized in a hierarchy. Instance i is a member of such concepts
c and this is described as $(i, \in, c)$ where $\in$ refers to the membership relation.
An ontology O consists of a terminology T and a set of instances I. Sometimes
we write (T, I) instead of O if we want to refer separately to the terminology

Table 6.1: The task of content-based recommendation

Input:	a user profile characterized as both a set of instance $I_{profile}$ and a set of concepts $C_{profile}$
Knowledge:	an ontology O = (T, I) consisting of a terminology T and an instance set I
Output:	a set of related concepts $(C^i \cup C^j \cup C^k)$ with C^i: Recommend$(I_{profile}$, O) = {(i, $\in$, c^i) | $\exists$i: i $\in I_{profile} \wedge$ i $\in c^i$} C^j: Recommend$(C_{profile}$, T) = {($c^j \sim$ c) | $\exists$c: c $\in C_{profile} \wedge c^j \sim$ c} C^k: Recommend$(C_{profile}$, O) = {($c^k \simeq$ c) | $\exists$c: c $\in C_{profile} \wedge c^k \simeq$ c $\wedge$ i $\in$ c $\wedge$ i $\in c^k$} and a set of related instances I' with I': Recommend$(C_{profile}, C^i, C^j, C^k$, O) = {(i', $\in$, c')| c' $\in (C_{profile} \cup C^i \cup C^j \cup C^k) \wedge$ i' $\in$ c'}

and the instance set of the ontology. In the case of CHIP, instances refer to artworks and each artwork is described with a number of concepts. Based on the semantically-enriched Rijksmuseum collection, we specify three different kinds of relations between artworks and concepts: (i) artwork feature, (ii) semantic relation, and (iii) implicit relation.

(i) Explicit relation (or called "artwork feature") between an artwork/instance and a concept, denoted as $(i, \in, c)$. For example, the artwork "The Night Watch" is related to the concept "Rembrandt van Rijn" via the artwork feature "*creator*". In CHIP, we apply three artwork features for recommendations: *creator, creationSite* and *subject*. Each of them has a reverse relation, e.g. *creatorOf, creationSiteOf* and *subjectOf*.

(ii) Explicit relation between two concepts with a direct link (or called "semantic relation"), denoted as $(c_i, \sim, c_j)$. In CHIP, most art concepts from the collection are mapped to the standard Getty vocabularies[4] (ULAN, AAT and TGN) and the Iconclass thesaurus[5], which provides a rich semantic structure for further inference (Wang et al. 2008b). Among various semantic relations between concepts, there are domain-specific relations within one vocabulary (e.g. *teacherOf*) and across two different vocabularies (e.g. *style*). Besides, there are also general relations within one vocabulary (e.g. *broader/narrower*).

(iii) Implicit relation between two concepts without a direct link, denoted as $(c_i, \simeq, c_j)$. This relation is built based on common artworks/instances these two concepts both describe, although there are no explicit/direct links between them. For example, concepts "Rembrandt van Rijn" and "Chiaroscuro" are not directly connected but they describe 8 artworks in common out of 34 artworks that are described by either one of these two concepts. Thus we could assume that these two concepts are in a way extensionally related. Surprisingly, this implicit relation is confirmed by domain experts, since: Chiaroscuro in Italian means strong

---

[4]http://www.getty.edu/research/
[5]http://www.iconclass.nl/
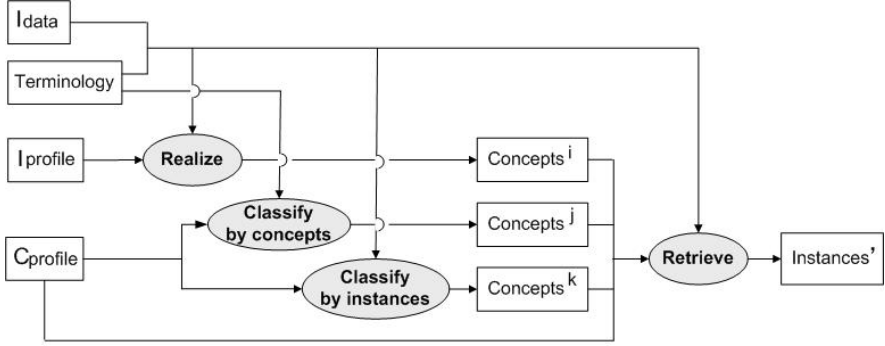
Figure 6.1: Inference steps for the task of content-based recommendation

contrast of light and dark shading. The Italian painter Caravaggio originally made
chiaroscuro his trademark and this effect is widely used in late 16th century by
many Dutch painters, such as Rembrandt van Rijn. In such a way, "Rembrandt
van Rijn" and "Chiaroscuro" are implicitly related. Another example is concepts
"Venus" and "Aphrodite", which share 4 artworks out of 6 artworks. Aphrodite
means the goddess of love and fertility in the Greek mythology and the goddness
is called "Venus" in Roman.

6.2.2   Decomposing the task into inference steps

To decompose the task of content-based recommendation, we identified four basic
inference steps (see Fig. 6.1): (i) Realization, (ii) Classification by concepts, (iii)
Classification by instances, and (iv) Retrieval. For each of them, we give a descrip-
tion, a signature (input and output datatypes), and a definition of the functionality
(relation between input and output).

   Realization is the task of finding a concept c describing the given instances i.
   • Definition: Find a concept c^i such that $O \vdash i \in c^i$
   • Signature: $i \times O \mapsto c^i$

   Classification by concepts is the task of finding a concept c^j which is directly
linked to the given concept c through a semantic relation $\sim$ in the hierarchy of
terminology T.
   • Definition: Find a related concept c^j through various semantic relations $\sim$
(e.g. *broader, narrower, teacherOf, birthPlace*, etc.) in the terminology such that
$T \vdash c \sim c^j$
   • Signature: $c \times T \mapsto c^j$

Classification by instances is the task of finding a concept c^k which shares sufficient common instances with the given concept c using the instance-based ontology matching $\simeq$.

- Definition: Find a concept c^k through the instance-based ontology matching $\simeq$ such that $O \vdash c \simeq c^k \wedge$ i $\in$ c $\wedge$ i $\in c^k$
- Signature: $c \times O \mapsto c^k$

Retrieval is the inverse of realization: determining which instance i' belong to the related concept c', where c' is a element of the unification of $C_{profile}$, C^i (Realization), c^j (Classification by concepts) and c^k (Classification by instances).

- Definition: Find an instance i' such that $i' \in c'$ where $c' \in (C_{profile} \cup C^i \cup C^j \cup C^k)$
- Signature: $c' \times O \mapsto i'$

Compared with the original definition of recommendation and its corresponding inference steps from van Harmelen and ten Teije (van Harmelen et al. 2009), we mainly extended the inference step of classification, which now consists of two components: classification by concepts and classification by instances. The original classification only determines where a given class should be placed in a subsumption hierarchy. It refers to the classification by concepts in our extended version, but we applied more semantic relations, e.g. the domain-specific relations (*teacherOf, style*) and the general relations (*broader/narrower*). In addition, we proposed a new component "classification by instances", which explores the implicit relations between concepts in the ontology.

6.3   Semantic-enhanced recommendation strategy

Following the inference steps, in this section we will explain how the system computes the prediction for related concepts and artworks based on the user's profile. As a general strategy, we apply the content-based recommendation (CBR) in CHIP, which analyzes item features/descriptions in order to identify items that are likely of interest to the user (Brusilovsky et al. 2007). Compared to other recommendation strategies (e.g. collaborative filtering), CBR performs well when there are sufficient features for items, even when there are only few user ratings (Burke 2002). Therefore it suits very well in the context of CHIP because the semantically-enriched collection could indeed provide us with rich metadata vocabularies, where artworks are connected to concepts via artworks features and concepts are linked with each other via various relations (Wang et al. 2008b).

Suppose the user likes the artwork "The Little Street", concepts "Rembrandt van Rijn" and "Venus", Fig. 6.2 shows how the CHIP system recommends related concepts and artworks based on the user profile by taking all four inference steps.

- **Realization**: Based on the artwork "The Little Street", it recommends
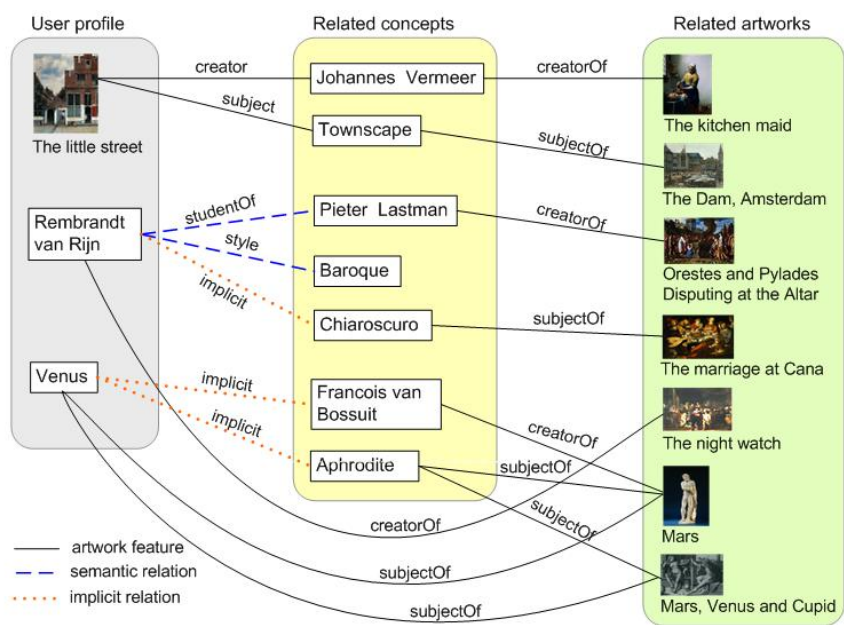
Figure 6.2: Example of semantically-enhanced recommendations

Table 6.2: Weights of explicit relations

Relation	creator	creation Site	subject	style	birth Place	death Place	teacher Of	aat Broader	tgn Broader	ic Broader
Weight	0.67	0.35	0.50	0.63	0.32	0.26	0.43	0.53	0.22	0.50
Inverse Relation	creator Of	creation SiteOf	subject Of	style Of	birth PlaceOf	death PlaceOf	student Of	aat Narrower	tgn Narrower	ic Narrower
Weight	0.68	0.31	0.54	0.61	0.28	0.21	0.44	0.55	0.16	0.52

the concept "Johannes Vermeer" via the artwork feature *creator* and the concept
"Townscape" via the artwork feature *subject*.

• **Classification by concepts**: Based on the concept "Rembrandt van Rijn",
it recommends the concept "Pieter Lastman" via the semantic relation *studentOf*
and the concept "Baroque" via the semantic relation *style*.

• **Classification by instances**: Based on the concept "Rembrandt van Rijn",
it recommends the concept "Chiaroscuro" because they share sufficient (by setting
the threshold) common artworks. Based on the concept "Venus", it recommends
concepts "Francois van Bossuit" and "Aphrodite" also because of the sufficient
common artworks they describe.

• **Retrieval**: Based on three sets of concepts: (i) rated concepts ("Rembrandt
van Rijn" and "Venus"); (ii) explicitly related concepts via artwork features and
semantic relations ("Johannes Vermeer", "Townscape", "Pieter Lastman" and
"Baroque"); and (iii) implicitly related concepts ("Chiaroscuro", "Francois van
Bossuit" and "Aphrodite"), it recommends artworks "The Kitchen Maid", "The
Dam, Amsterdam", "Orestes and Pylades Disputing at the Altar", "The Marriage
at Cana", "The Night Watch", "Mars" and "Mars, Venus and Cupid" via artwork
features *creatorOf* and *subjectOf*.

6.3.1   Computing the explicit value for the steps of realization and classification by concepts

In a previous user study (Wang et al. 2009c), we explored the use of various explicit
relations between artworks and concepts for recommendations. These relations
include: (i) artwork features between an artwork and concepts (e.g. *creator*); and
(ii) semantic relations between two concepts within one vocabulary (e.g. *broader*)
and across two different vocabularies (e.g. *style*).

Using the existing user ratings collected from this study, we investigated the
preliminary weights $W_{(r)}$ (see Table 6.2) for each explicit relation $R_{(i,j)}$, which
is either an artwork feature between an artwork i and a concept j or a seman-
tic relation between two concepts (i and j). For example, the relation between
artwork "The Little Street" and concept "Johannes Vermeer" is *creator*, denoted
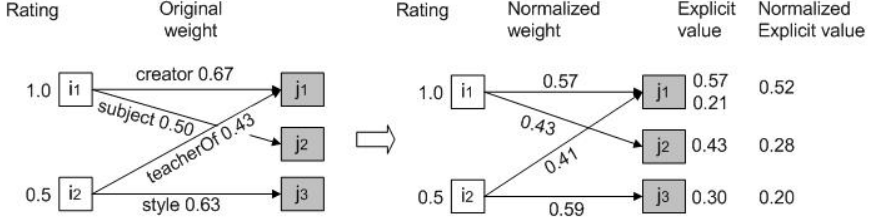
Figure 6.3: Example of calculating the normalized explicit value

as $R_{(TheLittleStreet, JohannesVermeer)} = creator$. From Table 6.2, we know that the weight of this relation $W_{(creator)}$ is 0.67. In the formulas below we write $W_{(i,j)}$ instead of $R_{(i,j)}$ and $W_{(r)}$.

Considering that a rated item (either an artwork or a concept) could be linked to multiple items via various explicit relations, we need to normalize the weight(s) for each related item. As shown in Fig. 6.3, the rated item i_1 is linked to items j_1 and j_2. The relation between i_1 and j_1 is *creator* and the corresponding weight of *creator* is denoted as $W_{(i_1,j_1)}$. From Table. 6.2, we know that $W_{(i_1,j_1)}$ (creator) is 0.67, $W_{(i_1,j_2)}$ (subject) is 0.50, $W_{(i_2,j_1)}$ (teacherOf) is 0.43, and $W_{(i_2,j_3)}$ (style) is 0.63.

$$NW_{(i,j)} = \frac{W_{(i,j)}}{\sum\limits_{j=1}^{J} W_{(i,j)}} \qquad \text{(Formula 1: Normalized weight)}$$

To normalize the weights, Formula 1 is applied. For example, based on i_1, the normalized weight of j_1: $NW_{(i_1,j_1)} = \frac{0.67}{0.67+0.50} = 0.57$ and the the normalized weight of j_2: $NW_{(i_1,j_2)} = \frac{0.50}{0.67+0.50} = 0.43$. In this way, we could calculate that based on i_2, normalized weight of j_1: $NW_{(i_2,j_1)} = \frac{0.43}{0.43+0.63} = 0.41$ and the normalized weight of j_3: $NW_{(i_2,j_3)} = \frac{0.63}{0.43+0.63} = 0.59$.

$$Exp_{(i,j)} = NW_{(i,j)} \times R_{(i)} \qquad \text{(Formula 2: Explicit value)}$$

Based on the normalized weights and user ratings, the next step is to compute the semantic value for related concepts, see Formula 2. Based on i_1, the semantic values of j_1 and j_2 are: $Exp_{(i_1,j_1)} = 0.57 * 1.0 = 0.57$, and $Exp_{(i_1,j_2)} = 0.43 * 1.0 = 0.43$. Based on i_2, $Exp_{(i_2,j_1)} = 0.41 * 0.5 = 0.21$, and $Exp_{(i_2,j_3)} = 0.59 * 0.5 = 0.30$.

$$NExp_{(j)} = \frac{\sum\limits_{i=1}^{I} Exp_{(i,j)}}{|\sum\limits_{i=1}^{I} \sum\limits_{j=1}^{J} Exp_{(i,j)}|} \qquad \text{(Formula 3: Normalized explicit value)}$$

Finally, we also need to normalize these semantic values for each related concepts, see Formula 3. $NExp_{j_1} = \frac{0.57+0.21}{|0.57+0.21+0.43+0.30|} = 0.52$; $NExp_{j_2} = \frac{0.43}{|0.57+0.21+0.43+0.30|} = 0.28$; and $NExp_{j_3} = \frac{0.30}{|0.57+0.21+0.43+0.30|} = 0.20$.

6.3.2 Computing the implicit value for the step of classification by instances

Sometimes there is no explicit relations between two concepts, however, they could be actually very similar or close to each other via some implicit relations. For example (see Fig. 6.2), "Rembrandt van Rijn" is famous for his technique using strong contrast of light and dark shading, which in Italian corresponds to "Chiaroscuro"; "Francois van Bossuit" often took "Venus" as a subject to paint; and "Venus" in Roman refers to "Aphrodite" in Greek. Compared with the "obvious recommendations" via explicit relations, these implicitly related concepts might be surprisingly new/unknown to users. The main challenge is to define how close these two concepts are in the collection.

To address this issue, Issaac et al. (Isaac et al. 2007) propose a method of instance-based ontology matching. The basic idea is that the more significant the overlap of artworks of two concepts is, the closer these two concepts are, and the level of significance is calculated by the corrected Jaccard measure, see Formula 4. In the formula, the set of instances described by a concept S is called the extension of S and abbreviate by S^i. The $JCcorr(S, T)$ measures the fraction of the refinement (by choosing the factor of 0.8) of instances described by both concepts relative to the set of instances described by either one of the concepts (Isaac et al. 2007).

$$JCcorr(S, T) = \frac{\sqrt{|S^i \bigcap T^i| \times (|S^i \bigcap T^i| - 0.8)}}{|S^i \bigcup T^i|} \qquad \text{(Formula 4: Corrected Jaccard measure)}$$

Adopting this method, we calculated the Corrected Jaccard values for all pairs of concepts in the collection. In general, the higher the Corrected Jaccard value is, the more common artworks these two concepts described. Below we give a brief look at the Corrected Jaccard values for some pairs of concepts:

0.96 (Sculptural studies – Terracotta models)
0.91 (unknown lacquerer – Lacquerware)
0.85 (Hermes – Mercury)

0.75 (Food and other objects – Still lifes with food)
0.63 (Militias – Militia paintings)
0.50 (Hinduism – Hindu deities)
0.40 (Still-life painting – Food and other objects)
0.30 (Drinking games – Sport and Games)
0.20 (Cupid – Love and Sex)
0.15 (Polychromy – Golden Legend)

0.10 (Rendering of texture – Woman)

There are in total 24249 pairs of concepts and the range of the Corrected
Jaccard value is between 0 and 1. Looking at these values and checking the cor-
responding number of artworks the pair of concepts describe in common, we set
0.20 as a preliminary threshold, which might needs more refinement in the future.
An example for the threshold 0.20 is "Cupid" and "Love and sex", which describe
8 artworks in common out of 40 artworks that are described by either one of these
two concepts. In comparison, the Corrected Jaccard value between "Rendering of
texture" and "Woman" is 0.10 and they describe 4 artworks in common out of 41
artworks.

After getting the Corrected Jaccard values for all concept pairs, we follow the
same steps (Formula 1, 2 and 3) as the calculation of the explicit semantic value
in Section 6.3.1. The only difference is that we use the Corrected Jaccard value to
replace the original weight between two concepts and then normalize the Corrected
Jaccard value in Formula 1. In the end, we will get a normalized implicit value
$NImp_{(j)}$ for each implicitly related concept j.

6.3.3 Combining the explicit and implicit values for the step of retrieval

Considering a related concept j could be linked to rated items via not only explicit
relations but also implicit relations, we need to combine values from these two
parts in order to get a final prediction $PreC_{(j)}$ for recommendation. Inspired by
the work from Mobasher et. al (Mobasher et al. 2004), we set a parameter α to
combine these two parts, see Formula 5. This combination parameter α measures
the strength of the explicit and implicit components with respect to the current
context. Taking two extreme examples: When α is 1, the system recommends items
purely based on explicit relations and this will work well if the collection is well
structured with rich semantic relations. When α is 0, it recommends items purely
based on implicit relations which is suitable for recommender systems working on
databases without semantic structures between concepts. Ideally, the parameter α
could be manually set by the user, or dynamically adapted by the system, which
enables the flexibility of the recommendation algorithm.

$$PreC_{(j)} = \alpha \times NExp_{(j)} + (1 - \alpha) \times NImp_{(j)} \qquad \alpha \in [0, 1]$$

$\big($Formula 5: Final prediction for related concepts$\big)$
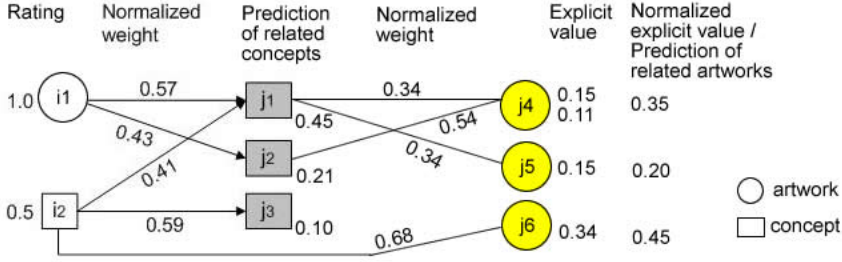
Figure 6.4: Example of calculating the prediction for related artworks

6.3.4 Computing the predication for related artworks

After getting the prediction for related concepts via both explicit and implicit relations, we calculate the prediction for related artworks. In the previous example (see Fig. 6.4), based on the user's ratings of artwork *i1* and concept *i2*, the system firstly finds related concepts *j1, j2, j3* via both explicit and implicit relations and computes the prediction for these three related concepts. Let's suppose $PreC_{(j1)}$ is 0.45, $PreC_{(j2)}$ is 0.21 and $PreC_{(j3)}$ is 0.10.

In the second step, the system finds related artworks *j4, j5, j6* via explicit relations between concepts and artworks. The relation between *j1* and *j4, j5* is *creatorOf* ($W_{(creatorOf)}$ is 0.68), the relation between *j2* and *j4* is *subjectOf* ($W_{(subjectOf)}$ is 0.54), and the relation between *i2* and *j6* is also *creatorOf*. Compared with the calculation for related concepts, we follow almost the same steps (Formula 1, 2 and 3) to compute the prediction for related artworks, except making one change in Formula 2: if an artworks is recommended based on a recommended concept (in the case of *j1, j2*), we use the prediction value $PreC_{(i)}$ of this concept instead of the rating value $R_{(i)}$. If this artwork is recommended based on a rated concept (in the case of *i2*), we still use the rating value $R_{(i)}$ of this concept. Considering that there is no implicit relations between concepts and artworks, Formula 4 is not needed. The final prediction value for related artworks is the normalized explicit value, which is compute from Formula 3.

6.4 Evaluation

In the evaluation, we use the existing user ratings collected from the previous user study (Wang et al. 2009c). There were 48 users who participated in this study. They used the CHIP Art Recommender to browse the semantically-enriched digital Rijksmuseum collection, which contains 729 artworks and 4320 art concepts. Each user rated 53 items (artworks and concepts) on average.

$$\text{MAE} = \frac{\sum\limits_{i=1}^{N} |r_i - p_i|}{n}$$

```
                          leave-one-out cross validation
for each user
   for each item i
      withhold rating r
      compute prediction p from other ratings of user
      compute  r - p, and remember
   end for
end for

add up all | r - p |
divide this sum by the number of | r - p | values
```
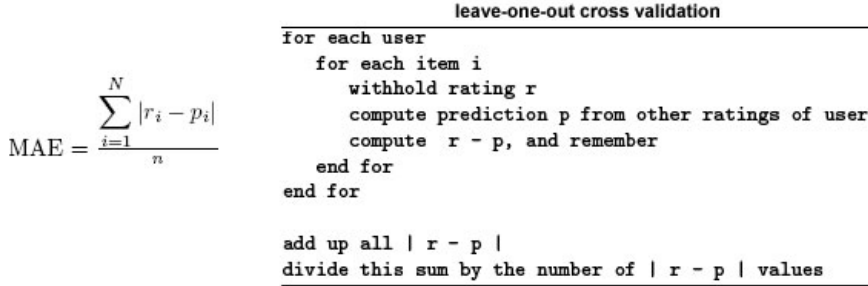
Figure 6.5: Compute the MAE by Leave-one-out cross validation

In the following sub-sections, we discuss how our approach behaviors in terms of (i) accuracy, (ii) serendipity, and (iii) explanation for recommendations, and we compare the results with the standard content-based recommendation strategy.

6.4.1   Influencing the recommendation accuracy

To measure the accuracy, we compute the standard Mean Absolute Error (MAE) by Leave-one-out cross validation (Herlocker et al. 2004). It measures the average absolute deviation between ratings and predictions, using a single observation from the original sample as the validation data, and the remaining observations as the training data. This is repeated such that each observation in the sample is used once as validation data, see Fig. 6.5. Note that ratings in the CHIP Art Recommender are based on a 5-star scale, which refers to -1, -0.5, 0, 0.5 and 1. Thus, the maximum possible value for MAE is 2 and the minimum value is 0. The lower MAE values represent the higher recommendation accuracy.

In order to see whether the semantic-enhanced content-based recommendation (SE-CBR) strategy in general improves or hamper the accuracy, we also measure the MAE for the standard content-based recommendation (CBR) strategy (Wang et al. 2008b), which was applied in the previous version of our CHIP Art Recommender. The standard CBR takes the inference steps of realization and retrieval, but no classification by concepts and instances, which means that based on user rated items, standard CBR only recommends items via artwork features.

Although there are a number of variables influencing the MAE (e.g. the parameter α, the weights for explicit relations and the threshold for the Corrected Jaccard value), in this evaluation, we only look at the impact of α on MAE in order to get a first insight and we leave the experiment with other variables to future work. Fig. 6.6 shows the impact of Alpha (α) on MAE for SE-CBR and the MAE for the standard CBR. From these preliminary results, we observe that:

(i) Compared with CBR (MAE is 0.4855), SE-CBR reaches a much lower MAE,

	Alpha	MAE
	0.0	0.3137
	0.1	0.3169
	0.2	0.3170
SE-CBR	0.3	0.3171
	0.4	0.3173
	0.5	0.3175
	0.6	0.3176
	0.7	0.3177
	0.8	0.3179
	0.9	0.3180
	1.0	0.3181
CBR	/	0.4855
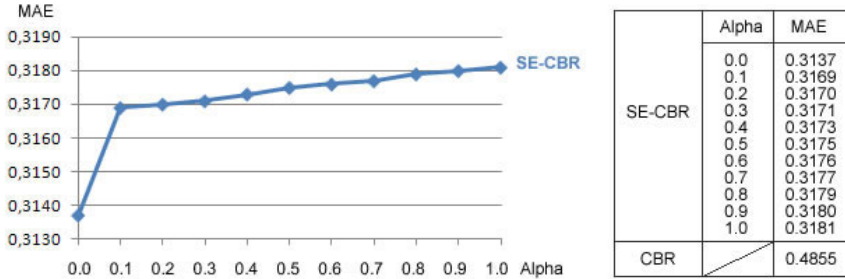
Figure 6.6: MAE for Semantic-enhanced CBR (SE-CBR) and CBR

which is in the range of 0.3137 (α is 0) and 0.3181 (α is 1). It shows that although recommending more items, SE-CBR does not sacrifice the recommendation accuracy, surprisingly, it even improves the accuracy compared with CBR.

(ii) The impact of α on MAE for SE-CBR is not significant, with a slight increase from 0.3137 (α is 0) to 0.3181 (α is 1). The reason could be that we set a very high threshold (0.20) for the Corrected Jaccard value when selecting implicitly related items. Among all 24249 pairs of concepts in the collection, only 4% (1175 pairs) has the Corrected Jaccard value above 0.20 and most of these pairs are either synonyms or very similar to each other, e.g. "Unknown lacquerer" - "Lacquerware" and "Food and other objects" - "Still lifes with food". The high similarity ensures a high accuracy for implicit recommendations. When α is 0, it only recommends implicitly related concepts which are kind of synonyms in our case and thus it reaches the lowest MAE value of 0.3137. Considering the majority (75%: 18186 concept pairs) has the Corrected Jaccard values between 0.01 and 0.10, if we set a threshold in a lower range, it will bring a lot of noisy recommendations, which might significantly decrease the recommendation accuracy. Besides the threshold for the Corrected Jaccard value, there are a number of parameters (e.g. weights for explicit relations) that influence the accuracy. We plan to try a machine learning based approach instead of the manual turning in future work.

6.4.2   Providing serendipitous recommendations

As many researchers have argued (Brusilovsky et al. 2007)(Herlocker et al. 2004), accuracy alone is not sufficient for selecting a good recommendation algorithm. A serendipitous recommendation helps a user find a surprising and new/unknown item that he/she might not have otherwise discovered. For example, if a user likes the famous Dutch painter "Rembrandt van Rijn", the standard CBR could only recommend the artwork "The Night Watch" via the artwork feature *creatorOf.*
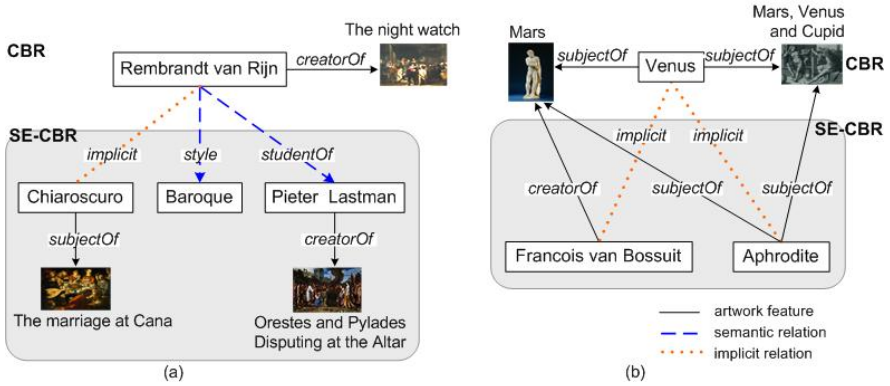
Figure 6.7: Semantic-enhanced recommendations and explanations

In comparison, the SE-CBR could recommend more items besides "The Night
Watch". As illustrated in Fig. 6.7.(a), following the semantic relations between
concepts, it finds two additional concepts "Baroque" (*style*) and "Pieter Lastman"
(*studentOf*); and based on instance ontology matching, it finds an implicitly related
concept "Chiaroscuro". Based on these concepts, it recommends more remotely-
related artworks "The Marriage at Cana" and "Orestes and Pylades Disputing at
the Altar".

Our previous study shows that compared with artwork features (e.g. *creator,
subject*), some specific semantic relations (e.g. *teacherOf, style*) offers surprisingly
interesting and new recommendations for users (Wang et al. 2009c). To follow-up
in this, it is indeed valuable to see, whether the implicitly related items are found
by users also surprisingly interesting and new.

6.4.3   Supporting more complete explanations

Besides the accuracy that affects user satisfaction, explanations of why an item
was recommended also helps users gain confidence in the system's recommenda-
tions (Herlocker et al. 2004). As shown in Fig. 6.7.(b), if a user likes "Venus", the
standard CBR recommends two artworks "Mars" and "Mars, Venus and Cupid"
via the artwork feature *subjectOf*. In the SE-CBR, it could find two implicitly re-
lated concepts "Francois van Bossuit" and "Aphrodite" based on instance ontology
matching. And these two additional concepts are also linked to the recommended
artworks: "Francois van Bossuit" is the *creatorOf* of "Mars", and "Aphrodite" is
the *subjectOf* of "Mars" and "Mars, Venus and Cupid".

In the explanation of "Why recommend", these relations between the user's
rated items and recommended items are automatically derived from the ontol-

ogy. We also found previously that explanations for "Why recommend" are useful, especially for indirectly related or serendipitous recommendations (Wang et al. 2009c). In general, for content-based recommender systems, this type of explanation proved to be preferred by most users (Herlocker et al. 2004). In such a way, the user could receive not only more recommended items, but also more complete explanations, which help them better understand the recommendations.

6.5 Discussion and future work

The main contribution of this chapter is to provide reusable inference steps and components for content-based recommender systems, which are based on semantically-enriched collections. Using classification by concepts and instances, our approach brings about three improvements: (i) retrieving more explicitly and implicitly related items without jeopardizing the recommendation accuracy; (ii) providing serendipitous recommendations, which users find new and interesting; and (iii) supporting more complete explanations for recommended items, which users consider useful.

For classification by concepts, we applied various explicit relations (artwork features and semantic relations) from the semantically-enriched museum collection in order to find explicitly-related concepts and artworks. We derived a preliminary weight for each relation from a previous study in Chapter 5 to compute an explicit value for each related concept.

For classification by instances, we adopted the method of instance-based ontology matching in order to find implicitly related concepts. Based on common instances, it builds an implicit relation between semantically-structured concepts and unstructured concepts. In this way, it bridges the vocabulary gap and provides serendipitous recommendations. We used the Corrected Jaccard value to compute an implicit value for each related concept.

To combine the explicit and implicit values for each related concept, we set a parameter α, which allows for the flexibility of the recommendation algorithm. In different domains or with different collections, the combination parameter α can be adjusted according to factors, such as how strong the semantic structure in the collection is, whether the user prefers more serendipitous recommendations than obvious recommendations, etc.

We regard our work as a first step towards a methodology for building recommender systems on the semantic web out of reusable knowledge elements. In future work, we would like to further investigate the impact of different variables (e.g. weights for different semantic relations, the threshold for the Corrected Jaccard value) on the outcome of recommendations in the evaluation.

Chapter 7

Collecting Distributed User Models for Interoperability

We introduced the museum domain ontology and the minimal user model in Chapter 2. In order to look at whether the user model is applicable to other applications with different types of data, we provide an example of reusing user interaction data (tags) in the same domain of art. In this example, user tagging about cultural events gathered by iCITY[1] is used to enrich the user model for generating content-based recommendations in the CHIP Art Recommender. To realize full tagging interoperability, we investigate the problems that arise in mapping user tags to domain ontologies. We propose additional mechanisms, such as the use of SKOS matching operators to deal with the possible mis-alignment of tags and domain-specific ontologies.

This chapter was published as "RSS-based Interoperability for User Adaptive Systems" in the Proceedings of the Adaptive Hypermedia and Adaptive Web-Based Systems (AH) Conference 2008 (Wang et al. 2008a) and was co-authored by Federica Cena, Francesca Carmagnola, Omar Cortassa, Cristina Gena, Natalia Stash and Lora Aroyo.

7.1 Introduction

The *Web 2.0* phenomenon introduced various social applications enabling online collaboration and encouraging the participation and contribution of spontaneous social networks. Users are increasingly involved in multiple Web 2.0 environments. However these applications are still "digital islands" in terms of personalized experience - not truly interconnected in a way which allows users to capitalize on the full potential of a distributed multi-application environment. Most of those services maintain a different identity, e.g. login information, preferences or profile of users with a limited integration of these data between different applications.

---

[1]http://iCITY.di.unito.it/dsa-dev/

However, tags inserted by users could be extremely useful for adaptive web applications (Brusilovsky et al. 2007), e.g. to enrich and extend the user model. The user usually tags to highlight and organize the items she is interested in, in order to retrieve them later. Thus the action of tagging can be analyzed in order to make interesting inferences on the user model (Carmagnola et al. 2007). The exploitation of tags for improving the user model, requires that systems could understand the semantics of the tags (e.g., applying suitable strategies borrowed from automatic Word Senses Disambiguation).

The focus of this chapter is to illustrate how existing fragments of user data in the form of tags can be brought together with the help of explicit semantics, and in this way allow for an adequate personalized experience across the boundaries of particular applications. This poses a considerable number of technological demands. Working in a distributed setting implies that personalization considers both *data-integration issues*, i.e. how the information from different applications is related, as well as *context-modeling issues*, i.e. in which space/time/mode the statements about a user are valid. In this chapter we look at the data-integration issue.Concretely, we provide a method for extracting, conceptualizing and linking user tags contained in public RSS files generated in the interaction of users with a social recommender system iCITY (Carmagnola et al. 2007). The tags are mapped to art-related concepts used in CHIP.

7.2   Related work

Users are increasingly involved in multiple Web 2.0 environments, such as Facebook[2], Flickr[3], YouTube[4], Del.icio.us[5] and Digg[6]. In each of them they maintain a different identity (e.g. login information, preferences). There is a *limited integration of these user profiles*, or if there exists it is not always under the control of the user, i.e. there is a lack of transparency in the use of personal data between different applications (Herlocker et al. 2004). As most of those services are relatively new and still aim at gaining critical mass of users, there is still not a methodological approach of how to assess the users' experience and improve in an *evolutionary* way the provision of the services.

Social tagging is of utmost relevance to the Cultural Heritage domain because it offers an opportunity for new relationships between cultural heritage institutions, collections and users. Social tagging may be of help: (1) to bridge the gap between the professional language defined by domain experts and the popular 'un-trained' language; (2) to encourage individuals to find personal meanings/perspectives in

---

[2]http://www.facebook.com/
[3]http://www.flickr.com/
[4]http://www.youtube.com/
[5]http://del.icio.us/
[6]http://www.digg.com/

public collections by labeling the artworks; and (3) to create public engagement with cultural heritage collections. Projects that explore this challenge, such as the Steve Museum (Trant and Wyman 2006), demonstrated the effectiveness of social tagging in engaging visitors with collections, and for the museum to understand what users consider as relevant. The Powerhouse Museum (Chan 2007), proved that user tagging and folksonomies can be used to improve navigation and discoverability through the museum collection.

Thus, user tags could be extremely useful for adaptive web applications, e.g. to enrich and extend the user model (Carmagnola et al. 2007). "Annotations can become part of his user profile as an indication of his perspective on the content collection and interest in the annotated object" (van Setten et al. 2006), thus, the systems can obtain from the tags the user has inserted, knowledge about preferences, interests, etc. Adaptive systems may use this "tag-enriched" profile for recommendations. Notice that tagging, and more generally annotating, can be considered as possible actions a user can perform on a social web site. As other kinds of *usage data* (Kobsa et al. 2001) (clicking, buying, etc.) these actions represent an important feedback from the user. In fact user usually tags in order to highlight and organize the items she is interested in, in order to retrieve them later. Thus the action of tagging is a stronger indicator (Kobsa et al. 2001) for user interests than simply clicking on a link, and therefore should be analyzed in order to make interesting inferences on the user model.

To be able to exploit tags for improving the UM, systems are required to understand the semantics of the tags. Suitable strategies for automatic Word Senses Disambiguation (WSD) are applied. This involves matching the context of a word instance with either information from an external knowledge source (knowledge-driven WSD), or information about the contexts of previously disambiguated instances derived from corpora (data-driven or corpus-based WSD) (Ide and Veronis 1998).

7.3 Usage scenario

In this chapter we show how two user-adaptive applications can realize a meaningful exchange of user data and in this way compensate for either lack of internal semantics or lack of user data.

In the scenario, Carlo uses *del.icio.us* to collect and share bookmarks, and *flickr* to store, share and retrieve pictures. He also uses iCITY to stay up-to-date about events in Turin. iCITY allows Carlo to store and tag events and retrieve them via his GPRS equipped mobile phone. Carlo has a strong interest in art. He is planning a weekend in Amsterdam, and he would like to visit the Rijksmuseum. He uses the *CHIP Tour Wizard* to prepare his visit. Information (maintained by iCITY) about the events he has seen, the tags he inserted and the topics he is interested in could be very useful for the CHIP system to quickly identify his focus

of interest and offer him a personalized visit to the museum. While registering in CHIP Carlo aligns his account with already existing accounts, e.g. iCITY, flickr, del.icio.us.

This is a typical example of re-use of user interaction data generated by one application into another in similar domains. In this way, we illustrate how can two user modeling problems be solved, i.e. (1) cold-start problem in CHIP, that can initialize the user model and start the recommendation from a point closer to user's interests, and (2) maintaining an integrated user profile, which reflects larger scope of user interests and activities.

7.4   iCITY-CHIP user interoperability architecture

The section illustrates the main characteristics of iCITY and CHIP user-adaptive systems and the interoperability aspects of their architectures (Fig. 7.1).
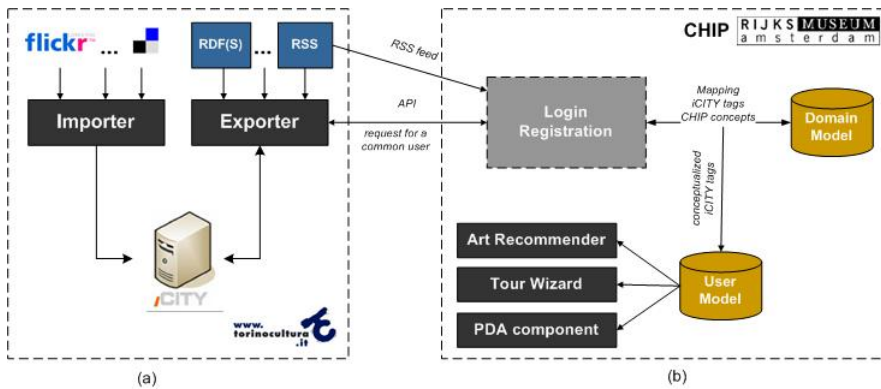
Figure 7.1: iCITY-CHIP Interoperability Architecture

7.4.1   iCITY tagging and recommender system

iCITY[7] is a social web-based, multi-device recommender system. It provides suggestions on cultural events in the city of Turin, and allows users to insert new events, to add information about events, to insert comments and tags. Recommendations are based on the user model enriched with tags, exploited to infer user features (see (Carmagnola et al. 2007) and (Carmagnola et al. 2008)). The

[7]Digital Semantic Assistant iCITY. Dept. of Computer Science of Turin, City of Turin and CSP - ICT innovation are project partners; http://iCITY.di.unito.it/dsa-dev/

iCITY user model is an overlay of the iCITY Event ontology, created as an RDFS transformation of the event classification in *TorinoCultura*[8], a web portal managed by the municipality of Torino for informing citizens about cultural ongoing events in the city. This ontology contains links both to WordNet[9] synsets and domains. iCITY has a modular architecture for extracting, maintaining, reasoning and exporting of user tags (Fig. 7.1), which can be shared with other applications via a RSS feed. The main components for interoperability (Fig. 7.1) are the Importer and Exporter Modules responsible respectively for the extracting the tags from external sources and making user tag available to other applications.

iCITY Importer Module. In the iCITY registration form, a part from username and password, the user can provide the tags that best describe her and her social web community. If the user provides web community accounts (e.g. flickr, del.icio.us), the Importer Module retrieves the available RSS files containing the set of tags provided by the user in those web communities (Carmagnola et al. 2007). The Importer Module is able to extracts these tags regardless of the format used to represent them. Once all the user tags have been extracted, they could give interesting information about user interest and knowledge. In order to understand their meaning, the system looks for correspondences between the tags and the *synsets* and the *domains* of the MultiWordet database[10].

If one or more correspondences are found, they are linked to the Event Ontology class/subclass[11]. Thus, the system can reason on tags and increases the level of inferred user's interests related to the class the tag belongs to. However, this approach is limited and suffers from several problems. The tags (both imported from other systems and inserted by users) could not be directly mapped on the concepts in iCITY ontology. For instance, if a user tags the event "Picasso exposition" with the term "Picasso", iCITY can infer a medium/high knowledge in Art, since the word belongs to the domain Painting linked to the class Art of the iCITY ontology. However, if a user tags the same event with the word "Spain", iCITY is not able to find a direct correspondence with a class/subclass of the ontology and the tag is not semantically analyzed. In this way iCITY loses some important information, since the tag could demonstrate that the user knows the painter and thus she has an medium/high knowledge in modern art, and her user model could be consequently updated.

iCITY Exporter Module. The Exporter Module generates, for every user, a RSS file with the list of the events tagged by the user. For every event, the file

---

[8]http://www.torinocultura.it/

[9]http://wordnet.princeton.edu/

[10]In MultiWordet (http://multiwordnet.itc.it/), and Wordnet, each synset is annotated with at least one domain label, selected from a set of about two hundred labels which constitute the so-called WordNet Domains.

[11]The classes and subclasses of the iCITY Event ontology are mapped on the corresponding synset and domains as a semantic enrichment step. This relations are in "one to more" cardinality, since a class/subclass of iCITY may correspond to one or more synset/domain of MultiWordet

stores: the title, the URL , the description, the reference to the event category and subcategory in the iCITY event ontology, the reference to the Wordnet synsets and domains linked to the subcategory, and finally the list of the tags associated by the user to the event. In the following, we report an example of the RSS file for a user.

```
...
<!-- TITLE OF EVENT -->
<title>What a click!</title>
<!-- LINK TO EVENT -->
<link>http://www.icity.di.unito.it/dsa-dev/events/1408</link>
<!-- DESCRIPTION TO THE EVENT -->
<description>...</description>
<!-- REFERENCE OF iCITY TAXONOMY -->
<category> Art--Interdisciplinary </category>
<!-- LINK TO WORDNET DOMAIN -->
<wn20schema:domain>Art</wn20schema:domain>
<!-- LINK TO WORDNET SYNSETID -->
<wn20schema:synsetId>00933420</wn20schema:synsetId>
<wn20schema:synsetId>n#00598038</wn20schema:synsetId>
<!-- USER -->
<dc:creator>carlo</dc:creator>
<dc:date>2008-01-19</dc:date>
<!-- TAGS -->
<dc:subject> violin, painting, sculpture, portrait </dc:subject>
...
```

As emerges from the above code, the list of tags ($< dc : subject >$) is expressed as a simple list of meaningless strings.

The only way to provide semantics to them for the receiver is exploiting the reference to the category and subcategory the event belongs to in the iCITY event ontology, and the reference to the WordNet domains and synsets[12]. In this way, a recipient system can import the RSS file containing the tags used by a particular user, and it can try to disambiguate the meaning of the tagged event thanks to the information, provided in the RSS file, about the event subclass they belong to and the references to WordNet domain and synset.

In the example, after tag disambiguation, the receiver could infer from the tags a user interest in Art, and in partilar painting, sculpture, and portrait (Carmagnola et al. 2007).

7.4.2 CHIP Art Recommender and Tour Wizard

CHIP system (Cultural Heritage Information Personalization) allows museum visitors to create their personalized experience in the Rijksmuseum both with the virtual collection on the museum Web site, and in the physical museum by quickly finding the right path covering her interests. To realize this, the CHIP demonstrator[13] provides a web-based virtual Tour Wizard and a Mobile Tour Guide used in

---

[12]We provide two different synset ID for every sysnet: the former is referred to the database location id shown in the online version of WordNet (http://wordnet.princeton.edu/perl/webwn); the latter refers to the ID given to the synsets in MultiWordet

[13]http://www.chip-project.org/demo

the physical museum (Wang et al. 2008c). CHIP takes an open Web and 'non-intrusive' user modeling approach for providing these personalized services on both in the virtual museum Web site and in the physical museum environment. As depicted in Fig. 7.2 the user data is stored into four user profile categories: *personal*, *social*, *ratings*, and *iteraction*.
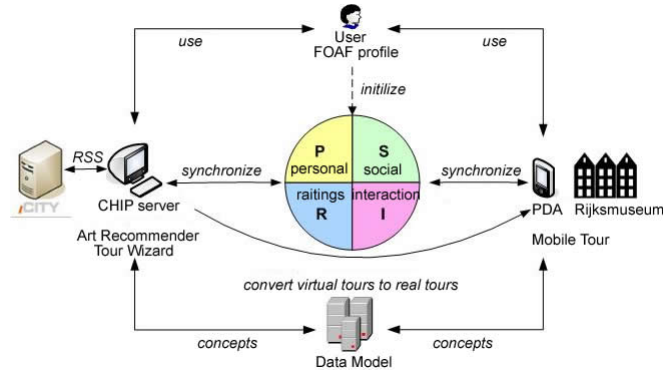
Figure 7.2: CHIP User Model Properties and link to iCITY exporter

The *personal category* stores the user stable characteristics, which in a typical open Web context could be initialized by either importing an existing FOAF RDF profile or via an OpenID channel linking the CHIP login data to an existing login information of third party Web application. Alternatively, the user can manually fill in the CHIP User Questionnaire to provide his or her personal data. The CHIP personal scheme contains two sets of concepts: (1) FOAF classes and properties, e.g. agent, person, givenname, surname, age, gender, homepage, mbox, img; and (2) CHIP concept properties, e.g. frequency, knowledge. The *social category* describes the user's social information also initialized by FOAF properties,e.g knows, openid, organization, group, member, OnlineAccount, accountName, accountServiceHomepage; and/or the CHIP User Questionnaire. The *interaction category* stores the user's interactions on the Web and in the physical museum (using the mobile guide), e.g. virtualTours (artworks created in the Tour Wizard on the Web), realTours (artworks visited in the real museum), time (time spent in the real museum tour). The core of the CHIP user profile is stored in the *ratings category*, which maintains the user's explicit ratings of artworks and topics in terms of VRA Core properties, e.g. work, creator, title, shortTitle, creationDate, creationSite, subject, materialSupport, materialMedium.

The CHIP Art Recommender provides in an interactive way to the user to: (1) express her art preferences; (2) to find quickly the artworks of interest and (3) in

the same time to build a user profile of these interests and preferences. This user profile is further used in the Tour Wizard to assist the user in (semi-)automatically creating museum tours, and in the Mobile Tour Wizard, where it is updated with the artworks seen and rated during the museum visit. Typical for the CHIP system is that recommendations are generated not only for artworks but also for various art-related topics.

7.5   iCITY-CHIP user tag interoperability

In this interoperability use case we use an open API to request and link the user data . Once the user personal (login) information is aligned between CHIP and iCITY (Fig. 7.1(b)), based on the RSS feed we maintain a dynamic mapping of iCITY user tags to the CHIP vocabulaty set (Rijksmuseum specific concepts, or shared domain vocabularies, such as Getty AAT, ULAN, TGN, and IconClass, or general purpose lexical data such as WordNet).

The main challenge in achieving the interoperability of user tags between the two systems iCITY and CHIP is to provide a dynamic mapping mechanism, which allows for a constant stream of user tags from iCITY to be interpreted (mapped) to concepts from the internal vocabularies of CHIP. This will allow to use iCITY tags to populate the user profiles of new users in CHIP and to be able to instantly generate a tour of recommended artworks in the Rijksmuseum.

To implement the tag mapping from iCity to CHIP, we use the Simple knowledge organization System (SKOS) Mapping Vocabulary Specification[14] created for linking thesauri to each other with relationships skos:exactMatch, skos:broadMatch, skos:narrowMatch and skos:relatedMatch. For this first stage alignment, the mappings are still based on the lexical match of tags. With a few additional simple restrictions by applying the type of tags, a lexical match gives more confidence to generate a semantic match as strong as owl:sameAs[15](Tordai et al. 2007). As an example, a mapping to *Amsterdam*, known to be the geographical name of a city in The Netherlands, can be made with owl:sameAs.

In Fig. 7.3 we illustrate how to realize an instance-based mapping of the iCITY tags *Amsterdam*, *Giovanni*, *photography* and *1400* from a particular user to art concepts in CHIP vocabularies[16].

The mapping is realized in two steps: (i) to identify the type (e.g. creator, place, material, etc.) of the tag as a simple restriction; (ii) to map the tags by using SKOS mapping relations. As shown in Fig. 7.3, the iCITY tag *"Giovanni"* results in two partial matches in the Getty Unified List of Artist Names (ULAN) with the type of `creator`: *"Piranesi, Giovanni Battista"* and *"Tiepolo, Giovanni Battista"*. The level of ambiguity with names as user tags could be quite high, especially if

---

[14]http://www.w3.org/2004/02/skos/mapping/spec/
[15]http://www.w3.org/TR/owl-features/
[16]Notice that tags are required to be provided in English.
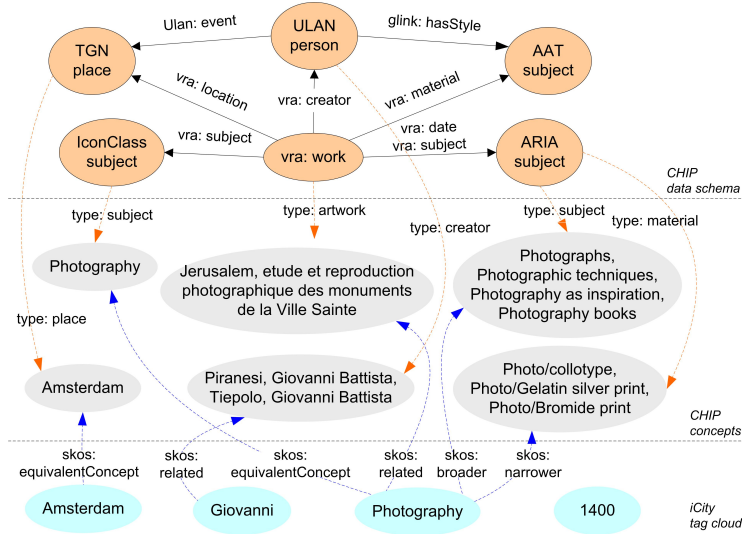
Figure 7.3: User Tag Mapping Scheme

the domain in not only limited to art. To confirm whether it is either of the two Getty ULAN artists or none of these two, we need further evidence from the tag cloud, e.g. the event (annotated with this tag) and related tags (used together to annotate this event). From the user's iCity RSS file, we know that *"Giovanni"* is used to annotate the event *"Why Africa?"* together with other tags *"Africa, exhibition, art, contemporary, Torino"* and actually the user means *"Giovanni Agnelli"*. Thus in this case, although we have good partial match, we use the semantic weak "`skos:relatedMatch`" mapping relationship with a low certainty for the type `creator`. Another example, the semantic equivalence between the iCITY tag *"Amsterdam"* and the Getty Thesaurus of Geographic Names (TGN) creationSite *"Amsterdam"* is expressed with skos:exactMatch with the type of `place`.

Compared with *"Giovanni"* (2 weak related matches), *"Amsterdam"* (1 exact match), and *"1400"* (no matches at all), the mapping of *"photography"* is more complex, which results in six partial matches in the CHIP vocabularies with four different SKOS mapping relations and three different tag types (see Fig. 7.3): (i) the `skosm:narrowMatch` for the type `material` points to hierarchical specialization in the Rijksmuseum ARIA vocabulary; (ii) the `skosm:broadMatch` for the type `subject` indicates hierarchical generalization in the Rijksmuseum ARIA Eny-

clopedia term list; (iii) the `skosm:exactmatch` for the type `subject` refers to the
IconClass concept vocabulary, which describes a semantic equivalence; and (iv)
again, a semantically weak relationship `skos:related` for the type `subject` is ap-
plied for the specialization in the Rijksmuseum ARIA collection.

The examples above give a good illustration of the semantic and syntactic
mappings we can provide between the iCITY tag cloud and the CHIP concepts.
Maintaining a certainty level for each mapping allows for tuning of those concept's
relevance. Further evidence from the tag cloud and/or the user model allows for a
good accuracy as well as the user's direct feedback/confrmation. However, some
problems still remain. Below we discuss these issues with the possible solutions.

- *Tags are messy.* The mapping is realized in two steps: first, to identify
  the type of the tag, e.g. whether *Giovanni* is a person, a kind of material
  or a place; and then to the map tags using SKOS mapping relations. The
  disambiguation of mapping can be delegated to the user for a varification, or
  to collect further evidences from the tag cloud.

- *Grammatical variation.* Often tags appear in various grammatical forms,
  which do not completely match the CHIP concept form, e.g. noun, verb,
  adjective and adverb formas. Maintaining additional relationships or dis-
  tances between the different term forms allows for clustering of all possible
  mappings for a given tag, e.g. sculpture - sculptural - sculptor, theater -
  theatrical. Using mapping to WordNet can facilitate this process efficiently.

- *Combined effect of tags for recommendations.* After mapping each tag to
  the CHIP vocabularies, we have to think about the combined effect of tags
  for generating recommendations in CHIP. Our idea is to treat the tags dif-
  ferently depending on their relations with the annotated events, which are
  described in the user's RSS file from iCity. For example, if "Giovanni" and
  "photography" are annotated with the same event, the CHIP system will
  search for and recommend artworks, which include both concepts with the
  higher priority; otherwise, the two tags can be used separately for generating
  different recommendations.

- *Ranks of recommendations.* To rank the recommendations based on different
  tags or tag groups, we are considering maintaining a dynamic weight for each
  tag, which could be defined by factors like *frequency* of use of this tag in the
  user profile, *uniqueness* of use of this tag in the whole system, and by all
  users, etc.

7.6   Discussion and future work

In this chapter we have presented an approach to exploit widely used tag anno-
tations to address two important issues in user-adaptive systems in the cultural

heritage domain: the cold-start problem and the integration of distributed user profiles. We have sketched a scenario, in which user tagging about cultural events gathered by iCITY is used to enrich the user profile for generating personalized recommendations of artworks in CHIP. To realize full tagging interoperability, we have investigated the problems that arise in mapping user tags to various ontologies, and we proposed additional mechanisms, such as the use of SKOS matching operators, to deal with the possible mis-alignment of tags and domain-specific ontologies. Issues that need to be addressed in future research are the loss of information that occurs when relating tags to event ontologies (iCITY) and the effective mapping between single or possibly multiple tags to the domain-specific ontologies as used in cultural heritage.

Last but not least, the CHIP user profile has to be exported back into the iCITY recommender in an interoperable format, in a way that the iCITY Importer module is able to import this information. The inferences made by the CHIP recommender could be useful to refine the iCITY user model and could also help the iCITY reasoning component to solve some of the disambiguation problem described in 7.4.1. On the other hand, the CHIP mapping component could be refined to utilize also the sysnets and the domains information exported in the iCITY RSS file in solving its disambiguation problems.

Chapter 8

Conclusions and Discussion

In this chapter we conclude the thesis by reflecting on the research questions posed in Chapter 1. For each of these questions we recall the results, draw conclusions, and discuss related issues. We finish with recommendations for future work.

8.1 Revisiting the research questions

RQ 1. *Can we acquire user information in a non-intrusive way?*

In order to provide personalized services, the traditional approach for recommender systems is to ask users provide personal information as much as possible beforehand. To minimize the intrusiveness, we built an interactive rating dialog to allow users (in particular, first-time users) immediately profit by getting recommendations in a way suiting their art preferences (Page 14).

This rating dialog is realized over representative samples of artworks from the museum collection. Users can give their feedback of likeness to both artworks and related concepts on a rating scale of 1 to 5 degrees, and meanwhile browse and search artworks in the collection. In this way, the system collects users' art preferences, which enables a quick instantiation of the user model for inferences and recommendations. Based on their ratings, the system predicts users' interest in other related artworks and concepts by deploying the semantic structure of the domain ontology. It also allows users give ratings to these recommended artworks and concepts, which are collected as feedback to refine the user model.

Because of the rich semantic structure in the domain ontology, even with a limited amount of user ratings, the system still may predict users' interest in related artworks and concepts for recommendations, which partially solves the cold-start and sparsity problems (Page 11). We deployed not only standard artwork features (e.g. *creator, creationSite* and *materialMedium*) between artworks and concepts, but also various semantic relations within one vocabulary (*teacherOf,*

borader/narrower) and across two different vocabularies (e.g. *styleOf, deathPlace*) (Page 57). In this way, it propagates user ratings in a wider range of explicitly and implicitly related artworks and concepts.

Although there is not enough evidences with certainty to prove the non-intrusiveness of our approach, users from the evaluations seem to have no problems with the acquisition of their art preferences using the rating dialog and were satisfied with the recommendations. From the first user study (Page 22) the results indicate that it helps users, especially novices, to elicit or clarify their art preferences from their implicit knowledge about the museum collection. In the second user study (Page 23) we compared different sequences of artworks (e.g. random, expert-sorted and self-selected) and different approaches for rating (e.g. rate only artworks, rate only concepts, and rate both artworks and concepts). The findings enable the system collected users' ratings in a range between 96 and 224 in a period of 5 minutes, and derive most representative samples for ratings, which contains 20 artworks and 45 well-distributed art concepts from the collection.

RQ 2. *What is a minimal user model to store user information in a recommender system?*

We have found in CHIP that a minimal user model needs to contain four parts: (i) user's personal information (e.g name, age and email); (ii) objects that the user has interacted with; (iii) user's activities over the objects (e.g. the user *rates* an object with a value and the user *views* an object during a tour); and (iv) corresponding contextual information (e.g. time, place and device).

To construct such a minimal user model ontology, we follow three design principles: i) light weight with both minimal structure and semantics; ii) share and reuse the data to other applications; iii) compatibility. Except for the application-dependent classes and properties, there is no intention to define redundant classes and properties that already exist in other standard RDF vocabularies. Following these design principles, we look at existing ontologies/vocabularies such as FOAF, SKOS, Dublin Core, SIOC and SEM, as a starting point to construct our CHIP user model. On one hand, these existing ontologies largely cover general descriptions of, as well as relations between persons, objects, and communities/organization, and, on the other hand, they are being adopted by a steadily increasing user community (Bojars et al. 2008).

Firstly, we built a user model as a specialization of FOAF to store users' personal data (Page 15). Main classes and properties from FOAF that we used in CHIP are *foaf:Person* and *foaf:holdsAccount*. The *foaf:Person* class is used to represent the information about a person who holds an account *chip:User* on a Web site. Account specific information is described by *chip:User* as a subclass of *foaf:OnlineAccount*. The *foaf:holdsAccount* property is used to link a *foaf:Person* to a *chip:User*. To store users' personal information which are collected from the user

questionnaire , we use existing FOAF properties (e.g. *foaf:firstName, foaf:age* and *foaf:gender*) as well as our self-defined properties in CHIP, such as *chip:profession, chip:rijksmuseumVisitFrequency* and *chip:museumVisitReason.*

Secondly, we defined a core class *chip:Rating* (called *chip:RatedRelation* in the first version) to describe the user's activities over the objects and these objects (Page 16). By using the definition of semantic N-ary relations, the *chip:Rating* class contains information in three arguments: who has rated (property: *chip:hasRated*), what is rated (property: *chip:rated*), and what value the rating gives (property: *chip:ratedValue*).

Thirdly, we mapped our user model to another existing event model SEM to store additional user activities during the museum tour (Page 49). We defined *chip:User* as a subclass of the *sem:Actor*, who participates in the *sem:Event*. In our case, there are three different types of events: rating, viewing, and taking a tour. In a rating event, the user rates a *sem:Object* with a *chip:ratedValue* in 5 degrees. The viewing events are usually part of the tour events, since the user views a *sem:Object* during the tour. In a tour event, the user adds a *sem:Object* into a particular tour. All of the objects added in the tour will be ordered in a sequence based on their locations in the museum, which are described using the *rdf:_n*[1] as a sub property of the *rdfs:member*.

RQ 3. *Can we use the semantic structure of collections to improve recommendation algorithms?*

For this question, our hypothesis is that by choosing specific semantic relations, it could help the recommender system retrieve more related items without decreasing the prediction accuracy and interestingness (Page 56).

To test this hypothesis, we take three steps and perform evaluations respectively. Firstly, we develop a content-based recommendation algorithm in the CHIP Art Recommender (Page 13). Based on the domain ontology, it recommends related artworks and concepts via basic artwork features, such as "Night watch" *creator* "Rembrandt van Rijn"; *creationSite* "Amsterdam" and *subject* "Militia". In addition, it provides users with explanations for recommended artworks and concepts by automatically extracting the features used in the inference step. In the evaluation (Page 22), we test the effectiveness of recommendations with real museum visitors. The results indicate that by providing recommendations with explanations, it significantly helps novices elicit their art interests, while there is a slight increase for experts.

Secondly, based on the domain ontology, we identify different types of semantic relations within one vocabulary (e.g. *broader/narrower*) and across multiple vocabularies (e.g. *hasStyle*) (Page 57). Besides the basic artwork features, these various semantic relations help retrieving more explicitly related items (artworks

---

[1]http://www.w3.org/TR/rdf-schema/

and concepts). However, not all related items are useful or interesting for users. In the evaluation, we test the Art Recommender with end users by applying both artwork features and semantic relations to recommend related concepts (Page 59). Using artwork features as a baseline, we compared recommendations via different relations in terms of accuracy and interestingness. The results demonstrate that by choosing specific semantic relations (e.g. *creator, hasStyle, teacher/studentOf*), the recommender system could retrieve more related items without jeopardizing the accuracy and interestingness. In comparison, semantic relations considering geographic locations (e.g. *tgn:broader/narrower, ulan:birth/ deathPlace*) score very low on both accuracy and interestingness.

Thirdly, we adopt an existing method of instance-based ontology matching to build implicit relations between concepts, and propose a hybrid approach combining both explicit and implicit relations for recommendations (Chapter 6, Page 74). On top of this, we define the task of personalized recommendation for semantically-enhanced recommender systems, and decompose this task into four inference steps: *realize, classify by concepts, classify by instances,* and *retrieve.* We evaluate the hybrid recommendation algorithm in terms of recommendation accuracy. Compared with the original content-based recommendation algorithm used in the first step, the hybrid algorithm not only significantly increases the recommendation accuracy, but also finds more related items for recommendations via both explicit and implicit relations. In addition, we show that the hybrid algorithm provides serendipitous recommendations, which users find new and interesting; and supports more complete explanations for recommended items, which users consider useful (Page 80).

RQ 4. *How can we present semantically-enhanced recommendations?*

All the first three research questions lead to the last question of presenting recommendation results for end users, in order to enhance their museum experiences in a more intensive, long-lasting and engaging way, by linking the museum experiences both online and on-site. Towards this goal, we develop three tools within the CHIP demonstrator in a coherent way, namely: *Art Recommender, Tour Wizard* and *Mobile Guide.*

- The *Art Recommender* helps users to discover their art interests in the museum collection and to store their art preferences in a corresponding user model (Page 13). To facilitate navigation and browsing, we adopt existing techniques like Spectacle[2] and Simile[3] in the Art Recommender in order to cluster multiple recommendations based on relations. In addition, the system automatically derives the relations which are applied to retrieve explicitly or

---

[2] http://www.aduna-software.com/products/spectacle/
[3] http://simile.mit.edu/

implicitly related concepts and artworks in order to explain the underlying recommendation inference to users.

- The *Tour Wizard* generates online museum tours containing interesting artworks recommended by the first tool, *Art Recommender*. We present artworks in the museum tours with different views such as a historical time-line and the museum map (Page 35).

- The *Mobile Guide* converts online museum tours (generated from the *Tour Wizard*) to on-site tours on handhelds (e.g. PDA, iPod), and assists the user in finding his or her way during the visit (Page 36). When the tour is finished, it sends the user's real behaviors to update the user model on the Web server. Besides, we also implement a real-time routing system in the Mobile Guide, using the coordinates of artworks and rooms in the museum (Page 48).

We evaluate the performance of the Art Recommender in terms of recommendation effectiveness (Page 22). Due to the constraints from the museum side (e.g. permission to use the real museum environment, the attachment of RFID tags to artworks in the current exhibition, and the availability of mobile devices and related hardware), it is difficult to perform empirical evaluations on museum tours with handhelds. Therefore, we augment the evaluation with a qualitative analysis of personalized museum tours provided by the Tour Wizard and the Mobile Guide (Page 37). Apart from the qualitative analysis, we measure the speed of the router in the Mobile Guide (Page 51). The results show that the sequence of recommended artworks follows an efficient route through the museum in a reasonable time that allows real-time interaction with the system.

8.2 Reflection and discussion

Looking back at what we have done in the CHIP project, we identify four limitations in our research, concerning respectively user modeling, recommendation algorithms, applying our method for other applications and increasing users' motivation. We describe each limitation in turn and discuss possible solutions and future work.

Towards a general user model ontology. In the last decades, a lot of work has been done on general user models (Kobsa 2001), which allows for domain independence and compatibility with different applications. The notion of ontology-based user models was first developed in OntobUM, which integrated user ontology, domain ontology and log ontology (Razmerita et al. 2003). GUMO[4] is

---

[4]http://www.gumo.org

another example, which allows uniform interpretation of distributed user models in intelligent environments (Heckmann et al. 2005). A more recent example is SIOC[5], which provides an ontology for representing rich data from the Social Web in RDF, and is commonly used in conjunction with the FOAF vocabulary for expressing personal profile and social networking information (Bojars et al. 2008).

Inspired by SIOC, we designed a minimal user model ontology as a specialization of FOAF to store user personal information and user ratings, and we manually mapped the user model schema to an external event model SEM to store user activities in the museum tours. As a first step towards user model data interoperability, we exchanged user ratings with external user tags of the iCITY project for recommending artworks based on these tags for individual users. However, the mapping between user tags to the domain ontology is done manually. In future work, we would like to continue our research on mapping algorithms between different user models, and to study related issues about user-identification, privacy and contextual information. As an inspiration, Berkovsky et al. proposed a framework that allows transformation of user modeling data between recommender systems (Berkovsky et al. 2007), such as dealing with the mediation between a trip planning system Trip@dvice(Ricci et al. 2002) and a personalized museum visitors guide PIL (Kuflik et al. 2007).

Explore hybrid recommendation algorithms. We started with the content-based recommendation algorithm in the CHIP Art Recommender because it work best when the collection contains a rich semantic structure, which can be exploited for inference and recommendations (Herlocker et al. 2004). In future work, we would like to study hybridizations of combining different recommendation algorithms such as content-based filtering (Morita and Shinoda 1994), collaborative filtering (Resnick et al. 1994) and knowledge-based recommendation (Burke 2002). As a first step towards this goal, we collaborated with the Kubadji[6] project to explore a hybrid algorithm combining both content-based and collaborative filtering algorithms. The main idea is to propagate the user's ratings to related but unrated objects in the user profile, and then based on the extension of overlaps between users' profiles, to compute the similarities between users for recommendations. Although the result is unmature and preliminary, it gives interesting insights for future work.

Apply our method to other domains. Because our approach is applied only on the Rijksmuseum collection in the cultural heritage domain, we can not make claims about how our method will be applied on other domains. However, we expect that some knowledge elements can be reused for other recommender systems within the same or different domains. For example, we defined in Chap-

---

[5]http://rdfs.org/sioc/spec/
[6]http://hum.csse.unimelb.edu.au/kubadji/

ter 6 the task of personalized recommendation and decomposed the task into four inference steps (realize, classify by concepts, classify by instances, retrieve). By setting the combination parameter α in the inference process, it enables both explicit and implicit recommendations based on collections between two extremes, a domain ontology with a semantic structure between objects and concepts, and a traditional database with a flat/weak structure. The results of our work produce novel experiences towards a methodology for building semantically-enhanced recommender systems.

Reusability is a key issue for the future work. Can we use our task model to describe other semantically-enhanced recommender systems, or will we end up inventing new task types and inference steps for every new recommender system? As an inspiration to our work, van Harmelen et al. analyzed most entries to the Semantic Web Challenge[7] events of the years 2005, 2006 and 2007 in order to test whether they could be properly described with the small set of primitive task types they defined (van Harmelen et al. 2009). Following their work, we would like to validate our task model, in particular, starting with ontology-based recommender systems, e.g. the collaborative filtering framework (Mobasher et al. 2004), Foafing the music (Celma 2006), and the hybrid recommendation model (Cantador et al. 2008). Although these recommender systems apply different recommendation algorithms, they all work with collections mapped to standard domain ontologies, which allows us to apply our task model also on these systems. We expect that the details of our definitions of the task type and inference steps may well have to be adjusted over time.

Increase users' motivation. Like Kelly et al. argued in the Museums and the Web conference, the various meanings of openness to museums pose opportunities as well as problems (Kelly et al. 2008): Open standards could enable maximum access, device independence, and interoperability, but they are often impeded by different interpretations of what they mean and by lack of mechanisms for enforcement. Open content could be essential for sharing cultural heritage knowledge, but who owns what content to share is not a simple matter. Open services could enable greater use, but it is not museum services that museums most want to ride along with, but rather services provided by commercial interests which, when opened, carry museums along on the wave (Bearman and Trant 2008).

In order to increase users' motivation to actively engage with museum services and generate their own content, museums did try to cooperate with social network sites, such as Facebook, YouTube, or Flickr. For example, until November 2007 Facebook did not allow institutional profiles, but the Brooklyn Museum[8] was able to link its ArtShare application to the Facebook API, giving it a degree of control and yet leveraging the social site. Museum curators want to give

---

[7] http://challenge.semanticweb.org/
[8] http://www.brooklynmuseum.org/

(some) users greater voice but do they want to give all users greater voice and what will be the impact of that? To study this question, Alexander et al. analyze the reports of five very different museums on what they found when they put museum videos on YouTube, and they find that overall, there was very little user-contributed video feedback, and indeed less user comment than anticipated, and YouTube exposure did not help drive traffic to the museum Web sites or to the physical museum (Alexander et al. 2008). Although these experiments with Web 2.0 social sites have had less impact, than anticipated, the results are helpful for the museum community. A more sophisticated approach might be needed for museums to actually motivate users to contribute to the museum Web sites and visit the museums.

8.3   Looking ahead

Let's suppose a user scenario: Robin likes art and she is a fan of the film director Peter Greenaway. She has a conference next week in Leiden (The Netherlands) but she wants to go sightseeing in Amsterdam. In order to plan such a trip, traditionally, you have to print out all the data, sort through it and then stand back and see if you can make the connections yourself, which can be quite stressful and time-consuming.

In Web 3.0, linked data sets are being released and in the activities of big companies like Microsoft[9], which acquired Powerset[10] in 2008, and Google[11], which acquired Metaweb[12] in 2010. In this environment, with the help of semantic websites, the realization for Robin's scenario becomes much easier. We could pull all these information forms together instantly, put them on the same map, let the semantic website(s) bridge the connections automatically and provide possible solutions or recommendations. For instance, using Tripit[13] or Kango[14], it can help Robin organize her travel plans such as airline tickets and hotel bookings in a way that's easy to share and access. Her interests of art and Peter Greenway could be extracted from her FOAF profile or Facebook[15], and then this information could be used in, e.g. the CHIP Art Recommender, for recommending artworks in the Rijksmuseum Amsterdam which are available for the exhibition during the period that she will stay in the Netherlands. In addition, based on her interest, movielens[16] could also recommends a film "Nightwatch" directed by Peter Greenaway, which leads people through Rembrandt's paintings into 17th century Amsterdam,

---

[9] http://www.microsoft.com
[10] http://www.powerset.com/
[11] http://www.microsoft.com
[12] http://www.metaweb.com/
[13] http://www.tripit.com/
[14] http://www.uptake.com/
[15] http://www.facebook.com/
[16] http://movielens.umn.edu/login

and she can easily find it at e.g. Amazon[17].

Is this a realistic scenario? Yes, we strongly think so, and it might be realized in a not too distant future. However, it requires that the results of various research projects mentioned in this thesis, such as the iCITY project and our own work in CHIP, are brought together and combined with ongoing research, such as the research of the NoTube[18] project, to fully profit from the semantically linked data. The NoTube project develops Web services for providing users with personalized and integrated content from broadcast, Web channels and social networks. They capture heterogenous user information from multiple sources, such as user activity of watching, favouriting and recommending, context of location and device, interests in cross-domain areas, and user identity (Palmisano et al. 2010). To effectively model this information, they propose a "layer cake" architecture, where each layer represents a different knowledge domain, such as temporal, spatial, geographic, music-specific and movie-specific, and they adopt a broad approach for the data schema, such as the mixing of SKOS conceptual schemes (for categories) with FOAF's RDF vocabulary for people-description (Schopman et al. 2010). Their work brings inspiration for supporting the scenario described above, in particular how to approach collecting and modeling distributed user information for personalization.

The scenario reminds me again of the metaphor Tim Berners-Lee used, "Let a thousand flowers bloom". Linked Data enables data to be opened up and connected on the Web. When connected via open standards, it enables various things to sprout from it. From my perspective, how to deal with users' social contexts for distributed personalization across applications and with multiple devices is one of the interesting topics for future research.

As a final remark, the research described in this thesis has explored the aspects of user modeling and personalized recommendation. However, it is just as start, much work lies ahead. As a PhD candidate with a Chinese background, I would like my research to yield results within the rich cultural heritage domain of China.

---

[17]http://www.amazon.com/
[18]http://notube.tv

Appendix A

CHIP User Model Example

Here we give a code fragment from an example of the CHIP user model. For each part, a short explanation is provided in the comment immediately preceding it.

```
@prefix : <http://www.chip-project.org/Rijksmuseum#>.
@prefix foaf: <http://xmlns.com/foaf/0.1/>.
@prefix sem: <http://semanticweb.cs.vu.nl/2009/04/event/>.
@prefix rdf: <http://www.w3.org/1999/02/22-rdf-syntax-ns#>.
@prefix rdfs: <http://www.w3.org/2000/01/rdf-schema#>.

<!-- A user participates in 8 events with the CHIP demonstrator -->

:Actor1 a sem:Actor;
    foaf:holdsAccount "rm_userID_1291214391418";
    foaf:accountServiceHomepage "http://www.chip-project.org/demo";
    sem:participatesIn :event_rm_userID_1291214391418_1 ,
                       :event_rm_userID_1291214391418_2 ,
                       :event_rm_userID_1291214391418_3 ,
                       :event_rm_userID_1291214391418_4 ,
                       :event_rm_userID_1291214391418_5 ,
                       :event_rm_userID_1291214391418_6 ,
                       :event_rm_userID_1291214391418_7 ,
                       :event_rm_userID_1291214391418_8 ;

<!-- Personal background about this user,
which is collected from the CHIP user questionnaire -->

    foaf:age "20-30";
    foaf:gender "female";
    :profession "student";
```

```
    :rijksmuseumVisitFrequency "once a year";
    :rijksmuseumCollectionFamiliarity "a bit familiar";
    :museumVisitFrequency "every few months";
    :museumVisitReason "participation" , "education" , "recreation";
    :artInterest "interested";
    :computerExperience "experienced";
    :recommenderWebsitesFrequency "every month or more".

<!-- Event 1: the user rates an artwork positively-->

:event_rm_userID_1291214391418_1 a sem:Event;
    sem:eventType :rating;
    :rated :artefactSK-C-251;
    :ratedValue "1".

<!-- Event 2: the user rates another artwork positively-->
:event_rm_userID_1291214391418_2 a sem:Event;
    sem:eventType :rating;
    :rated :artefactSK-A-718;
    :ratedValue "0.5".

<!-- Event 3: the user rates an artwork negatively-->

:event_rm_userID_1291214391418_3 a sem:Event;
    sem:eventType :rating;
    :rated :artefactBK-NM-11452;
    :ratedValue "-1".

<!-- Event 4: the user rates a concept positively-->

:event_rm_userID_1291214391418_4 a sem:Event;
    sem:eventType :rating;
    :rated <http://www.getty.edu/vocabularies/ulan#500115664>;
    :ratedValue "1".

<!-- Event 5: the user rates a concept negatively-->

:event_rm_userID_1291214391418_5 a sem:Event;
    sem:eventType :rating;
    :rated <http://www.cs.vu.nl/STITCH/pp/ic#not23E41>;
    :ratedValue "-0.5".
```

```
<!-- Event 6: the user participates a "Tour of favorites",
which includes all artworks that he/she rated positively -->

:event_rm_userID_1291214391418_6 a sem:Event;
     sem:eventType :tour;
     rdfs:label "Tour of favorites";
     a rdf:Seq;
     rdf:_1 :artefactSK-C-251;
     rdf:_2 :artefactSK-A-718.
     rdf:_1 rdfs:subPropertyOf rdfs:member.
     rdf:_2 rdfs:subPropertyOf rdfs:member.


<!-- Event 7: the user participates a "Tour of recommended artworks",
which includes all recommended artworks -->

:event_rm_userID_1291214391418_7 a sem:Event;
     sem:eventType :tour;
     rdfs:label "Tour of recommended artworks";
     a rdf:Seq;
     rdf:_1 :artefactSK-A-3981;
     rdf:_2 :artefactSK-A-384;
     rdf:_3 :artefactSK-A-182;
     rdf:_4 :artefactSK-A-383;
     rdf:_5 :artefactSK-C-230;
     rdf:_6 :artefactSK-C-229;
     rdf:_7 :artefactSK-A-385;
     rdf:_8 :artefactSK-C-150;
     rdf:_9 :artefactSK-A-4039;
     rdf:_10 :artefactSK-A-3059;
     rdf:_1 rdfs:subPropertyOf rdfs:member.
     rdf:_2 rdfs:subPropertyOf rdfs:member.
     rdf:_3 rdfs:subPropertyOf rdfs:member.
     rdf:_4 rdfs:subPropertyOf rdfs:member.
     rdf:_5 rdfs:subPropertyOf rdfs:member.
     rdf:_6 rdfs:subPropertyOf rdfs:member.
     rdf:_7 rdfs:subPropertyOf rdfs:member.
     rdf:_8 rdfs:subPropertyOf rdfs:member.
     rdf:_9 rdfs:subPropertyOf rdfs:member.
     rdf:_10 rdfs:subPropertyOf rdfs:member.


<!-- Event 8: the user views an artwork
during the "Tour of recommended artworks"-->
```

```
:event_rm_userID_1291214391418_8 a sem:Event;
     sem:eventType :viewing;
     sem:partOf :event_7;
     :viewed :artefactSK-A-3981.
```

Appendix B

Stages in the Development of the CHIP Tools

In this appendix we show a historical overview of the three CHIP tools (*Art Recommender, Tour Wizard* and *Mobile Tour Guide*) in different development stages, see Table B.

Tools	Main functionalities of the tools in different versions
Art Recommender	ver.1: rating artworks in a 2-degree scale, recommending topics with explanations and showing the user profile
	ver.2: rating and recommending both artworks and topics
	ver.3: rating artworks and topics in a 3-degree scale
	ver.4: rating artworks and topics in a 5-degree scale, adding artworks to museum tours
	ver.5: providing the new interface designed by Fabrique, searching artworks and topics
Tour Wizard	ver.1: viewing museum tours on the museum map
	ver.2: indicating how many artworks and which artworks from the tour are available in the museum and visualizing the tour also in a historical timeline
	ver.3: indicating the user's current location in the museum, providing the real-time adaptation of the tour route, and visualizing the location of artworks from the tour on Google Maps
Mobile Tour Guide	ver.1: RFID tag + PDA based prototype
	ver.2: iPod based prototype

Figure B.1: Art Recommender (ver.1) http://www.chip-project.org/demo1/ Main functionalities: presenting 9 artworks at the same time, rating artworks in a 2-degree scale, recommending topics (art concepts), providing explanations of "why recommended" for recommended topics and showing the user profile.
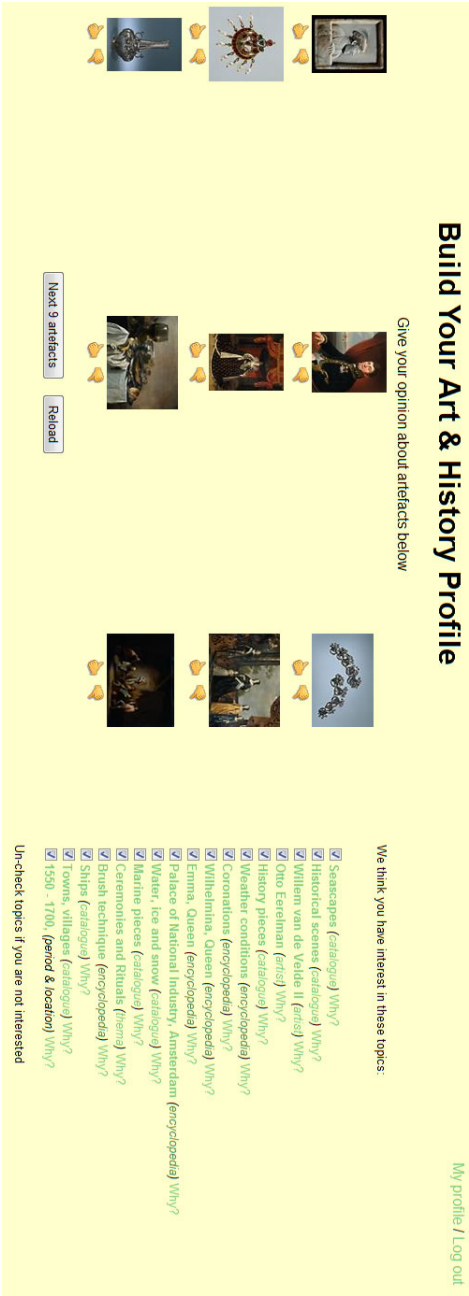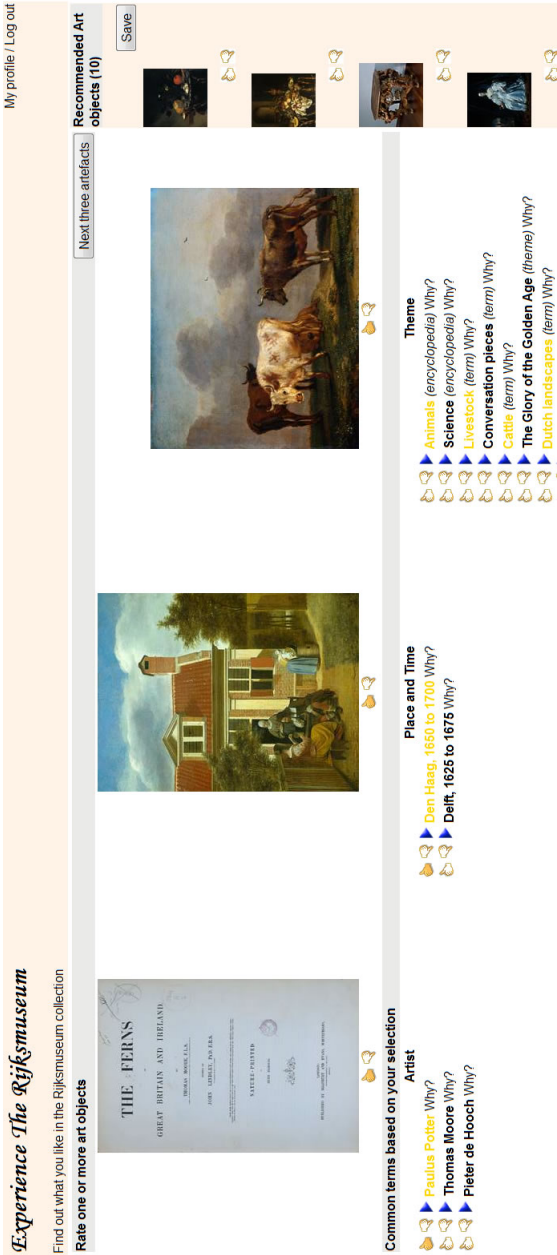
Figure B.2: Art Recommender (ver.2) http://www.chip-project.org/demo2/
Main functionalities: presenting 3 artworks at the same time, recommending both artworks and topics and grouping recommended topics in the categories of "Artist", "Place and time" and "Theme".
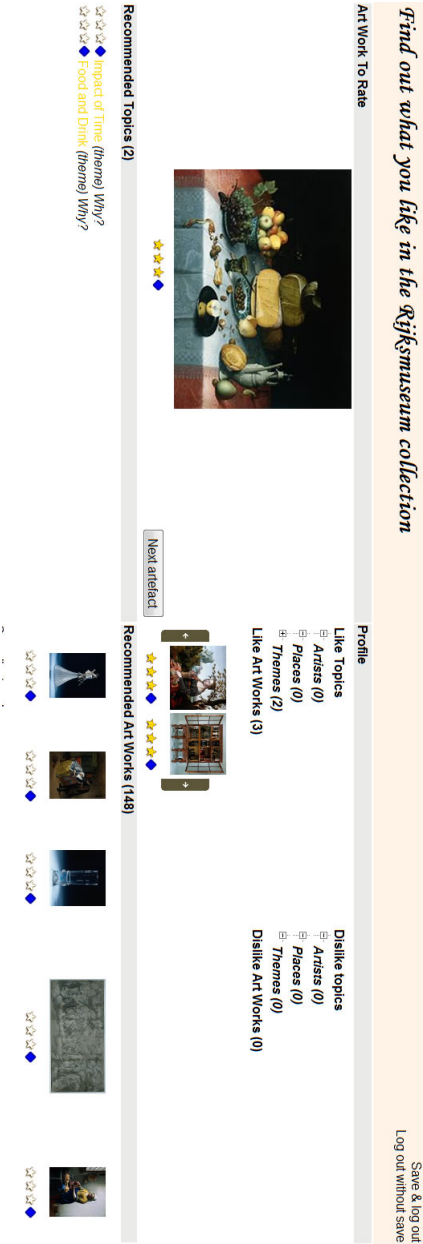
Figure B.3: Art Recommender (ver.3) http://www.chip-project.org/demo3/ Main functionalities: presenting 1 artwork, rating artworks and topics in a 3-degree scale and showing the user profile in the main interface.
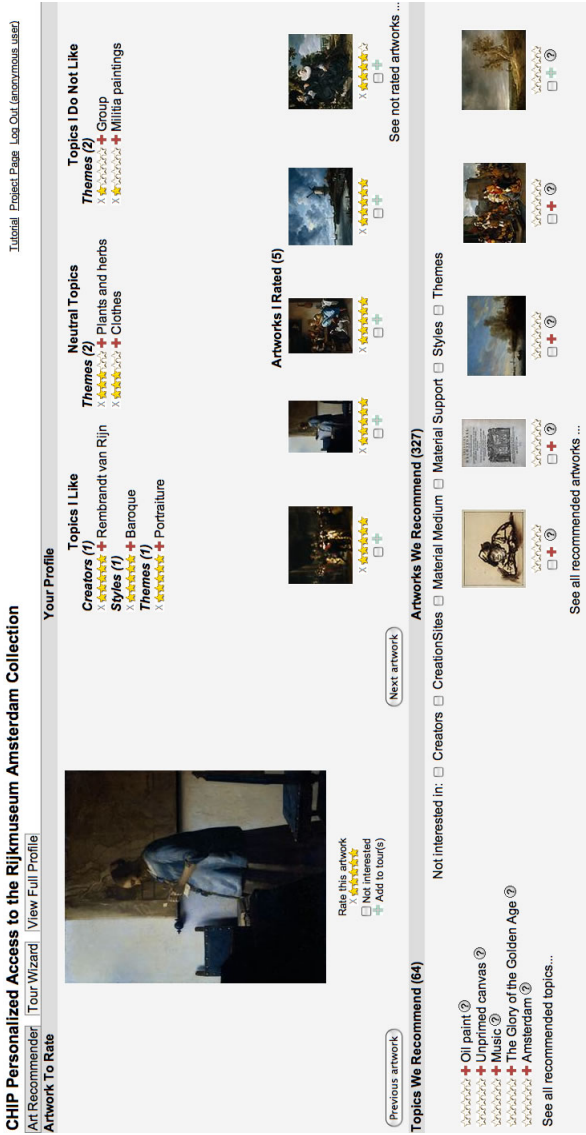
Figure B.4: Art Recommender (ver.4) http://www.chip-project.org/demo4/
Main functionalities: rating artworks and topics in a 5-degree scale, adding artworks to museum tours, excluding artworks for recommendations by selecting "Not interested" and filtering recommended topics in the categories of "Creators", "Creation-Sites", "Material Medium", "Material Support", "Styles" and "Themes".
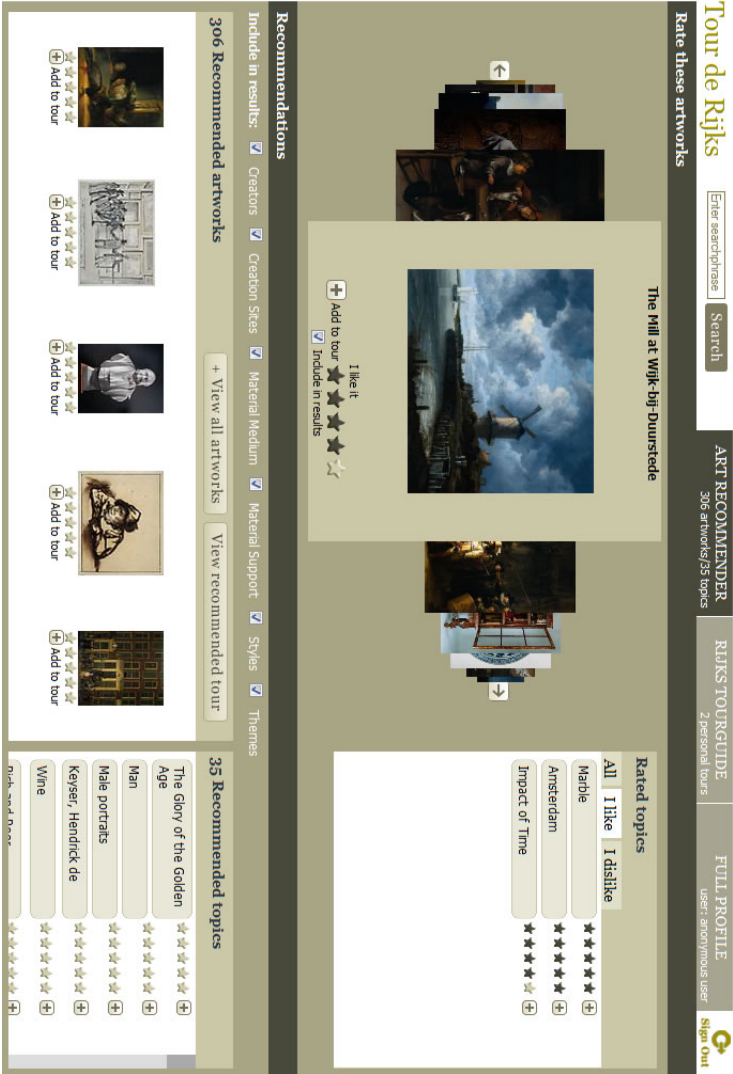
Figure B.5: Art Recommender (ver.5) http://www.chip-project.org/demo5/
Main functionalities: using the new interface designed by Fabrique, presenting the artworks for rating in an artwork carousel
and providing a search function.
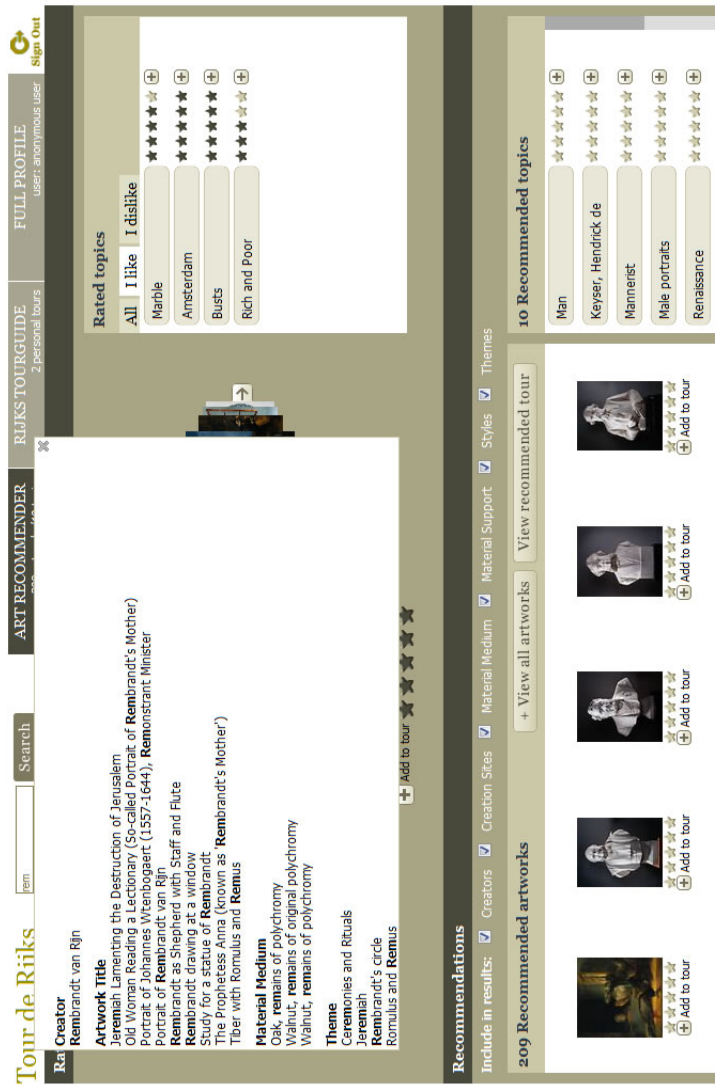
Figure B.6: Art Recommender (ver.5): Detailed view of using the search function.

Figure B.7: Art Recommender (ver.5): Detailed view of presenting an artwork.
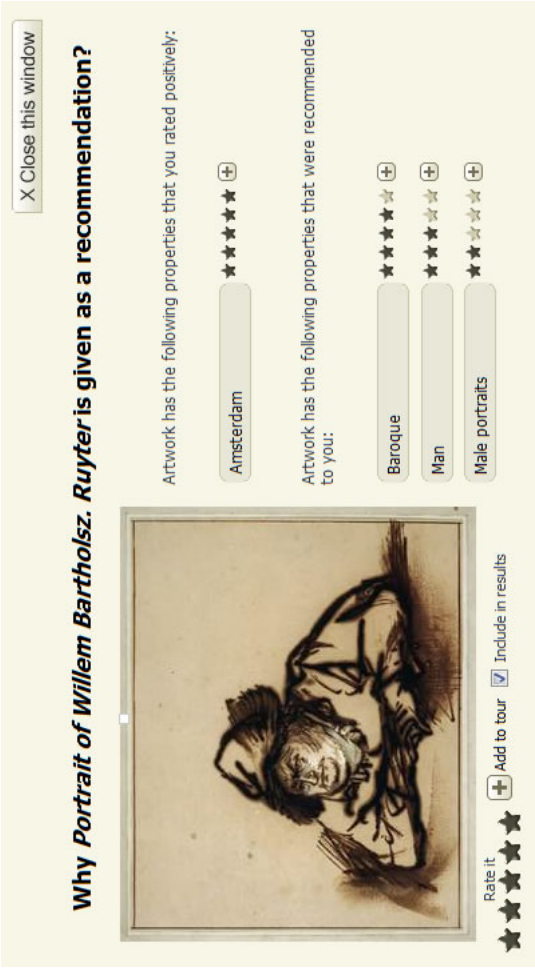
Figure B.8: Art Recommender (ver.5): Detailed view of "Why recommended" for a recommended artwork.
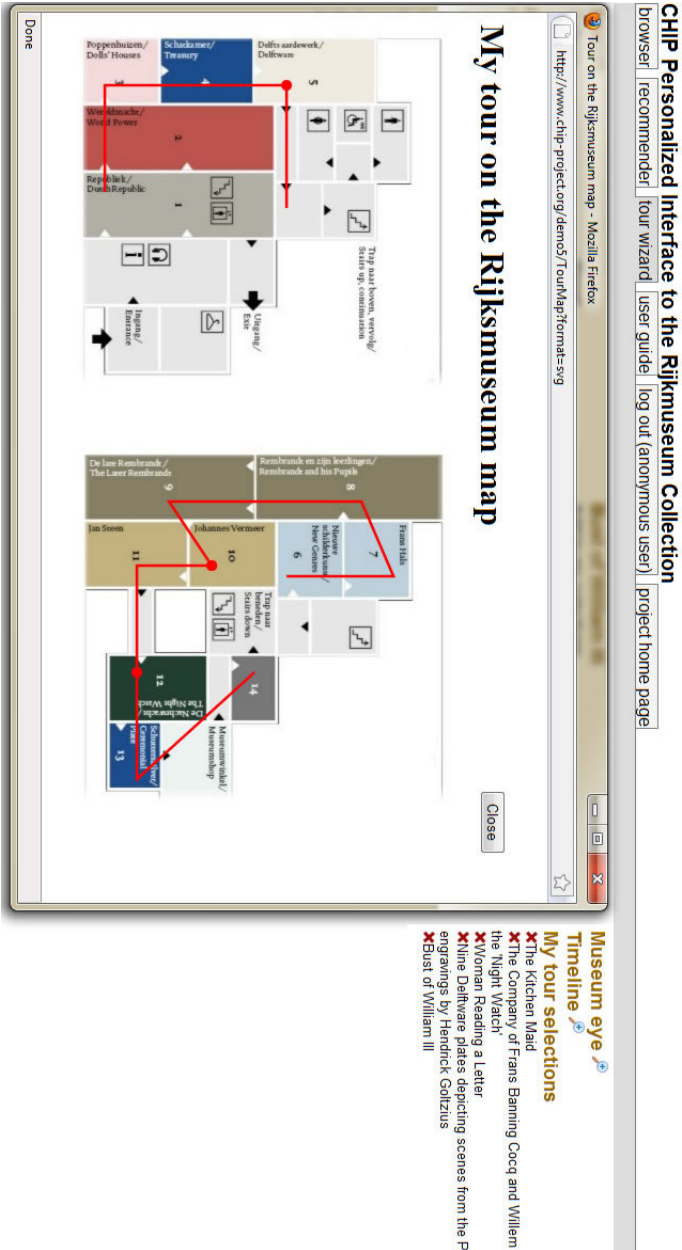
Figure B.9: Tour Wizard (ver.1) http://www.chip-project.org/demo4/tour_wizard.jsp
Main functionalities: visualizing museum tours on the museum map
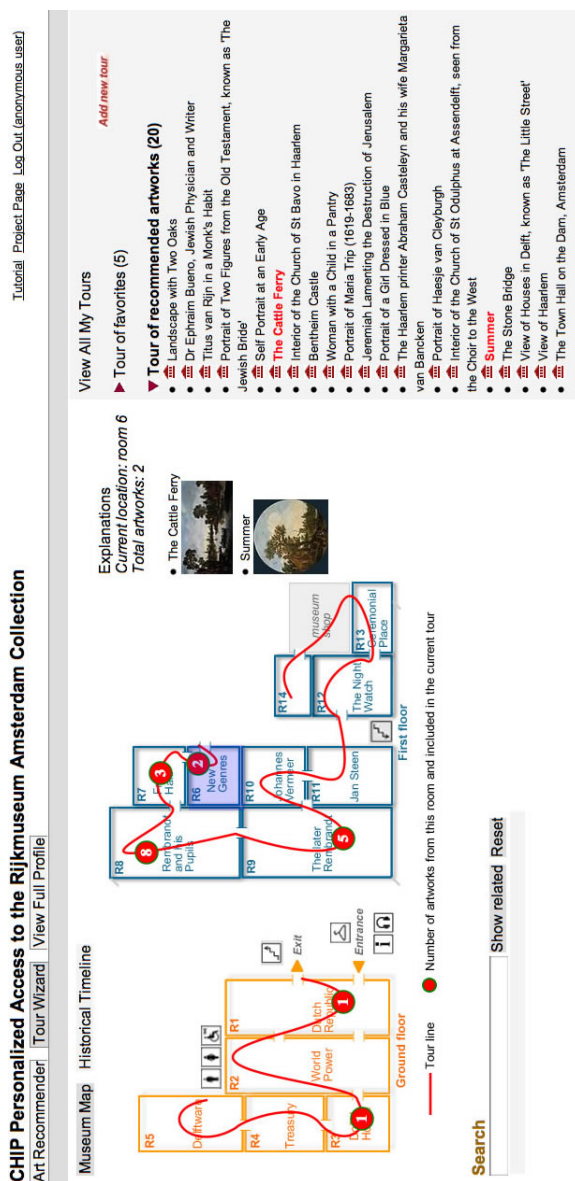
Figure B.10: Tour Wizard (ver.2) http://www.chip-project.org/demo4/tour_wizard.jsp
Main functionalities: indicating how many artworks and which artworks from the tour are currently available in the museum and visualizing the tour either on the museum map or in a historical timeline.
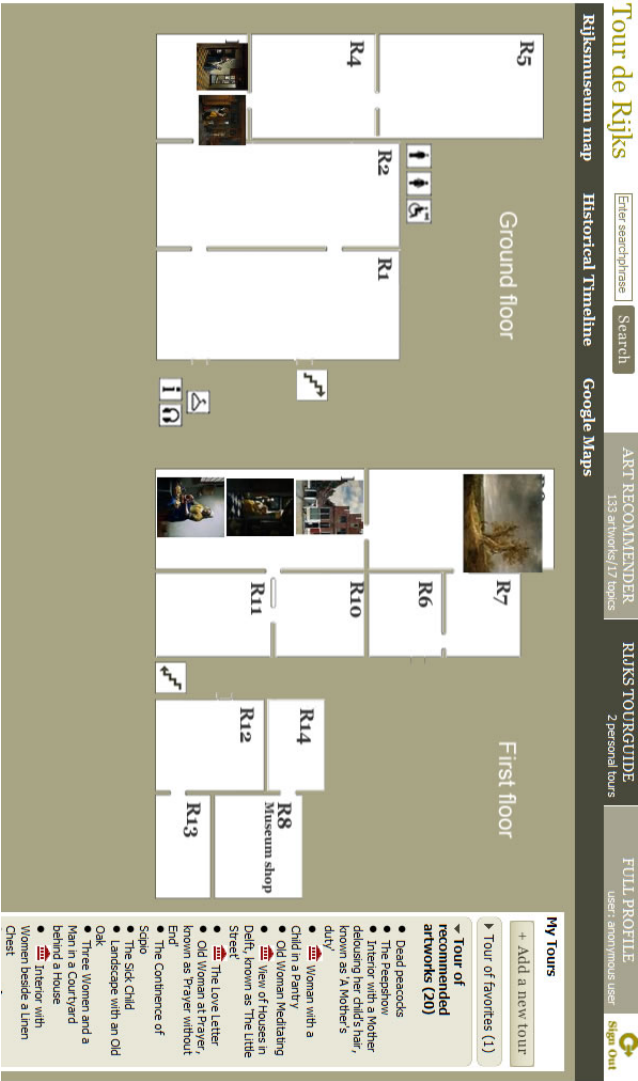
Figure B.11: Tour Wizard (ver.3) http://www.chip-project.org/demo5/tour_wizard.jsp
Main functionalities: showing available artworks from the tour in the museum exhibition rooms.
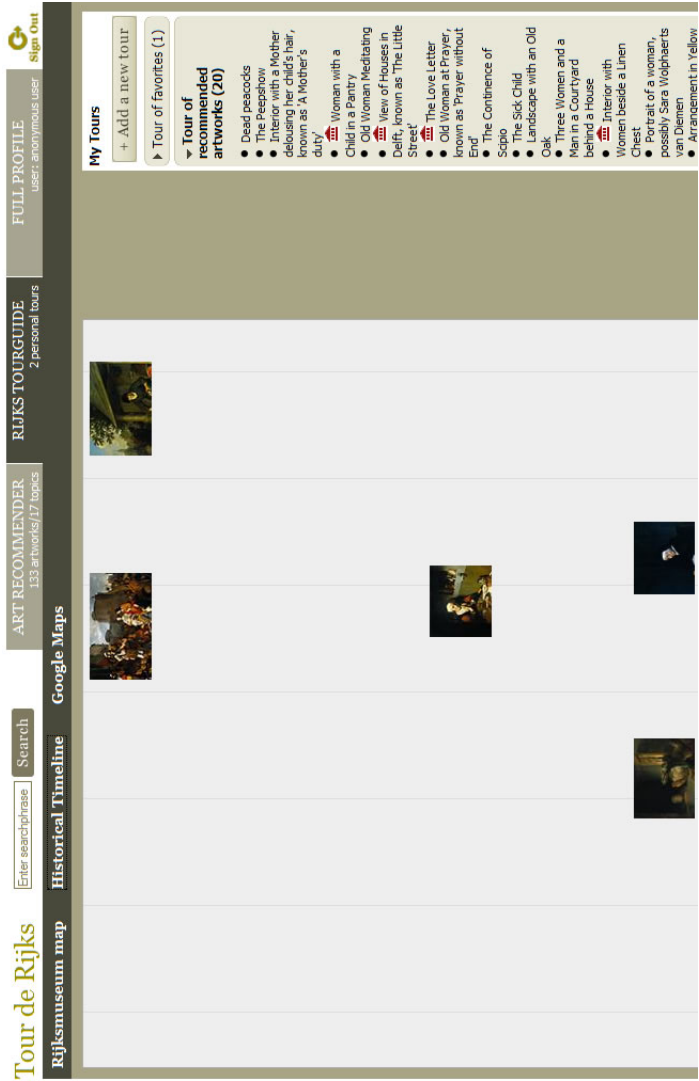
Figure B.12: Tour Wizard (ver.3): Detailed view of the museum tour on a historical time-line.
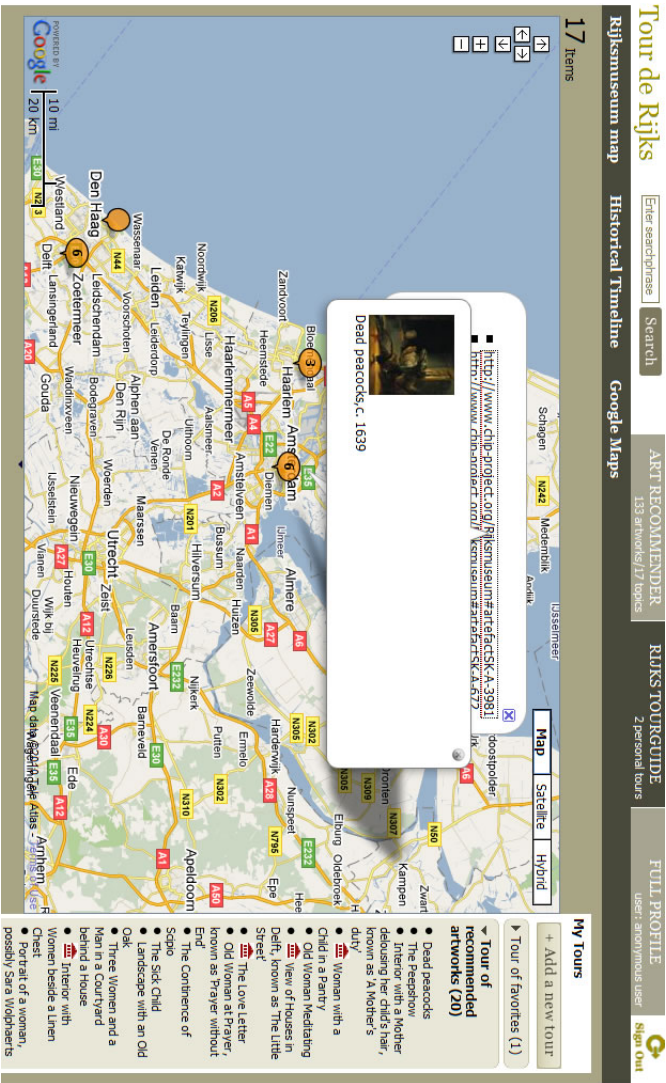
Figure B.13: Tour Wizard (ver.3): Detailed view of the creation sites of artworks from the tour on Google Maps.
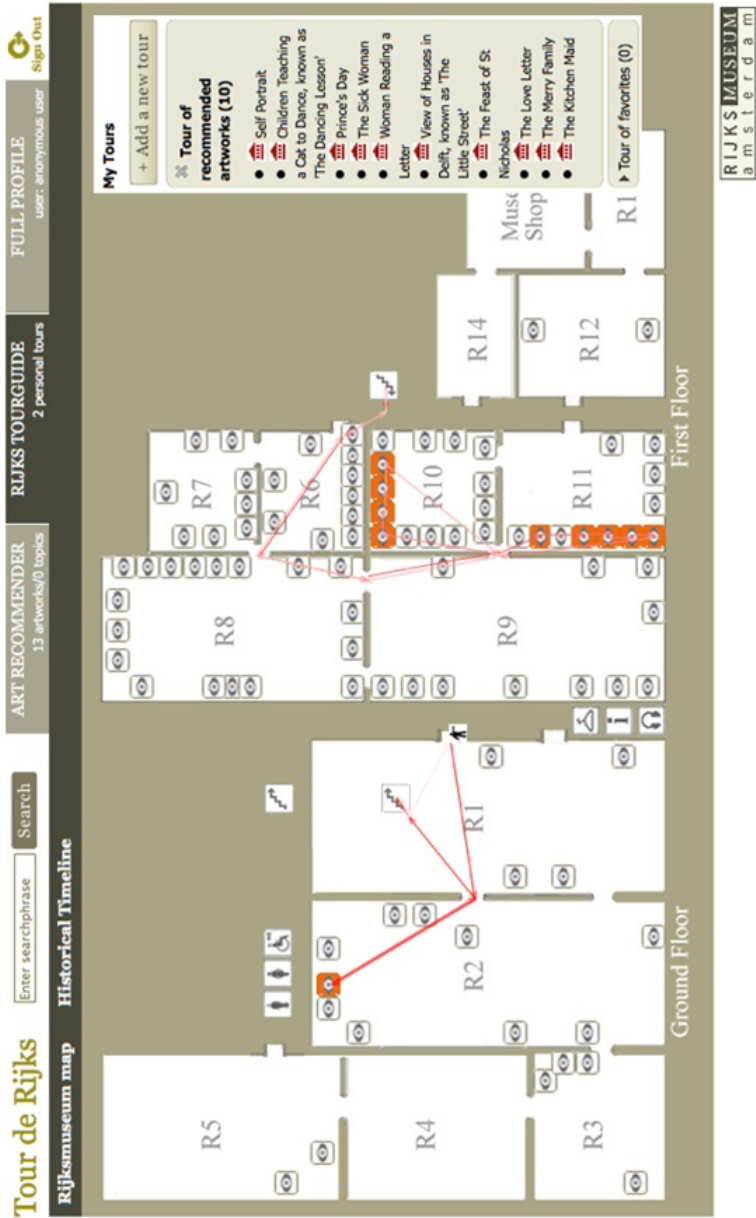
Figure B.14: Tour Wizard (ver.3): Detailed view of the tour, indicating the user's current location in the museum.
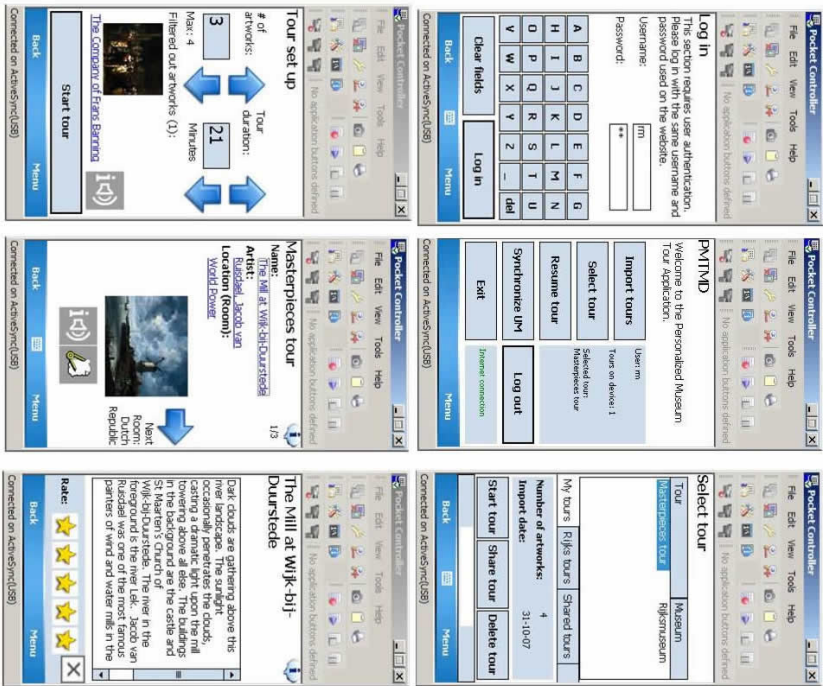
Figure B.15: Mobile Tour Guide (ver.1)Platform: RFID tag + PDA based prototype.
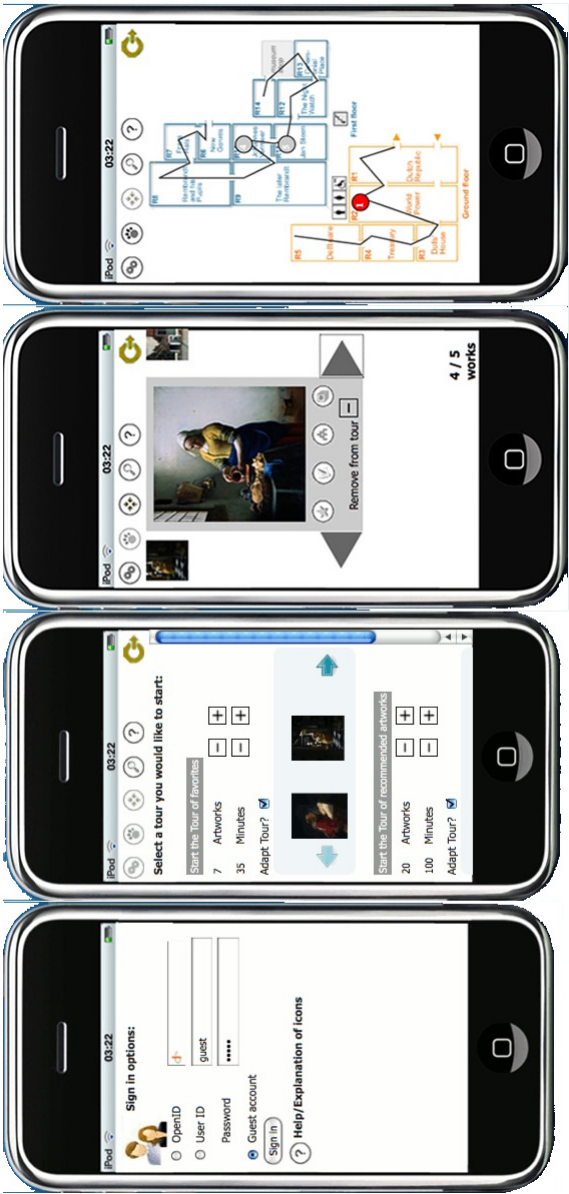
Figure B.16: Mobile Tour Guide (ver.2) http://www.chip-project.org/demo/mobileguide/index.jsp Platform: iPod based prototype.

References

Adomavicius, G. and A. Tuzhilin (2005). Toward the next generation of recommender systems: A survey of the state-of-the-art and possible extensions. *IEEE Transactions on Knowledge and Data Engineering 17*(6), 734–749.

Adrian, B., L. Sauermann, and T. Roth-Berghofer (2007). Contag: A semantic tag recommendation system. In *Proceedings of the New Media Technology and Semantic Systems*, pp. 297–304.

Albertoni, R., A. Bertone, and M. D. Martino (2004). Semantic web and information visualization. In *Proceedings of the 1st Italian Workshop on Semantic Web Application and Perspective*, pp. 108–114.

Alexander, C., A. Burnette, D. Dark, D. Hart, J. Rossi, and N. Minor (2008). Beyond launch: Museum videos on youtube. In *Proceedings of the International Conference on Museums and the Web (MW'08)*, pp. 49–57.

Anderson, P. and A. Blackwood (2004). Mobile and pda technologies and their future use in education. *JISC Technology and Standards Watch 04*(03).

Aoki, P.M., G. R.-H. A. S. M. T. J. W. A. (2002). Sottovoce: exploring the interplay of conversation and mobile audio spaces. In *Proceedings of the SIGCHI Conference on Human factors in computing systems (CHI'02)*, pp. 431–438.

Ardissono, L., A. Goy, G. Petrone, M. Segnan, and P. Torasso (2003). Intrigue: Personalized recommendation of tourist attractions for desktop and handset devices. *Applied Artificial Intelligence 17*(8), 687–714.

Aroyo, L., N. Stash, Y. Wang, P. Gorgels, and L. Rutledge (2007). Chip demonstrator: Semanticsdriven recommendations and museum tour generation. In *Proceedings of the 6th International Semantic Web Conference and 2nd Asian Semantic Web Conference (ISWC'07 + ASWC'07)*, pp. 879–886.

Baeza-Yates, R. and B. Ribeiro-Neto (1999). *Modern Information Retrieval.*

Barry, A. (2006). Creating a virtuous circle between a museum's on-line and physical spaces. In *Proceedings of the International Conference on Museums and the Web (MW'06)*.

Bartneck, C., A. Masuoka, T. Takahashi, and T. Fukaya (2006). The learning experience with electronic museum guides. *Psychology of Aesthetics, Creativity, and the Arts 1*, 18–25.

Bearman, D. and J. Trant (2008). Technologies, like museums, are social. In *Proceedings of the International Conference on Museums and the Web (MW'08)*, pp. 3–8.

Berkovsky, S., T. Kuflik, and F. Ricci (2007). Cross-domain mediation in collaborative filtering. In *Proceedings of the 11th International Conference on User Modeling (UM'07)*, pp. 355–359.

Berners-Lee, T., J. Hendler, and O. Lassila (2001). The semantic web. *Scientific American 284*(5), 34–43.

Berry, M., S. Dumais, and G. O. Brien (1995). Using linear algebra for intelligent information retrieval. *SIAM Review*, 573–595.

Blanco-Fernández, Y., J. J. P. Arias, A. Gil-Solla, M. Ramos-Cabrer, and M. Lopez-Nores (2008). Semantic reasoning: A path to new possibilities of personalization. In *Proceedings of the 5th European Semantic Web Conference (ESWC'08)*, pp. 720–735.

Bojars, U., J. G. Breslin, and A. Finn (2008). Using the semantic web for linking and reusing data across web 2.0 communities. *Web Semantics 6*(1), 21–28.

Bowen, J. P. and S. Filippini-Fantoni (2004). Personalization and the web from a museum perspective. In *Proceedings of the International Conference on Museums and the Web (MW'04)*.

Brusilovsky, P., A. Kobsa, and W. Nejdl (2007). *The Adaptive Web, Methods and Strategies of Web Personalization*, pp. 291–324.

Burke, R. (2002). Hybrid recommender systems: Survey and experiments. *User Modelling and User-Adapted Interaction (UMUAI) 12*(4), 331–370.

Cantador, I., A. Bellogin, and P. Castells (2008). A multilayer ontology-based hybrid recommendation model. *AI Communications 21*(2-3), 203–210.

Carmagnola, F., F. Cena, L. Console, O. Cortassa, C. Gena, A. Goy, I. Torre, A. Toso, and F. Vernero (2008). Tag-based user modeling for social multi-device adaptive guides. *User Modeling and User-Adapted Interaction (UMUAI) 18*(5), 497–538.

Carmagnola, F., F. Cena, O. Cortassa, C. Gena, and I. Torre (2007). Towards a tag-based user model: How can user model benefit from tags? In *Proceedings*

of the 11th International Conference on User Modeling (UM'07), pp. 445–449.

Carmagnola, F., F. Cena, O. Cortassa, C. Gena, Y. Wang, N. Stash, and L. Aroyo (2008). Tag interoperability in cultural web-based applications. In *Proceedings of the 19th ACM conference on Hypertext and hypermedia (Hypertext'08)*, pp. 221–222.

Carmagnola, F., F. Cena, and C. Gena (2007). User modeling in the social web. In *Proceedings of the International Conference on Knowledge-Based Intelligent Information and Engineering Systems (KES'07)*, pp. 745–752.

Celma, O. (2006). Foafing the music: Bridging the semantic gap in music recommendation. In *Proceedings of the 5th International Semantic Web Conference (ISWC'06)*, pp. 927–934.

Chan, S. (2007). Tagging and searching: Serendipity and museum collection databases. In *Proceedings of the International Conference on Museums and the Web (MW'07)*, pp. 87–100.

Chan, S. (2008). Towards new metrics of success for on-line museum projects. In *Proceedings of the International Conference on Museums and the Web (MW'08)*, pp. 13–22.

Chou, S.C., H. W.-G. F. L. S. N. (2005). Semantic web technologies for context-aware museum tour guide applications. In *Proceedings of the 19th International Conference on Advanced Information Networking and Applications*, pp. 709–714.

Claypool, M., A. Gokhale, T. Miranda, P. Murnikov, D. Netes, and M. Sartin (1999). Combining content-based and collaborative filters in an online newspaper. In *Proceedings of ACM SIGIR Workshop on Recommender Systems*.

Cramer, H., V. Evers, S. Ramlal, M. Someren, L. Rutledge, N. Stash, L. Aroyo, and B. Wielinga (2008). The effects of transparency on trust and acceptance in interaction with a content-based art recommender. *User Modelling and User-Adapted Interaction (UMUAI) 18*(5), 455–496.

Denaux, R., V. Dimitrova, and L. Aroyo (2005). Integrating open user modeling and learning content management for the semantic web. In *Proceedings of the 10th International Conference on User Modeling (UM'05)*, pp. 23–29.

Graziola, I., F. Pianesi, M. Zancanaro, and D. Goren-Bar (2005). Dimensions of adaptivity in mobile systems: personality and people's attitudes. In *Proceedings of the International Conference on Intelligent User Interfaces (IUI'05)*, pp. 223–230.

Greaves, M. and P. Mika (2008). Semantic web and web 2.0. *Web Semantics 6*(1), 1–3.

Gruber, T. (2008). Collective knowledge systems: Where the social web meets the semantic web. *Web Semantics 6*(1), 4–13.

Heckmann, D., T. Schwartz, B. Brandherm, M. Schmitz, and M. von Wilamowitz-Moellendorff (2005). Gumo the general user model ontology. In *Proceedings of the 10th International Conference on User Modeling (UM'05)*, pp. 428–432.

Herlocker, J., J. Konstan, A. Borchers, and J. Riedl (2004). Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems*, 5–53.

Hollink, L., G. Schreiber, and B. J. Wielinga (2007). Patterns of semantic relations to improve image content search. *Web Semantics 5*(3), 195–203.

Hook, K., J. Karlgren, N. D. A. Waern, C. G. Jansson, K. Karlgren, and B. Lemaire (1996). A glass box approach to adaptive hypermedia. *User Modeling and User-Adapted Interaction (UMUAI) 6*(2-3), 157–184.

Hyvonen, E., E. Makela, M. Salminen, A. Valo, K. Viljanen, S. Saarela, M. Junnila, and S. Kettula (2005). Museumfinland - finnish museums on the semantic web. *Web Semantics: Science, Services and Agents on the World Wide Web 3*(2), 224–241.

Ide, N. and J. Veronis (1998). Introduction to the special issue on word sense disambiguation: The state of the art. *Computational Linguistics 24*(1), 1–40.

Isaac, A., L. V. der Meij, S. Schlobach, and S. Wang (2007). An empirical study of instance-based ontology matching. In *Proceedings of the 6th International Semantic Web Conference and 2nd Asian Semantic Web Conference (ISWC'07 + ASWC'07)*, pp. 252–266.

Kelly, B., M. Ellis, and R. Gardler (2008). What does openness mean to the museum community? In *Proceedings of the International Conference on Museums and the Web (MW'08)*, pp. 195–204.

Kobsa, A. (2001). Generic user modeling systems. *User Modeling and User-Adapted Interaction (UMUAI) 11*(1), 49–63.

Kobsa, A., J. Koenemann, and W. Pohl (2001). Personalized hypermedia presentation techniques for improving online customer relationships. *The Knowledge Engineering Review 16*(2), 111–155.

Kruger, A., J. Baus, D. Heckmann, M. Kruppa, and R. Wasinger (2007). *Adaptive Mobile Guides*, pp. 521–549.

Kuflik, T., J. Sheidin, S. Jbara, G. oren Bar, P. Soffer, O. Stock, and M. Zancanaro (2007). Supporting small groups in the museum by context-aware communication services. In *Proceedings of the 12th International conference on Intelligent user interfaces (IUI'07)*, pp. 305–308.

Luyten, K., H. V. Loon, D. Teunkens, K. Gabriels, K. Coninx, and E. Manshoven (2006). Disclosing a museum by a socially-aware mobile guide. In *Proceedings of the 7th International Symposium on Virtual Reality, Archaeology and Cultural Heritage*.

Mobasher, B., X. Jin, and Y. Zhou (2004). Semantically enhanced collaborative filtering on the web. In *Proceedings of the 1st European Web Mining Forum (EWMF'03)*, pp. 57–76.

Morita, M. and Y. Shinoda (1994). Information filtering based on user behavior analysis and best match text retrieval. In *Proceedings of ACM SIGIR conference on Research and development in information retrieval*, pp. 272–281.

Navigli, R. and P. Velardi (2003). An analysis of ontology-based query expansion strategies. In *Proceedings of Machine Learning Conference*, pp. 42–49.

Oppermann, R. and M. Specht (1999). A nomadic information system for adaptive exhibition guidance. In *Proceedings of the International Conference on Hypermedia and Interactivity in Museums*, pp. 103–109.

O'Reilly, T. (2005). What is web 2.0. Website. `http://oreilly.com/web2/archive/what-is-web-20.html`.

Palmisano, D., M. Minno, D. Brickley, and L. Miller (2010). User and context model specification v2 (d3.1). Website. `http://www.notube.tv/images/stories/deliverables/d3.1.pdf`.

Pazzani, M. J. (1999). A framework for collaborative, content-based and demographic filtering. *Artificial Intelligence Review 13*(5), 393–408.

Peis, E., J. M. M. del Castillo, and J. A. Delgado-Lopez (2008). Semantic recommender systems. analysis of the state of the topic. In *Proceedings of Hipertext*, pp. 1–9.

Razmerita, L., A. Angehrn, and A. Maedche (2003). Ontology-based user modeling for knowledge management systems. In *Proceedings of the 9th international conference on User modeling (UM'03)*, pp. 213–217.

Resnick, P., N. Iacovou, M. Suchak, P. Bergstrom, and J. Riedl (1994). Grouplens: an open architecture for collaborative filtering of netnews. In *Proceedings of the ACM conference on Computer supported cooperative work*, pp. 175–186.

Ricci, F., B. Arslan, N. Mirzadeh, and A. Venturini (2002). Itr: a case-based travel advisory system. In *Proceedings of European Conference on Case-Based Reasoning*, pp. 613–627.

Rocchi, C., O. Stock, M. Zancanaro, M. Kruppa, and A. Kruger (2004). The museum visit: Generating seamless personalized presentations on multiple devices. In *Proceedings of the 9th International Conference on Intelligent User Interface (IUI'04)*, pp. 316–318.

Roes, I., N. Stash, Y. Wang, and L. Aroyo (2009). A personalized walk through the museum: The chip interactive tour guide. In *Proceedings of the 27th SIGCHI Conference on Human Factors in Computing Systems (CHI'09)*, pp. 3317–3322.

Ruotsalo, T. and E. Hyvonen (2007). An event-based approach for semantic metadata interoperability. In *Proceedings of the 6th International Semantic Web Conference and 2nd Asian Semantic Web Conference (ISWC'07 + ASWC'07)*, pp. 409–422.

Sakamura, K. (2003). Personalized digital museum assistant. Website. `http://www.um.u-tokyo.ac.jp/publish_db/2000dm2k/english/01/01-17.html`.

Schopman, B., D. Brickley, L. Aroyo, C. van Aart, V. Buser, R. Siebes, L. Nixon, L. Miller, V. Malaise, M. Minno, M. Mostarda, D. Palmisano, and Y. Raimond (2010). Notube: making the web part of personalised tv. In *Proceedings of the Extending the Frontiers of Society Online (WebSci'10)*.

Schreiber, G., A. Amin, L. Aroyo, M. van Assem, V. de Boer, L. Hardman, M. Hildebrand, B. Omelayenko, J. van Ossenbruggen, A. Tordai, J. Wielemaker, and B. Wielinga (2008). Semantic annotation and search of cultural-heritage collections: The multimediaN E-Culture demonstrator. *Web Semantics: Science, Services and Agents on the World Wide Web 6*(4), 243–249.

Schweibenz, W. (1998). The virtual museum: New perspectives for museums to present objects and information using the internet as a knowledge base and communication system. In *Proceedings of Internationalen Symposiums fur Informationswissenschaft (ISI'98)*.

Sparacino, F. (2002). The museum wearable. In *Proceedings of the International Conference on Museums and the Web (MW'02)*.

Sumi, Y., M.-K. (2004). Interface agents that facilitate knowledge interactions between community members. In *Proceedings of the International Conference on Cognitive Technologies Series*, pp. 405–427.

Tordai, A., B. Omelayenko, and G. Schreiber (2007). Thesaurus and metadata alignment for a semantic e-culture application. In *Proceedings of the 4th international conference on Knowledge capture (K-CAP'07)*, pp. 199–200.

Trant, J. and B. Wyman (2006). Investigating social tagging and folksonomy in art museums with steve.museum. In *Proceedings of the Collaborative Tagging Workshop at World Wide Web (WWW'06)*.

Tudhope, D., C. Binding, D. Blocks, and D. Cunliffe (2006). Query expansion via conceptual distance in thesaurus indexed collections. *Documentation* (4), 509–533.

van Assem, M., M. Menken, G. Schreiber, J. Wielemaker, and B. Wielinga (2004). A method for converting thesauri to rdf/owl. In *Proceedings of the International Semantic Web Conference (ISWC'04)*, pp. 17–31.

van Gendt, M., A. Isaac, L. van der Meij, and S. Schlobach (2006). Semantic web techniques for multiple views on heterogeneous collections: A case study. In *Proceedings of 10th European Conference on Research and Advanced Technology for Digital Libraries (ECDL'06)*, pp. 426–437.

van Hage, W. R., N. Stash, Y. Wang, and L. Aroyo (2010). Finding your way through the rijksmuseum with an adaptive mobile museum guide. In *Proceedings of the 5th European Semantic Web Conference (ESWC'10)*, pp. 46–59.

van Hage, W. R., J. Wielemaker, and G. Schreiber (2009). The space package: Tight integration of space and semantics. In *Proceedings of the 8th International Semantic Web Conference Workshop: TerraCognita*.

van Harmelen, F., A. ten Teije, and H. Wache (2009). Knowledge engineering rediscovered: towards reasoning patterns for the semantic web. In *Proceedings of the 5th International Conference on Knowledge Capture (K-CAP'09)*, pp. 81–88.

van Setten, M., R. Brussee, H. van Vliet, L. Gazendam, Y. van Houten, and M. Veenstra (2006). On the importance of "who tagged what". In *Proceedings of the Workshop on the Social Navigation and Community based Adaptation Technologies at AH'06*.

Wang, Y., L. Aroyo, N. Stash, and L. Rutledge (2007). Interactive user modeling for personalized access to museum collections: The Rijksmuseum case study. In *Proceedings of the 11th International Conference on User Modeling (UM'07)*, pp. 25–29.

Wang, Y., L. Aroyo, N. Stash, R. Sambeek, Y. S. G. Schreiber, and P. Gorgels (2009a). Cultivating personalized museum tours online and on-site. *Interdisciplinary Science Reviews 34*(2), 139–153.

Wang, Y., F. Cena, F. Carmagnola, O. Cortassa, C. Gena, N. Stash, and L. Aroyo (2008a). Rss-based interoperability for user adaptive systems. In *Proceedings of the Adaptive Hypermedia conference (AH'08)*.

Wang, Y., N. Stash, L. Aroyo, P. Gorgels, L. Rutledge, and G. Schreiber (2008b). Recommendations based on semantically-enriched museum collections. *Web Semantics: Science, Services and Agents on the World Wide Web 6*(4), 283–290.

Wang, Y., N. Stash, L. Aroyo, L. Hollink, and G. Schreiber (2009b). Semantic relations for content-based recommendations. In *Proceedings of the 5th International Conference on Knowledge Capture (K-CAP'09)*, pp. 209–210.

Wang, Y., N. Stash, L. Aroyo, L. Hollink, and G. Schreiber (2009c). Using semantic relations for content-based recommender systems in cultural heritage. In *Proceedings of the Workshop on Ontology Patterns (WOP) at the 8th International Semantic Web Conference (ISWC'09)*, pp. 1–13.

Wang, Y., R. van Sambeek, Y. Schuurmans, L. Aroyo, N. Stash, L. Rutledge, and P. Gorgels (2008c). Be Your Own Curator with the CHIP Tour Wizard. In *Proceedings of the International Conference on Museums and the Web (MW'08)*, pp. 133–140.

Wang, Y., S. Wang, N. Stash, L. Aroyo, and G. Schreiber (2010). Enhancing content-based recommendation with the task model of classification. In *Proceedings of the 17th Internaltional Conference on Knowledge Engineering and Management by the Masses (EKAW'10)*, pp. 431–440.

Zakaria, M. R., A. Moore, H. Ashman, C. D. Stewart, and T. J. Brailsford (2002). The hybrid model for adaptive educational hypermedia. In *Proceedings of the 2nd International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems (AH'02)*, pp. 580–585.

Ziegler, C. N. (2004). Semantic web recommender systems. In *Proceedings of the Joint ICDE/EDBT Ph.D. Workshop*, pp. 78–89.

Summary

In the Web 2.0 environment, institutes and organizations are starting to open up
their previously isolated and heterogeneous collections in order to provide visitors
with maximal access. Semantic Web technologies act as instrumental in integrat-
ing these rich collections of metadata by defining ontologies which accommodate
different representation schemata and inconsistent naming conventions over the
various vocabularies. Facing the large amount of metadata with complex semantic
structures, it is becoming more and more important to support visitors with a
proper selection and presentation of information. In this context, the Dutch Sci-
ence Foundation (NWO) funded the Cultural Heritage Information Personalization
(CHIP[1]) project in early 2005, as part of the Continuous Access to Cultural Her-
itage (CATCH[2]) program in the Netherlands. It is a collaborative project between
the Rijksmuseum Amsterdam[3], the Eindhoven University of Technology[4] and the
Telematica Instituut[5].

The problem statement that guides the research of this thesis is as follows: *Can
we support visitors with personalized access to semantically-enriched collections*?
To study this question, we chose cultural heritage (museums) as an application
domain, and the semantically rich background knowledge about the museum col-
lection provides a basis to our research. On top of it, we deployed user modeling
and recommendation technologies in order to provide personalized services for
museum visitors. Our main contributions are: (i) we developed an interactive
rating dialog of artworks and art concepts for a quick instantiation of the CHIP
user model, which is built as a specialization of FOAF[6] and mapped to an existing
event model ontology SEM[7]; (ii) we proposed a hybrid recommendation algorithm,

---

[1]http://www.chip-project.org/
[2]http://www.nwo.nl/catch
[3]http://www.rijksmuseum.nl/
[4]http://w3.tue.nl/
[5]http://www.novay.nl/en/
[6]http://www.foaf-project.org/
[7]http://semanticweb.cs.vu.nl/2009/11/sem/

combining both explicit and implicit relations from the semantic structure of the collection. On the presentation level, we developed three tools for end-users: Art Recommender, Tour Wizard and Mobile Tour Guide. Following a user-centered design cycle, we performed a series of evaluations with museum visitors to test the effectiveness of recommendations using the rating dialog, different ways to build an optimal user model and the prediction accuracy of the hybrid algorithm.

Chapter 1 introduces the research questions, our approaches and the outline of this thesis.

Chapter 2 gives an overview of our work at the first stage. It includes (i) the semantic enrichment of the Rijksmuseum collection, which is mapped to three Getty vocabularies[8] (ULAN, AAT, TGN) and the Iconclass thesaurus[9]; (ii) the minimal user model ontology defined as a specialization of FOAF, which only stores user ratings at that time, (iii) the first implementation of the content-based recommendation algorithm in our first tool, the CHIP Art Recommender.

Chapter 3 presents two other tools: Tour Wizard and Mobile Tour Guide. Based on the user's ratings, the Web-based Tour Wizard recommends museum tours consisting of recommended artworks that are currently available for museum exhibitions. The Mobile Tour Guide converts recommended tours to mobile devices (e.g. PDA) that can be used in the physical museum space. To connect users' various interactions with these tools, we made a conversion of the online user model stored in RDF into XML format which the mobile guide can parse, and in this way we keep the online and on-site user models dynamically synchronized.

Chapter 4 presents the second generation of the Mobile Tour Guide with a real time routing system on different mobile devices (e.g. iPod). Compared with the first generation, it can adapt museum tours based on the user's ratings artworks and concepts, her/his current location in the physical museum and the coordinates of the artworks and rooms in the museum. In addition, we mapped the CHIP user model to an existing event model ontology SEM. Besides ratings, it can store additional user activities, such as following a tour and viewing artworks.

Chapter 5 identifies a number of semantic relations within one vocabulary (e.g. a concept has a broader/narrower concept) and across multiple vocabularies (e.g. an artist is associated to an art style). We applied all these relations as well as the basic artwork features in content-based recommendations and compared all of them in terms of usefulness. This investigation also enables us to look at the combined use of artwork features and semantic relations in sequence and derive user navigation patterns.

Chapter 6 defines the task of personalized recommendations and decomposes the task into a number of inference steps for ontology-based recommender systems, from a perspective of knowledge engineering. We proposed a hybrid approach combining both explicit and implicit recommendations. The explicit relations include

---

[8]http://www.getty.edu/research/conducting research/vocabularies/
[9]http://www.Iconclass.nl/libertas/ic?style=index.xsl

artworks features and semantic relations with preliminary weights which are derived from the evaluation in Chapter 5. The implicit relations are built between art concepts based on instance-based ontology matching.

Chapter 7 gives an example of reusing user interaction data generated by one application into another one for providing cross-application recommendations. In this example, user tagging about cultural events, gathered by iCITY[10], is used to enrich the user model for generating content-based recommendations in the CHIP Art Recommender. To realize full tagging interoperability, we investigated the problems that arise in mapping user tags to domain ontologies, and proposed additional mechanisms, such as the use of SKOS matching operators to deal with the possible mis-alignment of tags and domain-specific ontologies.

We summarized to what extent the problem statement and each of the research questions are answered in Chapter 8. We also discussed a number of limitations in our research and looked ahead at what may follow as future work.

---

[10]http://iCITY.di.unito.it/dsa-dev/

Samenvatting

Met het opkomen van Web 2.0, zijn de verschillende instituten en organizaties begonnen hun eertijds geïsoleerde en heterogene collecties open te stellen ten einde het publiek optimaal toegang te verschaffen. Semantische Web technologieën zijn daarbij instrumenteel voor de integratie van rijke collecties van metadata en het definieëren van ontologieën die de variëteit aan representatie schemata omvatten en inconsistenties tussen de verschillende vocabulaires oplossen. Met het oog op de grote verscheidenheid aan metadata met een complexe semantische structuur is het van groot belang een zorgvuldige selectie te maken van de informatie die aan bezoekers wordt gepresenteerd. In deze context werkt het CHIP[11] project (Cultural Heritage Information Personalization), vanaf 2005, gesponsord door NWO (Nederlandse Organisatie voor Wetenschappelijk Onderzoek), als onderdeel van het Nederlandse CATCH[12] programma (Continuous Access to Cultural Heritage). CHIP is een samenwerkingsproject, met als partners het Rijksmuseum Amsterdam[13], de Technische Universiteit Eindhoven[14] en het Telematica Instituut[15].

De probleemstelling die aan het onderzoek dat in dit proefschrift beschreven wordt ten grondslag ligt luidt als volgt: *Kunnen we bezoekers ondersteunen met gepersonaliseerde toegang tot de semantisch verrijkte collecties*? Ten einde deze vraag te bestuderen en te beantwoorden hebben we gekozen voor cultureel erfgoed (en in het bijzonder musea) als domein van toepassing, en dient de semantisch rijke achtergrondkennis van de museum collecties als uitgangspunt voor ons onderzoek. Met die kennis als basis, hebben we daarenboven gebruik gemaakt van gebruikersmodellerings- en aanbevelings-technieken teneinde gepersonaliseerde diensten te kunnen ontwikkelen voor museumbezoekers. De balangrijkste onderzoeksbijdragen van dit proefschrift zijn, samengevat: (i) we hebben een interac-

---

[11]http://www.chip-project.org/
[12]http://www.nwo.nl/catch
[13]http://www.rijksmuseum.nl/
[14]http://w3.tue.nl/
[15]http://www.novay.nl/en/

tieve *rating*-dialoog voor kunstwerken en concepten ontwikkeld, om snel een CHIP gebruikersprofiel te kunnen instantiëren, dat gerealizeerd is als een specializatie van FOAF[16] en tevens afgebeeld kan worden op een bestaande *event model* ontologie SEM[17]; (ii) we hebben een hybride aanbevelingsalgoritme voorgesteld, waarin zowel expliciete als impliciete relaties die bestaan in de semantische structuur van de collectie gecombineerd worden. Ten behoeve van de presentatie, voor de eindgebruikers, hebben we drie *tools* ontwikkeld: *Art Recommender*, *Tour Wizard* en de *Mobile Tour Guide*. Met een *user-centered* ontwerp cyclus als leidraad, hebben we een reeks evaluaties met museum bezoekers uitgevoerd om de effectiviteit van aanbevelingen die voortkomen uit de *rating*-dialoog te testen, alsook de verschillende manieren waarop een gebruikers-profiel (*user model*) geconstrueerd kan worden en de voorspellings-accuraatheid van het hybride aanbevelingsalgoritme.

In hoofdstuk 1 worden de onderzoeksvragen en onze aanpak besproken, alsook de indeling van dit proefschrift.

In hoofdstuk 2 wordt een overzicht gegeven van het werk dat ter voorbereiding dient van het latere werk. Dit overzicht omvat (i) de semantische verrijking van de Rijksmuseum collectie, en de afbeelding daarvan op drie Getty vocabulaires[18] (ULAN, AAT, TGN) en de Iconclass thesaurus[19]; (ii) de minimale gebruikersprofiel (*user model*) ontologie, gedefiniëerd als een specialisatie van FOAF, waarin in eerste instantie alleen *ratings* opgeslagen worden; (iii) de eerste implementatie van het inhoud-gebaseerde (*content-based*) aanbevelingsalgoritme van de CHIP *Art Recommender*.

In hoofdstuk 3 worden de twee andere tools gepresenteerd: de *Tour Wizard* en de *Mobile Tour Guide*. Op basis van gebruikers-*ratings*, geeft de web-gebaseerde *Tour Wizard* aanbevelingen over mogelijke museum *tours* bestaande uit aanbevolen kunstwerken, die momenteel voorhanden zijn in de museum expositieruimte. De *Mobile Tour Guide* vertaalt deze aanbevelingen naar een formaat geschikt voor mobiele apparaten (bv. een PDA) die gebruikt kunnen worden in de (physieke) museum ruimte. Om te kunnen reageren op interactie door de gebruikers hebben we een vertaling gemaakt van het gebruikersprofiel zoals opgeslagen in RDF naar een XML formaat dat leesbaar is voor het mobiele apparaat, zodat een dynamische synchronizatie mogelijk is tussen de *online* en *on-site* gebruikersprofielen.

In hoofdstuk 4 wordt de tweede generatie van de *Mobile Tour Guide* gepresenteerd, waarin een *real time* routeringssysteem is opgenomen, geschikt voor verschillende apparaten (waaronder de iPod). In vergelijk met de eerste generatie, kan deze versie museum tours aanpassen op basis van de *rating* van kunstwerken en concepten door de gebruiker, de locatie waar de gebruiker zich in het (physieke) museum bevindt en de coordinaten van de kunstwerken en tentoonstellingszalen

---

[16]http://www.foaf-project.org/
[17]http://semanticweb.cs.vu.nl/2009/11/sem/
[18]www.getty.edu/research/conducting research/vocabularies/
[19]www.Iconclass.nl/libertas/ic?style=index.xsl

in het museum. Bovendien ondersteunen we in deze versie een afbeelding van het CHIP gebruikersprofiel naar een bestaande *event model* ontologie SEM. Behalve *ratings* laat dit model toe ook andere gebruikersactiviteiten op te slaan, zoals het volgen van een tour en het bekijken van kunstwerken.

In hoofdstuk 5 worden een aantal semantische relaties geïdentificeerd, zowel binnen een vocabulaire (bv. een concept heeft een generalizerend/specializerend concept) als over de grenzen van vocabulaires heen (bv. een kunstenaar is geassociëerd met een kunststijl). We hebben deze relaties in combinatie met elementaire kunstwerk eigenschappen toegepast in inhoud-gebaseerde aanbevelingen (*recommendations*) en de verschillende relaties met elkaar vergeleken in termen van bruikbaarheid (*usefulness*). Dit onderzoek geeft ons verder mede zicht op het gebruik van kenmerken van kunstwerken in combinatie met meervoudige semantische relaties, en de afleiding van typische patronen in gebruikersnavigatie.

In hoofdstuk 6 wordt de taak van gepersonaliseerde aanbevelingen nauwkeurig gedefiniëerd en deze taak vervolgens opgesplitst in een aantal inferentiestappen gegeneralizeerd naar ontologie-gebaseerde aanbevelingssystemen, vanuit een perspectief van *knowledge engineering*. We hebben een hybride aanpak voorgesteld, waarin zowel expliciete als impliciete aanbevelingen gecombineerd worden. Tot de expliciete relaties behoren kenmerken van kunstwerken alsook semantische relaties met initiële gewichten, afgeleid van de evaluatie beschreven in hoofdstuk 5. De impliciete relaties hebben betrekking op (kunst) concepten en zijn verkregen middels *instance-based ontology matching*.

In hoofdstuk 7 wordt een voorbeeld gegeven van het hergebruik van de uit een applicatie verkregen gebruikers interactie data in een andere applicatie, om op die wijze *cross-application* aanbevelingen te kunnen geven. In dit voorbeeld worden gebruikers *tags* over culturele *events*, verzameld door iCITY[20], gebruikt om het gebruikersprofiel van de CHIP *Art recommender* te verrijken. Om een volledige *tagging* interoperabiliteit te realizeren, hebben we de problemen die zich voordoen bij de afbeelding van gebruikers *tags* naar domein ontologieën nader onderzocht, en additionele mechanismes voorgesteld, zoals het gebruik van SKOS *matching* operatoren, om mogelijke *mis-alignments* tussen *tags* en domein-specifieke ontologiën op te vangen.

In hoofdstuk 8 is een samenvattende bespreking van de oorspronkelijke probleemstellingen opgenomen en de mate waarin de afzondere onderzoeksvragen beantwoord zijn.

---

[20]iCITY.di.unito.it/dsa-dev/

Curriculum Vitae

Yiwen Wang was born on January 9, 1980 in Dangtu, China. After finishing applied science in 1998 at the high school affiliated to Fudan University, she started studying at the East China University of Science and Technology, where she received her bachelor degree in computer science in 2002. Two years later, she moved to the Netherlands, to continue studying at the Free University Amsterdam and obtained her master degree with specialization in multimedia and culture in 2006.

Between 2006 and 2010 Yiwen Wang was a PhD candidate at the Eindhoven University of Technology, in the department of mathematics and computer science, under daily supervision of Dr. Lora Aroyo. During these four years, she worked within the CHIP project, of which the results are presented in this dissertation.

Her research interests include personalization, user modeling, recommender systems, semantic web technologies and multimedia applications in a wide cultural heritage domain.

SIKS Dissertatiereeks

1998

1999

2000

2001