

CHIP Demonstrator: Semantics-driven Recommendations and Museum Tour Generation

Lora Aroyo^{1,2}, Natalia Stash¹, Yiwen Wang¹,
Peter Gorgels³, and Lloyd Rutledge⁴

¹ Eindhoven University of Technology, Computer Science
{n.v.stash, y.wang}@tue.nl

² VU University Amsterdam, Computer Science l.m.aroyo@cs.vu.nl

³ Rijksmuseum Amsterdam p.gorgels@rijksmuseum.nl

⁴ Telematica Institute Lloyd.Rutledge@cwi.nl

1 Introduction

The main objective of the CHIP project is to demonstrate how Semantic Web technologies can be deployed to provide personalized access to digital museum collections. We illustrate our approach with the digital database ARIA of the Rijksmuseum Amsterdam⁵. For the semantic enrichment of the Rijksmuseum ARIA database we collaborated with the CATCH STITCH project⁶ to produce mappings to Iconclass⁷, and with the MultimediaN E-culture project⁸ to produce the RDF/OWL of the ARIA and Adlib databases. The main focus of CHIP is on exploring the potential of applying adaptation techniques to provide personalized experience for the museum visitors both on the Web site and in the museum. This resulted in three demonstrator components:

- Artwork Recommender - a Web-based rating dialog to build a user profile, based on semantics-driven recommendations.
- Tour Wizard - a Web-based tool using the user profile to generate automatically personalized museum tours for each user, and to (semi)-automatically generate various personalized routes through the digital Rijksmuseum collection.
- Mobile Tour - a PDA-based tool, which uses the results from the Tour Wizard and helps users navigate and discover artworks in the physical Rijksmuseum environment.

The online version of the CHIP demonstrator as well as a tutorial with a brief walk-through of the personalization functionality can be found at: <http://www.chip-project.org/demo/>.

⁵ <http://rijksmuseum.nl/aria/>

⁶ <http://www.cs.vu.nl/STITCH/>

⁷ <http://www.iconclass.nl/libertas/ic?style=index.xsl>

⁸ <http://e-culture.multimedien.nl/>

Further, we give a short introduction to the basic functionality of the Web-based parts of the CHIP demonstrator. Please note that the CHIP project collects feedback, on the functionality and usability of the demonstrator, on a regular basis from studies with museum visitors. Thus, the demonstrator changes over time as we are incorporating more functionalities and improvements to the interface.

2 Usage Scenario: You Rate - We Recommend

In fig. 1 we illustrates how we employ semantics in building user profiles and using them for generating recommendations to the user, as a way of guiding users through the museum collection. In the *Artwork Recommender*, the user rates an artwork and several properties:

- artwork *Night Watch* - 4 stars (i.e. ‘I like *Night Watch*’);
- creator property *Rembrandt* - 4 stars (i.e. ‘I like *Rembrandt*’);
- theme property *Landscape* - 4 stars (i.e. ‘I like *landscape*’);
- theme property *Self-portrait* - 1 star (i.e. ‘I hate *self-portrait*’).

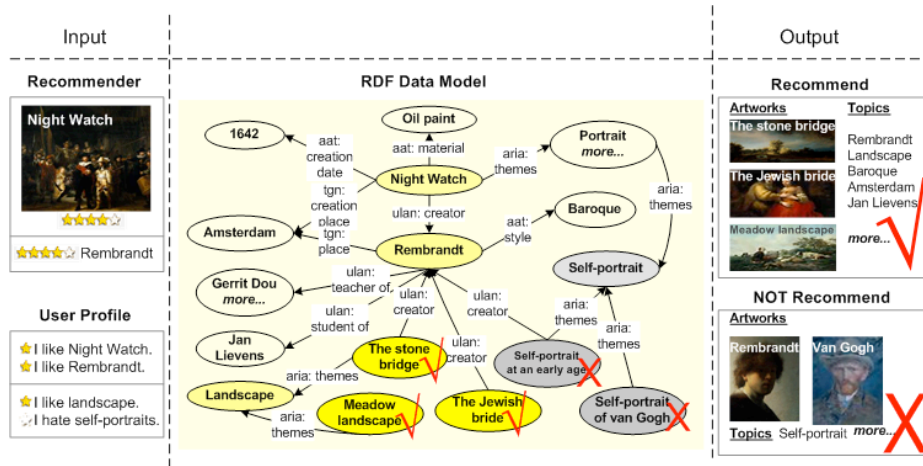


Fig. 1. Exploring semantic links in the Rijksmuseum Amsterdam collection

The User Profile stores the user’s ratings for generating recommendations. Now, let’s see how does the artwork and topic recommendation in CHIP work:

- Find all *Night Watch*-related properties, e.g. creator, creation place, creation year, material and themes
- Find all *Rembrandt*-related properties, e.g. style, teacher-of and student-of.

- Find all artworks with these properties, e.g. *The Jewish Bride* and *The Stone Bridge* by Rembrandt
- Include all artworks with property *Landscape*, e.g. *The Stone Bridge* and *Meadow Landscape*
- Exclude all artworks with property *Self-portrait*, e.g. *Self-portrait of Van Gogh* and *Self-portrait at an early age*

This results in two sets of recommendations:

- Result: recommend all artworks with the above positively rated properties. All recommended artworks are ordered by the number of matching properties, e.g. *The Stone Bridge* is the first one because it has both *Rembrandt* and *landscape*.
- Result: recommend all topics with the above positively rated properties, e.g. *Rembrandt*, *Landscape* and *Baroque*

Two more usage scenarios are give in the online tutorial.

3 CHIP Architecture

The demo is based on a Sesame [1] RDF store with SeRQL-based access to user modeling, recommendation and tour generation components. The tour generation component consists of two main parts: (1) a semantic-search facility for the user to search for themes or topics of a possible tour (e.g. a search for Rembrandt will result in a sub-set of Rembrandt artworks which are of interest to this user according to her user profile); and (2) my tours visualization on a historical timeline, museum map or as a list of artworks. In the latter the user can also manually create a tour by giving it a name and then continuing with the search option to find single artworks to include in the currently created tour holder. In fig. 2 we show the current CHIP architecture and its sub-components.

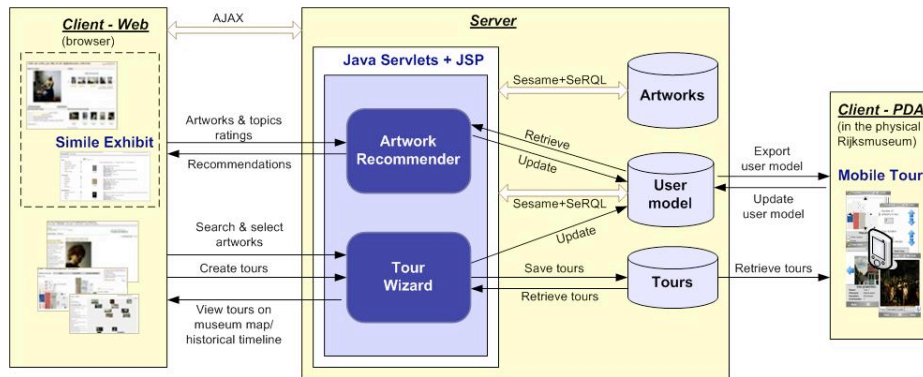


Fig. 2. CHIP Demonstrator Architecture

3.1 Rijksmuseum Amsterdam Collection and Shared Vocabularies

Currently, the demonstrator hosts four thesauri, namely the three Getty vocabularies⁹, i.e., the Art and Architecture Thesaurus (AAT), Union List of Artists Names (ULAN) and the Thesaurus of Geographical Names (TGN), as well as the subject classification Iconclass¹⁰. We use mappings to IconClass provided by the STITCH project [2]. We use the Getty thesauri conversion from their original XML format into an RDF/OWL representation done by the MultimediaN E-culture project [3,4]. The Getty thesauri are licensed¹¹. Following this approach we use mappings of the ARIA terminology to the AAT, ULAN, TGN and IconClass concepts. For example, the concepts for places in ARIA refer to location terms in TGN; styles in AAT are linked to artists in ULAN; birth places of artists in ULAN refer to location terms in TGN; subject themes in ARIA refer to subjects in IconClass; names of artists in ARIA refer to ULAN artists, etc. See fig. 3. We use the official ARIA collection of the Rijksmuseum in Amsterdam containing images of some 750 master pieces maintained at the Rijksmuseum Amsterdam website. However, we are now preparing for a migration to the main Oracle database of about 70,000 objects, extending the current RDF/OWL with not only more artworks but also shop, news and user comments items. The current client interface is developed in HTML+CSS and Ajax [5].

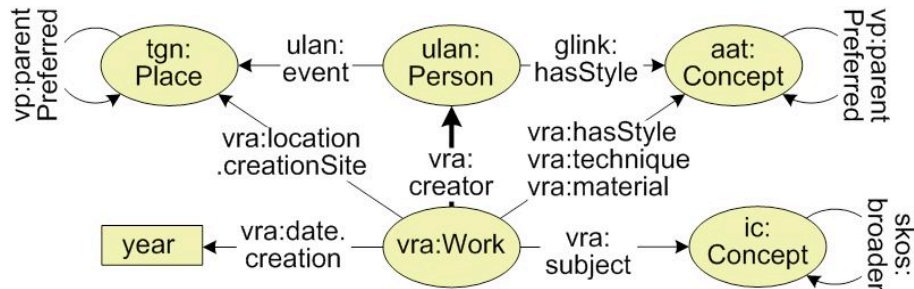


Fig. 3. CHIP Data Model and Vocabularies

4 Build Your Profile with Artwork Recommender

The user can start the exploration by first building a user profile in the Artwork Recommender component. This is driven by a rating dialog [6] for artworks from

⁹ http://www.getty.edu/research/conducting_research/vocabularies

¹⁰ <http://www.iconclass.nl/libertas/ic?style=index.xml>

¹¹ The partners in the project have acquired licenses for the thesauri. People using the demonstrator do not have access to the full thesauri sources, but can use them to annotate and/or search the collections.

the Rijksmuseum Amsterdam collection. The user can express her opinion using five stars, where the meaning of each star is shown when you hover the cursor over it. Next to a rating the user can indicate for each artwork whether to be used in further recommendations or not, by using the checkbox *Not interested in*. Rated artworks checked as *not interested in* will not influence the recommendations. The user can continue the process of rating artworks as long as she is satisfied with the state of the user profile shown on the right, or as long as the set of recommended artworks shown in the lower right part of the screen seem relevant. The user can skip rating artworks by pressing the *Next artwork* button. There is no average number of artworks the user needs to rate. In order to kick off the recommendations the user needs to give at least one positive (3, 4 or 5 stars) rating. The system would not be able to recommend artworks and topics based only on negative (1 or 2 stars) ratings.

Important here is that we recommend not only artworks but also topics (based on the semantic description of each artwork you have already rated). The user can provide her positive or negative feedback to each recommendation (both topics and artworks) by rating the empty set of stars associated with it. This would be recorded then in the user profile in order to increase the level of certainty for related properties and artworks.

The screenshot displays the 'CHIP Personalized Access to the Rijksmuseum Amsterdam Collection' interface. At the top, there are navigation links: 'Tutorial', 'Project Page', 'View Full Profile', and 'Log Out (total)'. Below this, the user is logged in as 'Art Recommender' and has access to a 'Tour Wizard'.

The main section is titled 'Your Profile' and features a large 'Artwork To Rate' on the left, which is a painting of a woman holding a platter with a man's head. Below the artwork are options to 'Rate this artwork' (with a star rating), 'Not interested', and 'Add to tour(s)'. Navigation buttons for 'Previous artwork' and 'Next artwork' are also present.

The 'Your Profile' section is divided into three columns:

- Topics I Like:**
 - Creation site (1): 5 stars, Haarlem
 - Material (2): 5 stars, Baroque; 5 stars, Oil paint
 - Themes (2): 5 stars, Food and Drink; 5 stars, Still lifes with objects
- Neutral Topics:**
 - Artists (1): 5 stars, Steen, Jan
 - Themes (2): 5 stars, Silver; 5 stars, Table- and kitchenware
- Topics I Dislike:**
 - Artists (1): 1 star, Heda, Willem; 1 star, Claesz
 - Themes (2): 1 star, Group; 1 star, Portraiture

Below the profile is a section for 'Artworks I Rated (6)', showing a row of five small artwork thumbnails with their respective star ratings and checkboxes.

At the bottom, there are two sections for recommendations:

- Topics We Recommend (51):** A list of topics with star ratings and question marks, including 'Brush technique', 'Still lifes with food', 'The Glory of the Golden Age', 'Amsterdam', and 'Light source'. A link 'See all recommended topics...' is provided.
- Artworks We Recommend (424):** A row of five artwork thumbnails with star ratings and question marks. A link 'See all recommended artworks ...' is provided.

Fig. 4. Screenshot of the CHIP Recommender

If the user is logged with a FOAF profile in the option to view the full user profile (right top) will show the user's personal and social network data. In the

full profile we also store the history interaction data about tours that this user has created and/or followed both on the Web and in the museum with the PDA-based mobile tour. Current investigations focus on including also social filtering [7] in order to include social aspects in recommendation and to optimize specific cold-start problems.

In this interactive way [8] the user has the opportunity to quickly get a feeling of the Rijksmuseum collection and give her opinion on the paintings presented. Based on a user rating the system finds artworks from the Rijksmuseum collection that are the most probable candidates for this user to like. See fig. 4 for a screenshot of the recommender user interface.

5 Create and Personalize Your Tours with Tour Wizard

The main rationale behind the CHIP demonstrator is based on the assumption that users more and more will spend time preparing their visits to museums and also reflecting on them after the visit. This is driven by the observation that users have an increasing choice of digitalized collections and related information. This leads us to our main goal to allow the **users to be their own curators**, e.g. selecting the artworks they want to see, influencing the order, the overall theme of the tour and the time to be spent in the museum. In order to realize this we maintain a common user profile for the user on the Web and on the mobile device, and in this way we keep track of users interaction history both in the virtual and in the physical museum.

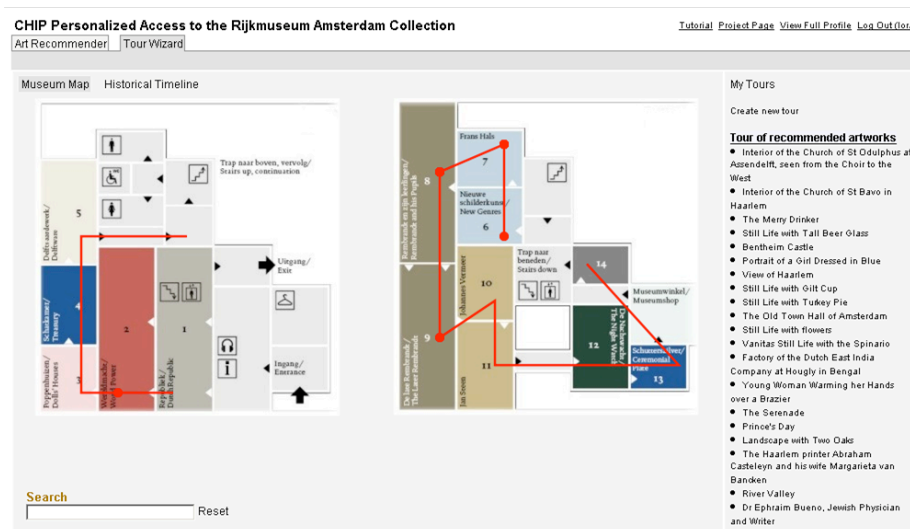


Fig. 5. Screenshot of *My Tours* on the museum map

In the Tour Wizard you have two main screens (tabs):

- **My Tours:** Here the user can see an overview of the current tours and can create a new tour if desired. She can (1) select a tour and then (2) search for a topic, artist or an artwork to include further new items in the currently selected tour. **Note** that if you go straight to the search option you will not be able to add new items to your tours; you will be able to only search for topics, artworks or artists and browse through the search results. Among the already existing tours we generate for each user a **Tour of Recommended Artworks**, which contains 20 Rijksmuseum master pieces selected according to user's current user profile. If further interaction with the demo appears, the user profile will be altered and in this way the tour of recommended artworks will be updated as well. Each tour can be viewed by the user as a list and on a **Museum Map** (see fig. 5) or on a **Historical Timeline** (see fig. 6).

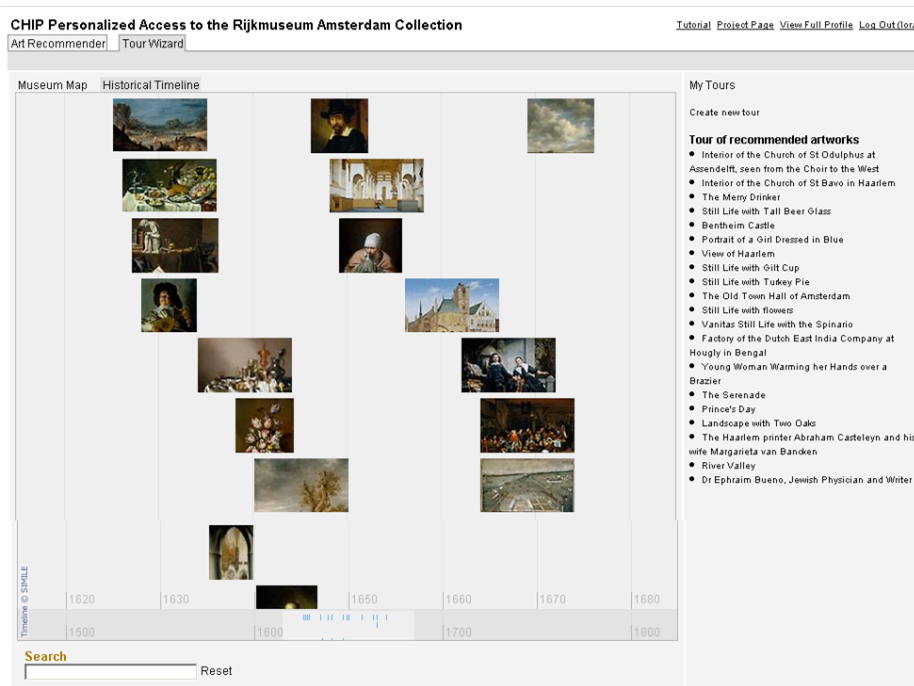


Fig. 6. Screenshot of *My Tours* on the historical timeline

- **Search Results:** Here the user can search the Rijksmuseum collection for topics, artists, locations, styles, artworks, etc. and browse the search results. Ideally these results will be altered (filtered with) according to the user profile, so that she will see the artworks related to her search query, which are of

relevance and interest to the user. However this option is still in work-in-progress, so it might not work properly yet. **Note** that if you first do search and then (while browsing the results) you want to add some of them to an existing tour, you will not be able to do so. Unfortunately, the current interface implementation restricts the user to **first select a tour and then perform a search for artworks to be included in this tour**. We are working on improving this interface limitation, so that you are able to add at any point of time an artwork from the search result (or even from your Artwork Recommender interface) to an existing tour.

Acknowledgements. The CHIP (Cultural Heritage Information Personalization) project is funded by the Dutch Science Foundation funded program CATCH (Continuous Access to Cultural Heritage) in the Netherlands.

References

1. Broekstra, J., Kampman, A., van Harmelen, F.: Sesame: A generic architecture for storing and querying rdf and rdf schema. (2002)
2. van Gendt, M., Isaac, A., van der Meij, L., Schlobach, S.: Semantic web techniques for multiple views on heterogeneous collections: A case study. In: ECDL'06: Proceedings of the International ECDL'06 Conference, Berlin / Heidelberg, Springer (2006)
3. van Assem, M., Menken, M., Schreiber, G., Wielemaker, J., Wielinga, B.: A method for converting thesauri to rdf/owl. In: ISWC '04: Proceedings of the International Semantic Web Conference, Berlin / Heidelberg, LNCS 3298 Springer (2006) 17–31
4. Schreiber, G., Amin, A., van Assem, M., de Boer, V., Hardman, L., Hildebrand, M., Hollink, L., Huang, Z., van Kersen, J., de Niet, M., Omelayenko, B., van Ossenbruggen, J., Siebes, R., Taekema, J., Wielemaker, J., Wielinga, B.: Multimedial e-culture demonstrator. In: ISWC '06: Proceedings of the International Semantic Web Conference, Berlin / Heidelberg, LNCS 4273 Springer (2006) 951–958
5. Garrett, J.J.: Ajax: A new approach to web applications. <http://www.adaptivepath.com/publications/essays/archives/000385.php> (February 2005)
6. Wang, Y., Aroyo, L., Stash, N., Rutledge, L.: Interactive user modeling for personalized access to museum collections: The rijksmuseum case study. In: UM'07: Proceedings of the International User Modeling Conference, Berlin / Heidelberg, Springer (2007)
7. van Setten, M.: Supporting people in finding information: Hybrid recommender systems and goal-based structuring. Telematica Instituut Fundamental Research Series, No.016 (TI/FRS/016). Universal Press. (2005)
8. Rutledge, L., Aroyo, L., Stash, N.: Determining user interests about museum collections. In: WWW'06: Proceedings of the International WWW'06 Conference (poster), Berlin / Heidelberg, Springer (2006)