Using Semantic Relations for Content-based Recommender Systems in Cultural Heritage

Yiwen Wang[1], Natalia Stash[1], Lora Aroyo[12], Laura Hollink[2], and Guus Schreiber[2]

[1] Eindhoven University of Technology, Computer Science
{y.wang,n.v.stash}@tue.nl
[2] VU University Amsterdam, Computer Science
{l.m.aroyo@cs.vu.nl,hollink,schreiber}

Abstract. Metadata vocabularies provide various semantic relations between concepts. For content-based recommender systems, these relations enable a wide range of concepts to be recommended. However, not all semantically related concepts are interesting for end users. In this paper, we identified a number of semantic relations, which are 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). Our goal is to investigate which semantic relations are useful for recommendations of art concepts and to look at the combined use of artwork features and semantic relations in sequence. These sequences of ratings allow us to derive some navigation patterns from users, which might enhance the accuracy of recommendations and be reused for other recommender systems in similar domains. We tested the CHIP demonstrator, called the Art Recommender with end users by recommending both semantically-related concepts and artworks features (e.g.*creator, material, subject*).

1 Introduction and Problem Statement

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[3], we have developed the CHIP Art Recommender: a content-based recommender system that recommend 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 Rijksmuseum InterActief (ARIA)[4] database [1] enables the opportunity to recommend a wide range of concepts via different semantic relations. These relations link concepts not only within one vocabulary (e.g. *teacher/studentOf, broader/narrower*), but also across two different

[3] http://www.rijksmuseum.nl
[4] http://www.rijksmuseum.nl/collectie/ontdekdecollectie

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 paper 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.

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 [2]. Peis et al [3] classified semantic recommender systems into three different types: (i) vocabulary or ontology based systems; (ii) trust network based systems constructed with FOAF[5]; and (iii) context-adaptable systems that use additional ontologies about e.g. the current time, place of the user. In this paper, 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 [4] recommends images based on semantic relations between selected images 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 [5], which extracts similar topics using the *broader/narrower* relations for recommendations. By knowing user

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

preferences, Blanco-Fernández [6] 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.

3 Metadata vocabularies and 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[6] and to the three Getty thesauri[7] (the Art and Architecture thesaurus (AAT), the Union List of Artists Names (ULAN) and the thesaurus of geographic Names (TGN)) [1] brings rich semantic structure to the Rijksmuseum collection and creates the opportunity to recommend a wide range of art concepts via various semantic relations. As shown in Fig. 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.

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 in 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

[6] http://www.iconclass.nl/libertas/ic?style=index.xsl
[7] http://www.getty.edu/research/conductingresearch/

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”.

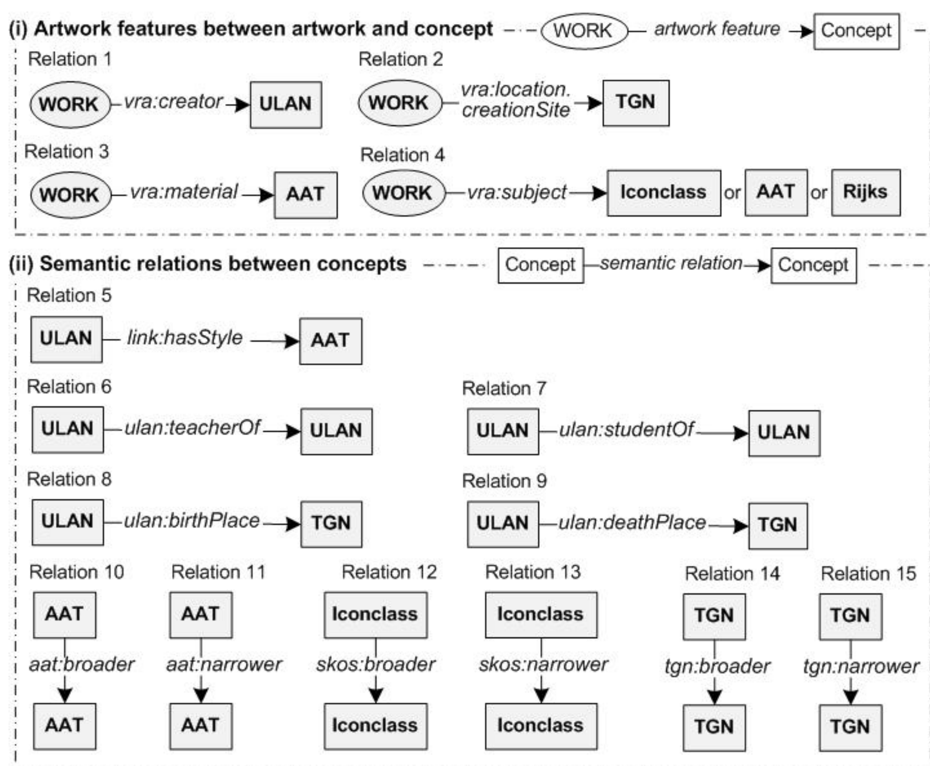

Fig. 1. Overview of artwork features and semantic relations based on the metadata vocabularies

4 Experiment

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.

4.1 Target System: The 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[8]. Considering artworks are recommended based on related/recommended art concepts, in order to get a clear insights, 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.

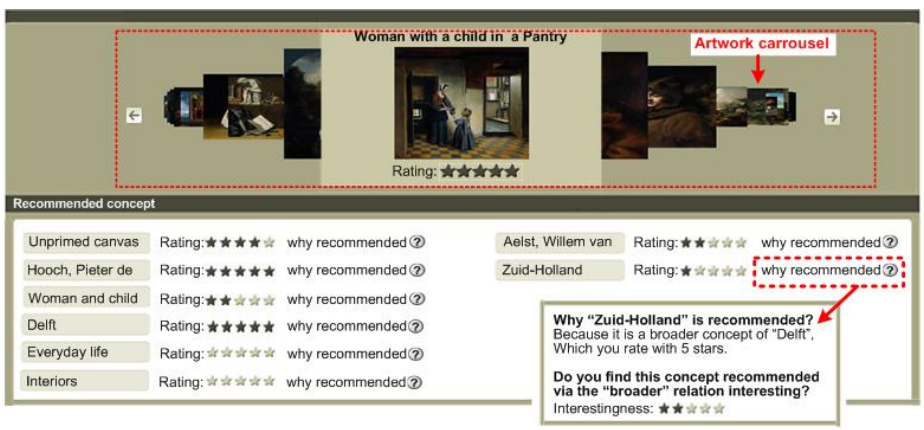

Fig. 2. User interface of the Art Recommender in the experiment

The user interface of the Art Recommender (see Fig. 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). 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.

[8] http://www.chip-project.org/demoUserStudy3/

In addition, in the "Why recommended" option (see Fig. 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).

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.

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. 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 [7, 8] and to measure the usefulness of semantic relations in query expansion for information retrieval systems [9–11]. *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 [8].

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 [7], relevance is subjective from an end user's standpoint and it often changes when the user gets explanations for recommendations. As Herlocker discussed in [8], 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 [7] (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 [8]. 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 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 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. 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.

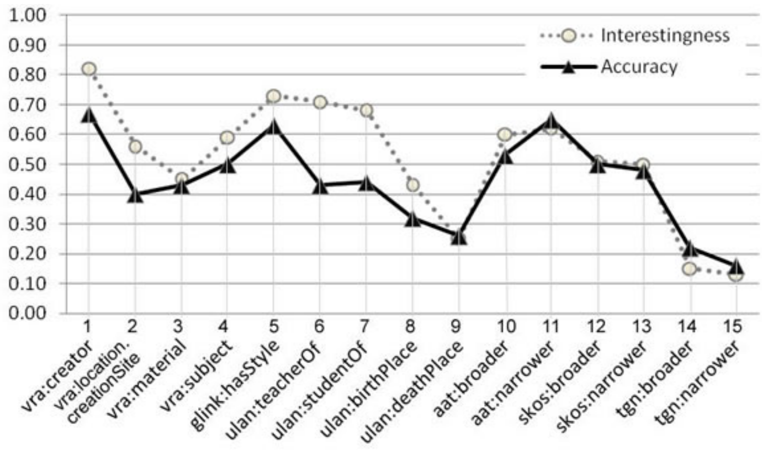

Fig. 3. Correlation between accuracy and interestingness

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 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 [9]. 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 [12], 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. 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. 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 artwork. The only exception is Step 4 in Pattern 2 (from teachers and

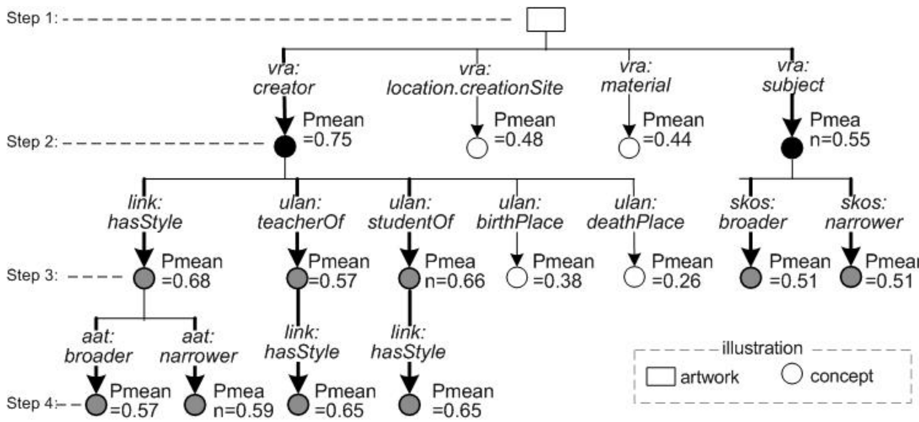

Fig. 4. Combining artwork features and semantic relations in sequence

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".

6 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 demonstrated 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 observation 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.

Generalizing, 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. In summary, 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 the preliminary results, 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 [13], Ruotsalo and Hyvönen [4], 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* [14], the times of user ratings of a particular artwork or concept, the semantic distance or similarity between two concepts by using latent semantic index (LSI) [15], etc.

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 [13] 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.

7 Acknowledgments

The CHIP project is funded by the Dutch Science Foundation funded program CATCH[9] (Continuous Access to Cultural Heritage). We would further like to thank all the participants in the experiment.

References

1. Wang, Y., Stash, N., Aroyo, L., Gorgels, P., Rutledge, L., Schreiber, G.: Recommendations based on semantically-enriched museum collections. Journal of Web Semantics (2008)
2. Berners-Lee, T., Hendler, J., Lassila, O.: The Semantic Web. (2001)
3. Peis, E., del Castillo, J.M.M., Delgado-Lopez, J.A.: Semantic recommender systems. analysis of the state of the topic. In: Hipertext. (2008)
4. Ruotsalo, T., Hyvönen, E.: A method for determining ontology-based semantic relevance. In: DEXA. (2007)
5. Adrian, B., Sauermann, L., Roth-Berghofer, T.: Contag: A semantic tag recommendation system. In: New Media Technology and Semantic Systems. (2007)
6. Blanco-Fernández, Y., Arias, J.J.P., et al.: Semantic reasoning: A path to new possibilities of personalization. In: ESWC. (2008) 720–735
7. Burke, R.: Hybrid recommender systems: Survey and experiments. Journal of User Model. User-Adapt. Interact. **12**(4) (2002) 331–370
8. Herlocker, J., Konstan, J., Borchers, A., Riedl, J.: Evaluating collaborative filtering recommender systems. ACM Transactions on Information Systems
9. Hollink, L., Schreiber, G., Wielinga, B.J.: Patterns of semantic relations to improve image content search. Journal of Web Semantics **5**(3) (2007) 195–203
10. Navigli, R., Velardi, P.: An analysis of ontology-based query expansion strategies. In: Machine Learning Conference. (2003)
11. Tudhope, D., Binding, C., Blocks, D., Cunliffe, D.: Query expansion via conceptual distance in thesaurus indexed collections. Journal of Documentation (2006)
12. Wang, Y., Aroyo, L., Stash, N., Rutledge, L.: Interactive user modeling for personalized access to museum collections: The rijksmuseum case study. In: User Modeling Conference. (2007)
13. Mobasher, B., Jin, X., Zhou, Y.: Semantically enhanced collaborative filtering on the web. Web Mining: From Web to Semantic Web (2004)
14. Baeza-Yates, R., Ribeiro-Neto, B. In:. Addison-Wesley, ACM Press (1999)
15. M.W. Berry, S.D., Brien, G.O.: Using linear algebra for intelligent information retrieval. SIAM Review

---

[9] http://www.nwo.nl/catch