Semantic Relations in Content-based Recommender Systems

Yiwen Wang
Natalia Stash
Eindhoven University of Technology
Computer Science
The Netherlands
{y.wang, n.v.stash}@tue.nl

Lora Aroyo
Laura Hollink
Guus Schreiber
VU University Amsterdam
Computer Science, The Netherlands
{l.m.aroyo, hollink, schreiber}@cs.vu.nl

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 both 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. 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*). Furthermore, we explored the use of combinations of artwork feature and semantic relations for recommendations.

Categories and Subject Descriptors

H.4.3 [**Information Systems Applications**]: Miscellaneous; H.3.3 [**Information Storage and Retrieval**]: Information Search and Retrieval; D.2.8 [**Software Engineering**]: Metrics—*performance measures*

General Terms

Experimentation, Measurement, Human Factors

Keywords

Semantic relations, metadata vocabularies, content-based recommender systems

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 a content-based recommender system, the "Art Recommender"[2]. Based on the user ratings of artworks, it recommends art-related concepts via the artwork features. For example, if a user gives the famous painting "Night watch" a high score, the system would recommend its *creator* "Rembrandt". Currently, the CHIP demonstrator works with the Amsterdam Rijksmuseum InterActief (ARIA)[3] database, containing images and descriptions. The mapping of metadata from ARIA to rich, well structured vocabularies [16] creates the opportunity to recommend a wider range of concepts via semantic relations. If the user likes artist "Rembrandt", it could recommend his *teacher* "Pieter Lastman", his *student* "Gerrit Dou".

Despite the promising examples, the use of semantic relations for recommendation also poses a problem, namely, not all related items are interesting for end users. Continuing the previous example of "Rembrandt", the system could also recommend his *death place* "Amsterdam", which might be not of interest for users. In the case of CHIP, this problem becomes even bigger because the semantic relations are not only within one vocabulary (e.g. *broader/narrower, teacher/studentOf*) but also across two different vocabularies (e.g.*hasStyle, birth/deathPlace*). In the experiment, we tested the Art Recommender with end users by applying both artwork features and semantic relations. Using artwork features as a baseline, we compared the results of recommended concepts via different semantic relations in terms of accuracy and interestingness.

The paper is organized as follows: Section 2 presents related work about recommender systems and the use

[1] http://www.rijksmuseum.nl
[2] http://www.chip-project.org/demo/
[3] http://www.rijksmuseum.nl/collectie/ontdekdecollectie

of semantic relations for recommendations. Section 3 gives a brief overview of the metadata vocabularies. In Section 4 we identify a number of semantic relations as well as artwork features. Section 5 explains the design of the experiment and Section 6 discusses the results. We conclude and discuss future work in Section 7.

2. RELATED WORK

Recommender systems provide advice to users about items they might be interested in [4]. As discussed in [4, 5], collaborative filtering and content-based recommendations are the most mature and widely implemented recommendation strategies.

Collaborative filtering recommender systems (e.g. MovieLens[4], Ringo[5]) assess the similarity between multiple users in order to recommend unseen items to a user. It works best for a user with many other users of similar taste, however, it tends to offer poor results when there are not enough user ratings (*cold-start* problem) and the number of items to rate far exceeds what a user could rate (*sparsity* problem) [5, 13].

Content-based recommender systems (e.g. Last.fm[6] and Amazon[7]) analyze item features/descriptions to identify items that are likely of interest to the user. It performs well when there are sufficient features for items. They also suffer from the *cold-start* problem, and additionally, they have the problem of *over-specialization* [5, 9], which means that the user is restricted to getting recommendations which bear a strong resemblance to those he already knows or defines in the user profile.

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]. As Peis discussed [10], semantic recommender systems can be classified as three different types: (i) vocabulary or ontology based systems; (ii) trust network based systems constructed with FOAF[8]; 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 can be used for recommendations.

Metadata vocabularies or domain ontologies are so far mainly used for content-based recommender systems, but not often used for collaborative filtering recommender systems. An exception is Mobasher's work [8], which demonstrated how to extract and populate items for similarity computation, by using structured semantic knowledge in collaborative filtering recommender systems. Regarding content-based recommender system using vocabularies, a lot of work has been done: the CULTURESAMPO portal [12] 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 for recommendations. A similar approach is adopted in the ConTag project [1], which extracts similar topics using the *broader/narrower* relations for recommendations. By knowing user preferences, Blanco-Fernández [3] inferred semantic associations between user preferences and relevant instances from the domain ontology in order to provide personalized recommendations of TV programs.

[4]http://www.movielens.org/
[5]http://www.ringo.com/
[6]http://www.last.fm
[7]http://www.amazon.com/
[8]Friend of A Friend: http://www.foaf-project.org/

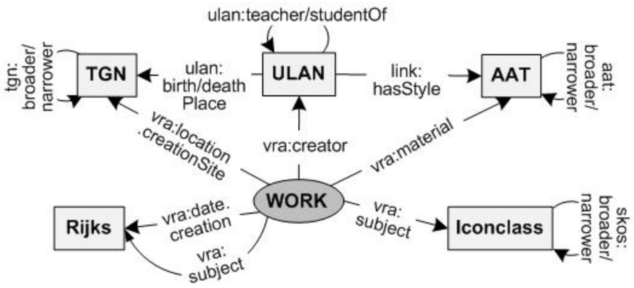

Figure 1: CHIP Metadata vocabularies

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 more rich semantic relations: not only the 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

The current CHIP Art Recommender works with the Rijksmuseum ARIA database. As discussed in [16], ARIA has two main problems: (i) inconsistent descriptions, and (ii) flat structure, which brings a severe obstacle for content-based recommendation inference.

For the semantic enrichment of heterogenous museum collections [16], we collaborated with the STITCH[9] project

[9]http://www.cs.vu.nl/STITCH/

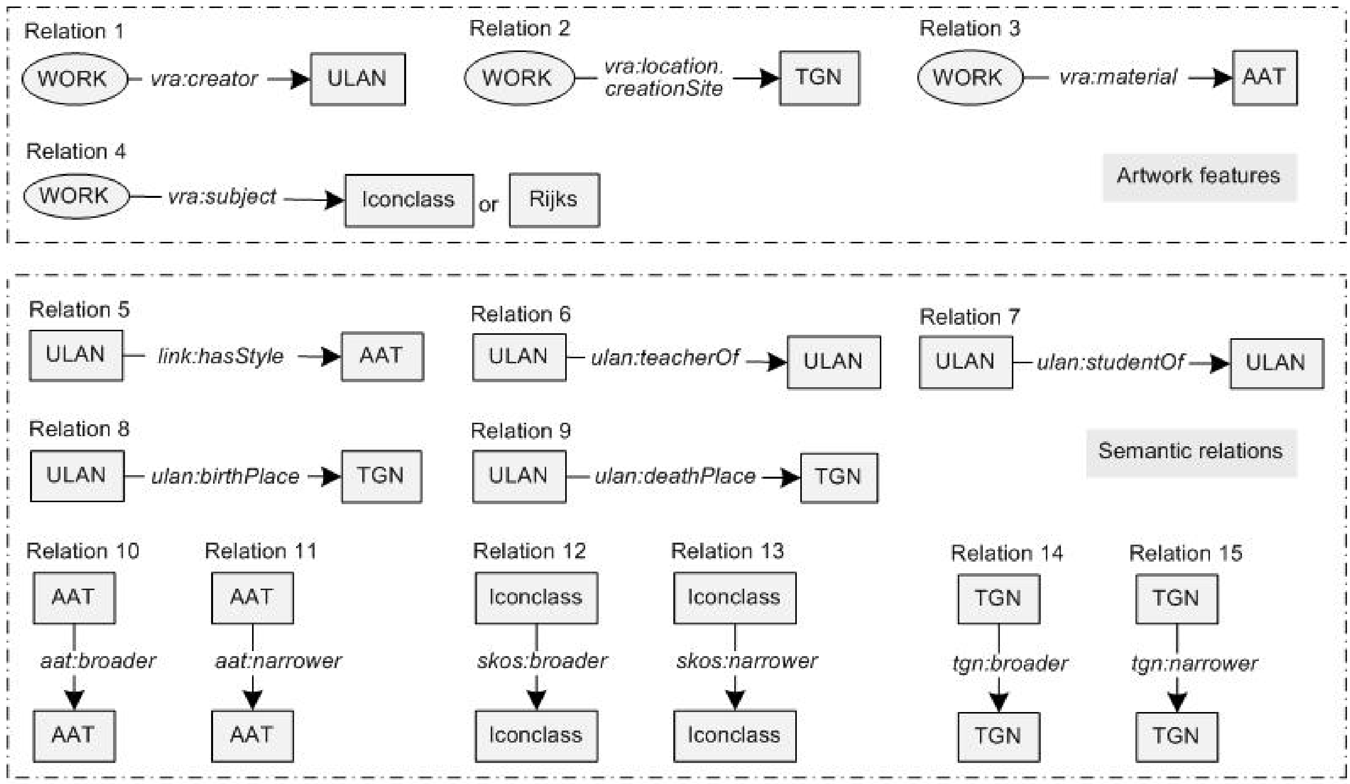

Figure 2: Overview of artwork features and semantic relations

to produce mappings to the Iconclass[10] vocabulary, and use the RDF/OWL representation (provided by the MultimediaN E-culture[11] project) to the three Getty[12] vocabulary, namely, the Art and Architecture thesauri (AAT), the Union List of Artists Names (ULAN) and the thesauri of geographic Names (TGN).

The mapping to metadata vocabularies bring rich semantic structure to the Rijksmuseum collection. Fig. 1 presents a top-level overview of the RDF Schema used in CHIP. VRA Core[13] is interpreted here to be a specialization of Dublin Core[14] for describing works of art and images of works of art. Creators are mapped to artists in the ULAN vocabulary. They are associated with other artists within ULAN mainly via *ulan:teacher/ studentOf* relations. Terms for creation sites in ARIA refer to the geographic locations in the TGN vocabulary. These geographic locations in TGN are connected with each other via the hierarchical relations *tgn:broader/narrower*. Terms of material and art styles from ARIA are mapped to concepts in the AAT vocabulary, which also use the hierarchical relations *aat:broader/ narrower*. ARIA subjects refer to concepts in the Iconclass vocabulary and the concepts are connected with each other using the *skos:broader/narrower* relations. The unmapped themes remain in ARIA (or called "Rijks" in all figures).

In addition, with the help of domain experts, we defined specific relations to map concepts across two different vocabularies, for example, artists in ULAN are linked to art styles in AAT via the *link:hasStyle* relation [16]. The Getty vocabularies provide relations (*ulan:birth/deathPlace*) between artists in ULAN and geographic locations in TGN where the artist was born in or died in.

4. SEMANTIC RELATIONS

The rich metadata vocabularies bring in total eleven semantic relations within one vocabulary and across two different vocabularies. As shown in Fig. 2, we list all these semantic relations (Relations 5-15) as well as four basic artwork features (Relations 1-4).

Relations 1-4 are artwork features, which have already been implemented in the official 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 Rijks.

[10]http://www.iconclass.nl/libertas/ic?style=index.xsl
[11]http://e-culture.multimedian.nl/
[12]http://www.getty.edu/research/conductingresearch/
[13]http://www.vraweb.org/resources/datastandards/vracore3/
[14]http://dublincore.org/

Figure 3: Interface of the Art Recommender in the experiment

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 in ULAN to an art styles in AAT (e.g. "Rembrandt" has an art style "Baroque"). Relations 6 and 7 are the *ulan:teacher/studentOf* relations 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, across the ULAN and TGN vocabularies, e.g. "Rembrandt" was born in "Leiden" and died in "Amsterdam".

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. EXPERIMENT

The goal of the experiment is to investigate which semantic relations are useful for recommendations and what is the added value of applying these semantic relations for content-based recommender systems in comparison with the basic artwork features. We tested the CHIP demonstrator, the "Art Recommender"[15] by applying both artwork features and semantic relations.

As shown in Fig. 3, the interface of the Art Recommender used in the experiment contains two parts: the top part is the artwork carrousel, which allows the user to browse artworks in the collection and give ratings to artworks with 1-5 stars (I hate it, I dislike it, neutral, I like it, and I like it very much). Based on the ratings of artworks with 4 or 5 stars, it will recommend concepts via artwork features, as shown in the bottom part of the interface. The user rates recommended concepts in the same 1-5 stars as feedback for the accuracy of recommendations. The list of recommended concepts will be dynamically updated based on the ratings of the last (current) artwork or concept.

In addition, the user could click on the "why recommended" icon next to each recommended concept. In the pop-up window (see Fig. 3), an explanation about which feature or relation is used for the recommendation is provided and the user is asked to give feedback about whether he/she finds the concept recommended via this feature or relation interesting (*interestingness*) on a scale of 1-5 stars (Not interesting at all, Not interesting, neutral, Interesting, Very interesting). In comparison with the direct ratings of concepts about "do you like it", the dimension of interestingness puts concepts back into the context, which helps user to understand the inference of recommendations by using a particular artwork feature or semantic relation.

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

At any point, the user could stop rating recommended concepts and go to select another artwork from the artwork carrousel. It will repeat the same process for each rated artwork.

5.1 Method

At the beginning of each session, participants were asked to fill out a questionnaire about: (i) 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 (from "artwork -*artwork features*-▶ concepts -*semantic relations*-▶ more concept"). Using the Art Recommender, users were asked to follow two steps, as shown in Fig. 3.

`Step 1 (Pre-task)`: to find an artwork that he/she likes from the artwork carrousel by giving 4 or 5 stars. As a baseline, it will produce the first set of recommended concepts via the basic artwork features.

`Step 2 (Main task)`: to rate recommended concepts. Based on the ratings of concepts with 4 or 5 stars, it will produce a new set of recommended concepts via semantic relations, which also allows users to rate. Besides ratings, at any point, the user could click on the "why recommended" icon for each recommended concept and give feedback on interestingness. As the main focus in the experiment, Step 2 gave us an insight in the performance of concepts recommended via semantic relations in comparison with 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.

5.2 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`*: it directly asks the user whether he/she likes this recommended concept, which is shown as "Ratings" in the Art Recommender in Fig. 3.
- *`Interestingness`*: by giving the explanations of why recommended, it asks the user whether he/she finds the concept recommended via this particular semantic relation or artwork feature interesting.

In order to get the ratio of relevant concepts to total recommended concepts, we used the metric of *precision* to measure both accuracy and interestingness. For accuracy, relevant concepts refer to recommended concepts that the user likes and for interestingness, relevant concepts refer to recommended concepts that the user finds interesting. As generally known, *precision* and *recall* are most popular metrics to evaluate recommender systems [5, 6] and also to measure the usefulness of semantic relations in query expansion for information retrieval systems [7, 11, 14]. As defined in [6], *precision* represents the probability that a recommended item is relevant and *recall* represents the probability that a relevant item will be recommended. However, in our case, determining the total number of relevant items for recall is difficult because the relevance is subjective from an end user's standpoint and it often changes when the user gets explanations for recommendations [5]. Thus in the experiment we only measure precision but not recall for both accuracy and interestingness. Below we give the formula of precision and explain the variables.

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

`Total Rec.Rated` is the total number of recommended concepts that have been rated by the user in terms of accuracy and interestingness respectively. To avoid the data sparsity problem [5] (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. We do not use the number of "Total Rec." which contains all recommended concepts with and without user ratings, because we do not have user feedback about those concepts without ratings [6]. 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.

`Correct Hits` is the total number of relevant concepts recommended in terms of accuracy and interestingness respectively. Both accuracy and interestingness have a rating scale of 1-5 stars. To measure the standard precision, we transformed the 5-star scale into a binary scale by converting 4 and 5 stars to "relevant" and 1-3 stars to "not relevant".

6. RESULTS

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

Table 1: Evaluation results for artwork 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.creationSite	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

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 his/her art interests better (79%)

Table 1 gives an overview for both 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. For accuracy and interestingness, they get an average score.

In order to see whether there are some correlations between accuracy and interestingness, in Fig. 5, we plotted these two dimensions for artworks features and semantic relations. Interestingly, there is a strong pos-

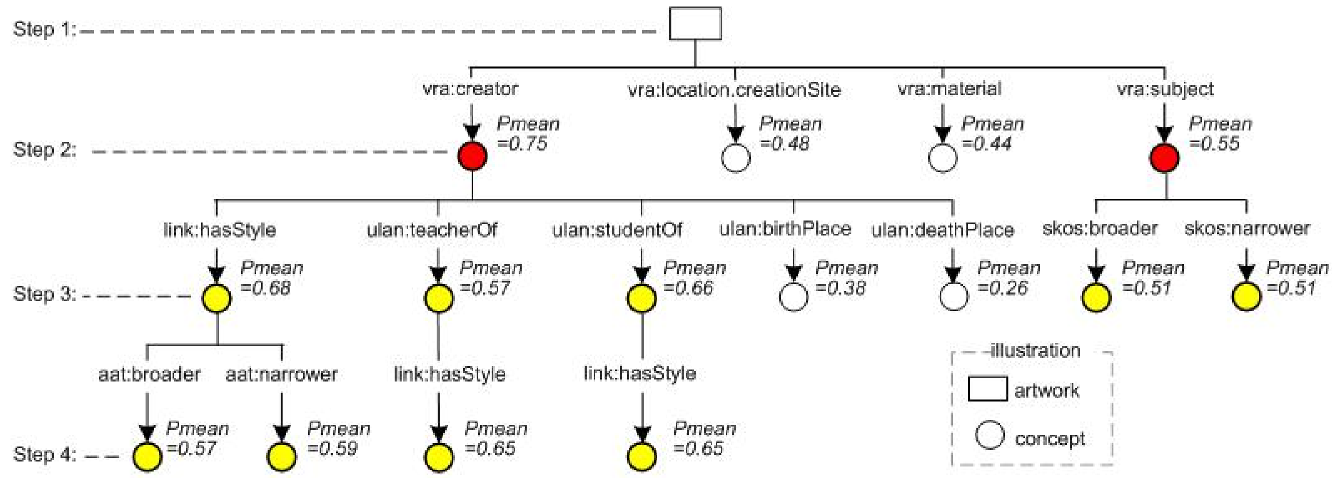

Figure 4: Combining artwork features and semantic relations in sequence

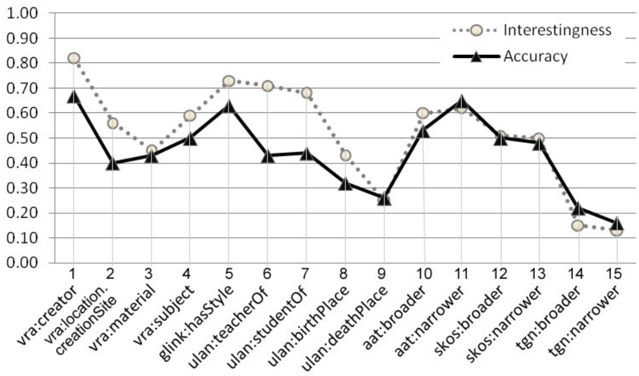

Figure 5: Correlation between Accuracy and Interestingness

itive 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 a bit low (0.43, 0.44), the interesting precisions for these relations are very high (0.71, 0.68). This gives a good explanation of how semantic relations can be used to partially solve the over-specialization problem (see Section 2) by recommending surprising or interesting items, although the recommendations are not always very accurate.

The setup of the experiment using the Art Recommender gives us an opportunity to look at the use of artwork features and semantic relations in sequence. A 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 [7]. These sequences of ratings allow us to examine the quality of recommendations based on sequences of semantic relations and artwork features.

We first remove all sequences for which we have less than 10 user ratings. From our previous user studies [15], 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 red and semantic relations (used in Step 3 and 4) that have a P_{mean} greater than 0.49 in yellow. Interestingly, we found three potentially useful patterns of combined artwork feature and semantic relations:

(i) artwork -► creator -► style -► broader/narrower styles
(ii) artwork -► creator -► teacher/student -► styles
(iii) 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 to the 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".

7. 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 patterns of combined features and relations with multiple intermediate concepts. These patterns can potentially be used to effectively recommend remotely linked concepts without asking the user's explicit feedback for the intermediate concepts.

In summary, we can conclude that the use of 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. Our conclusions about which semantic relations are most beneficial to recommendations can also be used for collaborative filtering recommender systems. Inspired by Mobasher's work [8] (see section 2), the beneficial semantic relations can be used to populate more related items in order to compute the similarity between users. This might be helpful to partially solve the cold-start and sparsity problems for collaborative filtering recommender systems.

Furthermore, we plan to investigate the weights for different semantic relations based on the user rating data we collected in the experiment. These weights can be used in the function for 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.

8. ACKNOWLEDGMENTS

The CHIP project is funded by the Dutch Science Foundation funded program CATCH[16] (Continuous Access to Cultural Heritage). We would further like to thank all the participants in the experiment.

9. REFERENCES

[1] B. Adrian, L. Sauermann, and T. Roth-Berghofer. Contag: A semantic tag recommendation system. In *New Media Technology and Semantic Systems*, 2007.

[2] T. Berners-Lee, J. Hendler, and O. Lassila. *The Semantic Web*. 2001.

[3] Y. Blanco-Fernández, J. J. P. Arias, and et al. Semantic reasoning: A path to new possibilities of personalization. In *ESWC*, pages 720–735, 2008.

[4] P. Brusilovsky, A. Kobsa, and e. Wolfgang Nejdl. *The Adaptive Web, Methods and Strategies of Web Personalization*. Springer, 2007.

[5] R. Burke. Hybrid recommender systems: Survey and experiments. *Journal of User Model. User-Adapt. Interact.*, 12(4):331–370, 2002.

[6] J. K.-A. B. Herlocker, J.L. and J. Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems.*

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

[8] B. Mobasher, X. Jin, and Y. Zhou. Semantically enhanced collaborative filtering on the web. *Web Mining: From Web to Semantic Web*, 2004.

[9] M. Pazzani and D. Billsus. Content-based recommendation systems. In *The Adaptive Web, Methods and Strategies of Web Personalization*, 2007.

[10] E. Peis, J. M. M. del Castillo, and J. A. Delgado-Lopez. Semantic recommender systems. analysis of the state of the topic. In *Hipertext*, 2008.

[11] R.Navigli and P. Velardi. An analysis of ontology-based query expansion strategies. In *Machine Learning Conference*, 2003.

[12] T. Ruotsalo and E. Hyvönen. A method for determining ontology-based semantic relevance. In *DEXA*, 2007.

[13] J. B. Schafer, D. Frankowski, J. L. Herlocker, and S. Sen. Collaborative filtering recommender systems. In *The Adaptive Web*, pages 291–324, 2007.

[14] D. Tudhope, C. Binding, D. Blocks, and D. Cunliffe. Query expansion via conceptual distance in thesaurus indexed collections. *Journal of Documentation*, 2006.

[15] Y. Wang, L. Aroyo, N. Stash, and L. Rutledge. Interactive user modeling for personalized access to museum collections: The rijksmuseum case study. In *User Modeling Conference*, 2007.

[16] Y. Wang, N. Stash, L. Aroyo, P. Gorgels, L. Rutledge, and G. Schreiber. Recommendations based on semantically-enriched museum collections. *Journal of Web Semantics*, 2008.

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