

EXPERT ADVICE AND REGRET FOR SERIAL RECOMMENDERS

Anton Eliëns
Faculty of Sciences
VU University Amsterdam
email: eliens@cs.vu.nl

Yiwen Wang
Dept. of Math. and Comp. Sc.
Eindhoven University of Technology
email: y.wang@tue.nl

KEYWORDS

recommender systems, expert advice, decision theory, personalization, guided tours, digital dossier, cultural heritage

ABSTRACT

In this paper we propose a tentative framework (R3) for adapting a sequence of predictions (guided tour) generated by what we call a *serial recommender*. The R3 framework (*rate*, *recommend*, *regret*) is applied to the construction of personalized guided tours, based on expert advice, in the domain of cultural heritage, in particular *digital dossiers* about contemporary art. Guided tours are in first instance obtained by tracking expert users. Our proposal is based on a variant of decision theory, that uses a regret function to measure the difference between a proposed decision and a finite collection of expert decisions. In our framework, personalization may then be seen as a minimization problem over a weighting scheme, expressing the relative importance of experts of which tours are available. Our aim in this paper is to arrive at a formalization of the recommendation of sequences (guided tours) that allows for adaptation to individual user preferences by a revision of the weight attached to a particular advice based on user feedback.

INTRODUCTION

Leaving all responsibility for interaction to the user is usually not a good choice, in particular when an information system contains complex, highly interrelated information. Despite the wealth of recommendation systems, it still seems to be an open problem how to generate a related collection of recommendations, that is an organized sequence of recommended items that may be used as a guided tour, for example an overview of artworks and related information from a museum collection.

In Eliens et al. (2006b), Wang et al. (2006), van Riel et al. (2006) we describe the *3D digital dossier* format, in which we presented the information of respectively the Dutch-Serbian artist Marina Abramovic¹ and the

Australian artist Jeffrey Shaw², contemporary artists with a variety of work, ranging from video to art installations. The *digital dossier* supports navigation using a concept graph and allows for presenting media-rich material, including 3D models of artwork installations. The digital dossiers have been implemented using X3D/VRML³ to allow for deployment on the web.

Recently we have explored *guided tours* in digital dossiers, van Riel et al. (2006), which actually automate user interaction, by mimicking user actions through events generated by a script. Although this provides an easy way to create guided tours, this does not solve the problem of what to select as elements in the guided tour, or how to personalize these tours in an intelligent manner.

In this paper, we discuss techniques from decision theory as a means to aid the construction of guided tours by consulting an advice function based on tracking the navigation behavior of expert users. We will also indicate how a similar advice function can be used for personalizing tours in cooperation with a recommender system for artworks, by altering the weight given to particular properties.

More in general, our aim is to arrive at a formalization of the mechanics underlying the recommendation of sequences (guided tours) that allows for adaptation to individual user preferences by a revision of the weight attached to a particular advice based on user feedback. Moreover we will give an indication how to generalize our approach to include the refinement of content-based ratings from which sequences are generated, by adapting weight attached to specific attributes of items featured in the guided tour. We opt for the phrase *serial recommender*, to stress on the one hand that the recommendation concerns sequences and not individual items, and on the other hand what one may call the compulsive nature of the recommendations, due to the fact that they are originally generated by experts. Mind that our approach has been primarily motivated by the need to support guided tours in digital dossiers. As we discuss in more detail in the paper, digital dossiers, and in particular the concept graph as a navigation paradigm, adhere to specific constraints that do not apply in general. As a consequence, it might be hard to

¹www.few.vu.nl/~dossier05

²www.few.vu.nl/~casus05

³www.web3d.org

generalize the approach to other domains where guided tours are useful. However, by including ratings based on content and an appropriate distance function between recommended items, it seems that the R3 framework introduced here is applicable to a wider class of (serial) recommenders.

structure The structure of this paper is as follows. First we will give a brief overview of recommender systems, after which we will give a short introduction to decision theory. Then we will describe the *abramovic* dossier, and discuss how techniques from decision theory can be applied to the construction of guided tours in digital dossiers, followed by a discussion of how to realize expert advice functions in digital dossiers. We will then illustrate how to apply decision theory for the personalization of tours in a more conventional cultural heritage application, sketch a formal model for (serial) recommender systems, introduce a distance function for item recommendations, and indicate how to deal with user feedback discrepancy. Finally, we will give our conclusions and indicate directions for future research.

RECOMMENDER SYSTEMS – BRIEF OVERVIEW

There is a great wealth of recommender systems, and a daunting number of techniques for producing recommendations, based on content, user behavior or social groups. See the AAAI 2004 Tutorial⁴ on recommender systems and techniques for an (extensive) overview.

In Van Setten (2005) a distinction is made between the following types of prediction techniques:

- social-based – dependent on (group) rating of item(s)
- information-based – dependent on features of item(s)
- hybrid methods – combining predictors

Social-based prediction techniques include collaborative filtering (CF), item-item filtering, popularity measures, etcetera. Information-based prediction techniques include information filtering, case-based reasoning and attribute or feature comparison. Finally, as hybridization techniques, Van Setten (2005) distinguishes between weighted combination, switching, mixed application and meta-approaches such as feature combination and cascaded application.

The approach we present in this paper, the R3 framework, has aspects of social-based as well as information-based methods and may be characterized as hybrid since it uses a weighting scheme to select between experts for advice.

For clarity, it is worthwhile to delineate briefly what we understand by the phrases *rate*, *recommend*, *regret*, and

how the R3 framework fits within the wider scope of recommendation techniques:

- *rating* – a value representing a user’s interest
- *recommendation* – item(s) that might be of interest to the user
- *regret* – a function to measure the accuracy of recommendations

In our approach, we (initially) proceed from the assumption that a rating is already present, and more in particular a rating that implies a sequential order on the presentation of a (limited) number of items. Later, however, we will explore how to relax this assumption and apply the R3 framework to sequences that are generated on the basis of content-based user preferences, to allow for an incremental adaptation of recommendations.


MATHEMATICAL PRELIMINARIES – DECISION THEORY

Before discussing how to realize guided tours in digital dossiers using user tracking and expert advice, we will give a very brief introduction to decision theory, more in particular a variant of decision theory introduced in Cesa-Bianchi and Lugosi (2006), that provides a mathematical foundation for our approach.

In classical prediction theory a prediction is a sequence of elements $x_1, x_2, \dots$ that results from a stationary stochastic process. The risk of the prediction is taken to be the expected value of the accumulated *loss* function, measuring the discrepancy between predicted values and actual outcomes. Cesa-Bianchi and Lugosi (2006) introduce a variant of prediction theory in which no assumption is made with respect to the nature of the source of predictions. Instead, the *forecaster* is considered to be an entity that gives a prediction for an element based on *advice* of one or more *experts*. These experts might be actual sequences stored in a database. The deviation of the forecaster with the actual outcome is measured using a *regret* function, and the prediction task may hence be formulated as minimizing the *regret* function by choosing the best expert for advice for each element of a prediction sequence.

For example, for the prediction of a bitstring of length n , the forecaster is a vector of n expert indices, that give advice for the bitvalue, 0 or 1 , in that position. In the general case, in which we have no information on the error rate of the experts’ advice, we may use a weighting factor $0 \leq \beta_i \leq 1$ for each expert i , to indicate the credibility of the experts’ advice. After each prediction, obtained by taking the majority decision of the experts, according to the weighting scheme, we may verify which experts fail to give the right advice, and decrease their weight, thus eliminating the influence of their advice in the long run.


⁴www.dfki.de/~jameson/aaai04-tutorial



THE ABRAMOVIC DOSSIER.

As a user interface for navigating the *abramovic* dossier, we created a concept graph, fig. 1, that represents arbitrary information structures in a hierarchical way. The concept graph allows the user to detect relations and search for information. Unlike the 3D cone tree, Robertson and MacKinlay (1991), where the complete hierarchical structure is presented, only a subset of the hierarchy is shown - three levels deep.

Presentation is an essential part of the digital dossier but is separated from navigation. The digital dossier contains different presentation facilities for 2D and 3D content. For 2D media content we need to be able to present video, images or textual information. This is implemented as a presentation gadget with three windows, fig 2. In each of the three windows the user can view either text, image or video content. The windows are positioned in such a way that the user can inspect the information simultaneously. In our experience, three views can be presented at the same time without much visual distortion.




usage scenario: When starting the dossier, it loads the concept graph that is used to navigate through the available information. In the center of the concept graph, a shining star is shown to illustrate the root of the information hierarchy, which is used as the start object.

When clicked, a star structure spreads and child objects appear surrounding the center star object.

Clicking on the *Interviews* node gives an overview of all interview fragments, then going back clicking on the information node *Artworks* and then on *China Ring* will bring the node for *China Ring* into focus. When clicking on the center node *China Ring*, a content presentation environment appears. which has three windows to present different types of information, grouped into the categories text, pictures and video. If desired, the user can focus on any window by using a zoom function. When the presentation of media content is finished, clicking on the close button will result in going back to the concept graph. Alternatively, the home function of the tool bar may be used to return directly to where we started: the original shining star.

An important feature of our digital dossier is the possibility to include 3D models of artwork installation. For example, in fig. 3, the installation *Terra della Dea Madre* is shown, which allows for interactive manipulation, such as rotation and positioning as a means to experiment with exhibition parameters in (virtual) space.



GUIDED TOURS IN DIGITAL DOSSIERS

In digital dossiers, we explored the use of guided tours as a means to present the information in a story-like way, relieving the user of the often cumbersome task to interact, van Riel et al. (2006b). Guided tours, in the digital dossier, may take one of the following forms:

- automated (viewpoint) navigation in virtual space,
- an animation explaining, for example, the construction of an artwork, or
- the (narrative) presentation of a sequence of concept nodes.

In practice, a guided tour may be constructed as a combination of these elements, interweaving, for example, the explanation of concepts, or biographic material of the artist, with the demonstration of the positioning of an artwork in an exhibition space.

A pre-condition for the construction of guided tours based on user tracking is that navigation consists of a small number of discrete steps. This excludes the construction of arbitrary guided tours in virtual space, since it is not immediately obvious how navigation in

virtual space may be properly discretized. In this case, as we will discuss later, a guided tour may be constructed using a programmed agent showing the user around.

For navigation in the concept graph, as well as for the activation of the media presentation gadget, the discretization pre-condition holds, and a guided tour may be composed from a finite number of discrete steps, reflecting the choice of the user for a particular node or interaction with the presentation gadget.

For example, in the *abramovic* dossier, the user has the option to go from the *Main* node to either *Artworks*, *Video Installations* or *Interviews*, and from there on further to any of the items under the chosen category. Tracking the actual sequences of choices of a user would suffice to create a guided tour, simply by re-playing all steps.

To obtain more interesting tours, we may track the navigation behavior of several experts for a particular task, for example retrieving information about the installation *Terra degli della Madre*. In case the experts disagree on a particular step in the tour, we may take the majority decision, and possibly correct this by adjusting the weight for one or more experts. When we have a database of tours from a number of experts, we may offer the user a choice of tours, and even allow to give priority to one or more of his/her favorite experts, again simply by adjusting the weighting scheme.

As a technical requirement, it must be possible to normalize interaction sequences, to eliminate the influence of short-cuts, and to allow for comparison between a collection of recordings. For the actual playback, as a guided tour, a decision mechanism is needed that finds the advice at each decision point, from each expert, to select the best step, according to a decision rule that takes the weighting scheme into account.

PERSONALIZATION BY EXPERT RATING

In a more mathematical way, we may state that for each node n we have a successor function $S(n)$, that lists the collection of nodes connected with n , which we may write as $S(n) = n_1, \dots, n_k$, where the suffix $i \leq k$ is an arbitrary integer index over the successor nodes. To take a history of navigation into account, we let $\bar{p}$ be a string of integers, representing the choices made, encoding the navigation path. So, for a node $n_{\bar{p}}$, with history $\bar{p}$, the collection of successor nodes is $S_{\bar{p}}(n) = n_{\bar{p}1}, \dots, n_{\bar{p}k}$. Now assume that we have a weight function w , that assigns to each expert e_i a weight $0 \leq \beta_i \leq 1$, indicating the relevance of expert i . Then for a particular node n we may assume to have an advice $\alpha_i = x$, with weight β_i and x in $S(n)$. If an expert has no advice for this node, we may simply assume its weight to be 0. For a collection of experts, the final advice will

be $\alpha(n) = \alpha_i(n)$ with weight β_i and $w(e_i) > w(e_j)$ for $i \neq j$. If no such advice $\alpha_i(n)$ exists, we may query the user to decide which expert has preference, and adapt the weights for the experts accordingly. This procedure can be easily generalized to nodes $n_{\bar{p}}$ with history $\bar{p}$.

To cope with possible shortcuts, for example when a choice is made for a node at three levels deep, we must normalize the path, by inserting the intermediate node, in order to allow for comparison between experts.

Now assume that we have expert navigation paths with cycles, for example $n_{\bar{p}} \rightarrow n_{\bar{p}1} \rightarrow n_{\bar{p}13}$, where actually $n_{\bar{p}} = n_{\bar{p}13}$, which happens when we return to the original node. In general such cycles should be eliminated, unless they can be regarded as an essential subtour. However, in this case, they could also be offered explicitly as a subtour, if they have length ≥ 4 . When offering guided tours for which several variants exist, we may allow the user to simply assign weights to each of the experts from which we have a tour, or allow for incrementally adjusting the weight of the experts, as feedback on the actual tour presented.

INTELLIGENT GUIDANCE — REALIZATION

Our aim is to arrive at a general framework for artist's digital dossiers, that provide intelligent guidance to both the expert user, responsible for the future re-installation of the work(s), and the interested layman, that wishes to get acquainted with a particular work or collection of works. In general, there are two techniques that we can apply to provide such guidance:

- filtering the information space according to the user's perspective, and
- intelligent agents, that (pro) actively aid the user in searching the information space.

Filtering the information space may be used to restrict the concept graph that defines the navigation structure, by stating assumptions with respect to the relevance of particular categories from a user's perspective.

Intelligent agents is an approach stemming from artificial intelligence which allows for providing guidance in a variety of ways, possibly even in an embodied form using a face or humanoid figure to give suggestions to the user on what interactions to perform, an approach that we will discuss later on.

For selecting the items to be presented in a guided tour, the most obvious way is to pre-define a sequence based on user profiles. Very likely this can be done in a more flexible way in a rule-based manner, applied to a template tour. More interesting, however, is to generate guided tours dynamically based on tracking actual user interaction of (expert) users, using techniques from prediction theory, as explained in the previous sections.




Fig. 4: Construction tool

A special case of a guided tour is the tool environment constructed for the *Revolution* installation of Jeffrey Shaw, which allows for experimenting with the (de-) construction of the installation, fig. 4, and exhibition parameters, fig. 5.




Fig. 5: (a) Light (b) Material

Tracking interaction with such 3D models is, given the limitations imposed by the tool environment, relatively simple, and can be used for creating a repository of navigation sequences. More difficult, however, is to find proper normalizations for these interactions, and so in this case we may possibly have to rely on expert weighting only.

agent technology In Hoorn et al. (2004) we have investigated the use of embodied agents in a digital dossier for the artist Marinus Boezem, fig. 6. To allow for a discrete mode of navigation we have used a map, displaying the interesting parts of the atelier, which contains locations where relevant information can be obtained, such as a filmprojector, for displaying interviews, a cabinet that contains biographical material and textual descriptions of the artworks, and an exhibition environment that displays (3D models of the) artworks. To construct a guided tour, we deployed a humanoid agent that shows the user around.




Fig 6. Overview Interview Agent

In a user evaluation test we found that humanoid agents were instrumental in providing information about the re-installation of artworks, but interestingly also that believability was positively affected by the degree of

realism of the agent, Van Vugt et al. (2006). However, in creating guided tours for the current generation of digital dossiers, using concept graphs for navigation instead of a spatial metaphor, we will not use humanoid agents. Our agent technology, however, can be used in a fruitful way.

In the I-GUARD⁵ project (Intelligent Guidance in Archives and Dossiers), we investigate how to realize *advice functions*, implemented using agent technology, Eliens et al. (2002), based on actual navigation paths obtained by tracking expert users, that offer the user at any navigation point a choice of continuations and/or a selection of guided tours, focussing on a topic of interest.

INCREMENTAL ADAPTATION OF RECOMMENDATIONS

In the CHIP⁶ project (Cultural Heritage Information Personalization), the aim is to develop a recommender system that generates a collection of artworks in accordance with the users' preferences based on the rating of a small sample of artworks. The properties on which the recommendation is based include *period*, *artist*, and *genre*. The recommender system will also be used to generate guided tours, where apart from the already mentioned properties the *location* (the proximity in the actual museum) will be taken into account.

Using a weighting scheme on the properties, that is a difference metric on the properties, a graph can be created, giving a prioritized accessibility relation between each artwork and a collection of related artworks. By changing the weight for one of the properties, for example *location*, in case the tour is generated for the actual museum, the priority ordering may be changed, resulting in a different tour.

In contrast to the successor function for nodes in the concept graph of the digital dossier, we may assume to have a weighted successor function $S_w(n) = (n_1, \omega_1), \dots, (n_k, \omega_k)$, with $\omega_i = w(n_i)$ the weight defined by the relevance of the node n_i , with respect to the attributes involved. In a similar way as for the digital dossier, user tracking may be deployed to incrementally change the weight of the arcs of the graph, reflecting the actual preference of the user when deviating from an existing guided tour. In the remainder of this paper we will give the outline of a recommender model supporting the incremental adaptation of preferences by user feedback.

⁵www.cs.vu.nl/~eliens/i-guard.html

⁶www.chip-project.org

SERIAL RECOMMENDER MODEL

Admittedly not the best way to do research, although common practice, we found a good starting point for modelling recommender systems, by googling on *serial recommender*, in a paper from Microsoft Research on privacy in distributed recommender systems, Oard et al. (2006). The model introduced in Oard et al. (2006), distinguishes between:

U = user
 I = item
 B = behavior
 R = recommendation
 F = feature

and allows for characterizing observations (from which implicit ratings can be derived) and recommendations, as follows:

- observations – $U \times I \times B$
- recommendations – $U \times I$

In a centralized approach the mapping $U \times I \times B \rightarrow U \times I$ provides recommendations from observations, either directly by applying the $U \times I \rightarrow I \times I$ mapping, or indirectly by the mapping $U \times I \rightarrow U \times U \rightarrow I \times I$, which uses an intermediate matrix (or product space) $U \times U$ indicating the (preference) relation between users or user-groups. Taken as a matrix, we may fill the entries with distance or weight values. Otherwise, when we use product spaces, we need to provide an additional mapping to the range of $[0,1]$, where distance can be taken as the dual of weight, that is $d = 1 - w$.

In a decentralized approach, Oard et al. (2006) argue that it is better to use the actual features of the items, and proceed from a mapping $I \times F \rightarrow U \times I \times R$. Updating preferences is then a matter of applying a $I \times B \rightarrow I \times F$ mapping, by analyzing which features are considered important.

For example, observing that a user spends a particular amount of time and gives a rating r , we may apply this rating to all features of the item, which will indirectly influence the rating of items with similar features.

B = [time = 20sec, rating = r]
 F = [artist = rembrandt, topic = portrait]
 R = [artist(rembrandt) = r , topic(portrait) = r]

Oard et al. (2006) observe that B and R need not to be standardized, however F must be a common or shared feature space to allow for the generalization of the rating of particular items to similar items.

With reference to the CHIP project, mentioned in the previous section, we may model a collection of artworks by (partially) enumerating their properties, as indicated below:

$A = [p_1, p_2, \dots]$
 where $p_k = [f_1 = v_1, f_2 = v_2, \dots]$

with as an example

$A_{nightwatch} = [\text{artist}=\text{rembrandt}, \text{topic}=\text{group}]$
 $A_{guernica} = [\text{artist}=\text{picasso}, \text{topic}=\text{group}]$

Then we can see how preferences may be shared among users, by taking into account the (preference) value adhered to artworks or individual properties, as illustrated in fig. 7.




Fig 7. Users, artworks and properties

As a note, to avoid misunderstanding, Picasso's Guernica is not part of the collection of the Rijksmuseum, and does as such not figure in the CHIP studies. The example is taken, however, to clarify some properties of metrics on art collections, to be discussed in the next section.

CONTENT METRICS

To measure similarity, in information retrieval commonly a distance measure is used. In mathematical terms a distance function $d : X \rightarrow [0,1]$ is distance measure if:

$$\begin{aligned} d(x, y) &= d(y, x) \\ d(x, y) &\leq d(x, z) + d(z, y) \\ d(x, x) &= 0 \end{aligned}$$

From an abstract perspective, measuring the distance between artworks, grouped according to some preference criterium, may give insight in along which dimension the grouping is done, or in other words what attributes have preference over others. When we consider the artworks

a_1 = [artist = rembrandt, topic = self-portrait]
 a_2 = [artist = rembrandt, name = nightwatch]
 a_3 = [artist = picasso, topic = self-portrait]
 a_4 = [artist = picasso, name = guernica]

we may, in an abstract fashion, deduce that if $d(a_1, a_2) < d(a_1, a_3)$ then $r(topic) < r(artist)$, however if $d(a_1, a_3) < d(a_1, a_2)$ the reverse is true, that is then $r(artist) < r(topic)$. Somehow, it seems unlikely that a_2 and a_4 will be grouped together, since even though their topic may be considered to be related, the aesthetic impact of these works is quite different, where *selfportraits* as a genre practiced over the centuries indeed seem to form a 'logical' category. Note that we may also express this as $w(artist) < w(topic)$ if we choose to apply weights to existing ratings, and then use the observation that if $d(a_1, a_3) < d(a_1, a_2)$ then $w(artist) < w(topic)$ to generate a guided tour in which a_3 precedes a_2 .

For serial recommenders, that provide the user with a sequence of items $\dots, s_{n-1}, s_n, \dots$, and for s_n possibly alternatives $a_1, a_2, \dots$, we may adapt the (implied) preference of the user, when the user chooses to select alternative a_k instead of accepting s_n as provided by the recommender, to adjust the weight of the items involved, or features thereof, by taking into account an additional constraint on the distance measure. Differently put, when we denote by $s_{n-1} \mapsto s_n/[a_1, a_2, \dots]$ the presentation of item s_n with as possible alternatives $a_1, a_2, \dots$, we know that $d(s_{n-1}, a_k) < d(s_{n-1}, s_n)$ for some k , if the user chooses for a_k . In other words, from observation B_n we can deduce R_n :

$$\begin{aligned} B_n &= [\text{time} = 20\text{sec}, \text{forward} = a_k] \\ F_n &= [\text{artist} = \text{rembrandt}, \text{topic} = \text{portrait}] \\ R_n &= [d(s_n, a_k) < d(s_n, s_{n+1})] \end{aligned}$$

leaving, at this moment, the feature vector F_n unaffected. Together, the collection of recommendations, or more properly revisions R_i over a sequence S , can be solved as a system of linear equations to adapt or revise the (original) ratings. Hence, we might be tempted to speak of the $R4$ framework, *rate*, *recommend*, *regret*, *revise*. However, we prefer to take into account the cyclic/incremental nature of recommending, which allows us to identify revision with rating.

MEASURES FOR FEEDBACK DISCREPANCY

So far, we have not indicated how to process user feedback, given during the presentation of a guided tour, which in the simple case merely consists of selecting a possible alternative. Before looking in more detail at how to process user feedback, let us consider the dimensions involved in the rating of items, determining the eventual recommendation of these or similar items. In outline, the dimensions involved in rating are:

- positive vs negative
- individual vs community/collaborative
- feature-based vs item-based

Surprisingly, in Wang et al. (2007) we found that negative ratings of artworks had no predictive value for an explicit rating of (preferences for) the categories and properties of artworks. Leaving the dimension *individual vs community/collaborative* aside, since this falls outside of the scope of this paper, we face the question of how to revise feature ratings on the basis of preferences stated for items, which occurs (implicitly) when the user selects an alternative for an item presented in a guided tour, from a finite collection of alternatives.

A very straightforward way is to ask explicitly what properties influence the decision. More precisely, we may ask the user why a particular alternative is selected, and let the user indicate what s/he likes about the selected alternative and dislikes about the item presented by the recommender. It is our expectation, which must however yet be verified, that negative preferences do have an impact on the explicit characterization of the (positive and negative) preferences for general artwork categories and properties, since presenting a guided tour, as an organized collection of items, is in some sense more directly related to user goals (or educational targets) than the presentation of an unorganized collection of individual items. Cf. Van Setten (2005).

So let's look at $s_{n-1} \mapsto s_n/[a_1, a_2, \dots]$ expressing alternative selection options $a_1, a_2, \dots$ at s_n in sequence $S = \dots, s_{n-1}, s_n$. We may distinguish between the following interpretations, or revisions:

- neutral interpretation – use $d(s_n, a_k) < d(s_n, s_{n+1})$
- positive interpretation – increase $w(\text{feature}(a_k))$
- negative interpretation – decrease $w(\text{feature}(s_{n+1}))$

How to actually deal with the revision of weights for individual features is, again, beyond the scope of this paper. We refer however to Eliens (2000), where we used feature vectors to find (dis)similarity between musical fragments, and to Schmidt et al. (1999), on which our previous work was based, where a feature grammar is introduced that characterizes an object or item as a hierarchical structure, that may be used to access and manipulate the component-attributes of an item.

CONCLUSIONS

In this paper we have shown how to adapt guided tours based on tracking expert users by modifying the weights attached to the experts that contributed to the construction of this tour. The application of these techniques requires that choices are discrete and hence do not apply to arbitrary navigation in virtual environments, unless we find proper ways to encode such navigation as a small finite collection of discrete steps. Also in the discrete case, however, we must be able to normalize navigation paths, in order to compare and weigh the contribution of the experts involved.

We have generalized our approach to a wider class of serial recommenders, and indicated how to apply the revision of ratings in an incremental fashion to adapt an existing tour to personal preferences, reflecting the actual navigation behavior of users.

As future work, we wish to investigate how we can use both positive and negative user feedback to revise and refine ratings for the actual features involved. Additionally, we would like to study features not directly related to artworks, but for example to group norms or personal likes and dislikes.

ACKNOWLEDGEMENT(S) We thank the reviewer who valued our original submission, which was meant to be a short paper, so positively that we were invited to write an extended paper. We took the challenge and hopefully not to the disappointment of our anonymous reviewer.

AUTHOR BIOGRAPHY

ANTON ELIENS is coordinator of the multimedia curriculum at the computer science department of the Faculty of Sciences of the Vrije Universiteit Amsterdam. He has written numerous papers and published books on distributed logic programming and object oriented software engineering.

YIWEN WANG is Ph.D. researcher in the CHIP project from the Technische Universiteit Eindhoven. She does her research at the Rijksmuseum Amsterdam, together with the other members of the CHIP team.

REFERENCES

- Aroyo, L., Rutledge, L., Brussee, R., de Bra, P., Gorgels, P., Stash N., Veenstra, M. (2005), *Personalized Presentation and Navigation of Cultural Heritage Content*, Multimedia and Expo, ICME 2005. IEEE International Conference, 2005
- Cesa-Bianchi N. and Lugosi G. (2006), *Prediction, Learning, and Games*, Cambridge University Press
- Eliens A. (2000), *Principles of Object-Oriented Software Development*, Addison-Wesley Longman, 2nd edn.
- Eliens A., Huang Z., and Visser C. (2002), *A platform for Embodied Conversational Agents based on Distributed Logic Programming*, In: *Proc. AAMAS Workshop – Embodied conversational agents - let's specify and evaluate them!*, Bologna 17/7/2002
- Eliens A., van Riel C., Wang Y. (2006), *Navigating media-rich information spaces using concept graphs – the abramovic dossier*, In *Proc. InSciT2006*, 25-28 Oct. 2006, Merida, Spain
- Hoorn J., Eliens A., Huang Z., van Vugt H.C., Konijn E.A., Visser C.T. (2004). *Agents with character: Evaluation of empathic agents in digital dossiers*, In: *Proc. AAMAS 2004*, Emphatic Agents, New York 19 July - 23 July, 2004
- Oard D.W., Leuski A. and Stubblebine S. (2006), *Protecting the Privacy of Observable Behavior in Distributed Recommender Systems*, In: *Proc. SIGIR 2003*
- Riel C. van, Eliens A., Wang Y. (2006), *Exploration and guidance in media-rich information spaces: the implementation and realization of guided tours in digital dossiers*, In: *Proc. InSciT2006*, 25-28 Oct. 2006, Merida, Spain
- Riel C. van, Wang Y. and Eliens A. (2006b), *Concept map as visual interface in 3D Digital Dossiers: implementation and realization of the Music Dossier*, In: *Proc. CMC2006*, Costa Rica, Sept 5-8 2006
- Robertson G.G. and MacKinlay J.D. (1991), *Cone trees: animated 3D visualizations of hierarchical information*, In: *Proc. of the SIGCHI 1991*, pp. 189-194
- Schmidt, A.R. Windhouwer M.A., Kersten M.L, (1999). *Indexing real-world data using semi-structured documents*, CWI Report INS-R9902
- Subrahmanian V.S. (1998), *Principles of Multimedia Databases*, Morgan Kaufmann
- Van Setten M. (2005), *Supporting People in Finding Information – Hybrid recommender Systems and Goal-based Structuring*, Ph.D. Thesis, Telematica Institute Netherlands
- Van Vugt, H. C., Konijn, E. A., Hoorn, J. F., Keur, I., & Eliens, A. (2006). *Realism is not all! User Engagement with Task-Related Interface Characters*, *Interacting with Computers*, 2006
- Wang Y., Eliens A., van Riel C. (2006), *Content-oriented presentation and personalized interface of cultural heritage in digital dossiers*, In: *Proc. InSciT2006*, 25-28 Oct. 2006, Merida, Spain
- Wang Y., Aroyo L., Stash N., Rutledge L. (2007), *Interactive User Modeling for Personalized Access to Museum Collections – The Rijksmuseum Case Study*, accepted for *11th Conf. on User Modeling, UM 2007*, June 25-29, Corfu, Greece