

Using Interest and Transition Models to Predict Visitor Locations in Museums

Fabian Bohnert^{a,*}, Ingrid Zukerman^a,
Shlomo Berkovsky^{a,b,**}, Timothy Baldwin^b,
and Liz Sonenberg^b

^a Monash University,
Clayton, VIC 3800, Australia
E-mail: {fabianb,ingrid}@csse.monash.edu.au

^b The University of Melbourne,
Parkville, VIC 3010, Australia
E-mail: {shlomo,tim}@csse.unimelb.edu.au
E-mail: l.sonenberg@unimelb.edu.au

Museums offer vast amounts of information, but a visitor's receptivity and time are typically limited, providing the visitor with the challenge of selecting the (subjectively) interesting exhibits to view within the available time. Mobile, electronic handheld guides offer the opportunity to improve a visitor's experience by recommending exhibits of interest, and adapting the delivered content. The first step in this personalisation process is the prediction of a visitor's activities and interests. In this paper we study non-intrusive, adaptive user modelling techniques that take into account the physical constraints of the exhibition layout. We present two collaborative models for predicting a visitor's next locations in a museum, and an ensemble model that combines the predictions of these models. The three models were trained and tested on a small dataset of museum visits. Our results are encouraging, with the ensemble model yielding the best performance overall.

Keywords: Collaborative user model, location prediction, museum, physical space

1. Introduction

Museums offer vast amounts of information, but a visitor's receptivity and time are typically limited. The possibility of information overload is evident, as the visitor is confronted with the challenge of selecting the

personally interesting exhibits to view within the available time. However, appropriate advance information about the exhibits in different areas of the museum is often not readily available, and visitors may change their mind about what they want to see after viewing some exhibits they thought they would be interested in. A personal human guide, aware of a visitor's interests and time limitations, could solve these problems, but the provision of personal guides is outside the resource limitations of most museums.

Advances in mobile, context-aware computing and user modelling point towards an alternative solution: personalised electronic handheld guides. Electronic guides have the potential to (1) make recommendations about items of interest, and (2) personalise the content delivered for these items, based on tracked visitor behaviour. The Kubadji project (<http://www.kubadji.org>) is developing user modelling and language technologies to support the creation of such guides. A first step is to infer a visitor's interests and activities non-intrusively from his/her behaviour within the museum, and to store the acquired information in models of the user. The physicality of the domain poses practical challenges for such user modelling [13]. For example, the spatial layout of the environment influences the curator's decisions about the positioning of the exhibits, and both influence a visitor's decisions about which exhibits to view and in which order. Hence, the spatial arrangement of items is an input that should improve the accuracy of predictions of a visitor's behaviour. To our knowledge, this factor has not been considered to date.

In this paper, we describe a first step in the recommendation and personalisation process, i. e., the prediction of a visitor's interests and locations in a museum on the basis of observed behaviour. Specifically, we consider two collaborative predictive models, *Interest* and *Transition*, and an ensemble model that combines their predictions. The *Interest Model* predicts exhibits to be viewed by a visitor on the basis of his/her observed viewing times in the context of the viewing times of other museum visitors. The *Transition Model*

*Corresponding author: Fabian Bohnert, Faculty of Information Technology, Monash University, Clayton, VIC 3800, Australia. E-mail: fabianb@csse.monash.edu.au.

**Now affiliated with: Tasmanian ICT Centre, CSIRO, Hobart, TAS 7001, Australia. E-mail: shlomo.berkovsky@csiro.au.

predicts the next exhibit (location) based on the transitions between exhibits previously recorded for other visitors. These models are employed to predict the next K exhibits ($K = 1$ and $K = 3$) using two prediction approaches: *set*, which predicts a set of exhibits, and *sequence*, which predicts an ordered sequence.

We trained and evaluated our models on a small dataset of museum visits collected at the Marine Life Exhibition in Melbourne Museum. Our results show that the *Transition Model* outperforms the *Interest Model*, indicating that the layout of a physical space with homogeneous exhibits (e. g., marine theme) is a significant factor influencing visitor behaviour. However, the ensemble model yielded the best performance overall, demonstrating the importance of considering also a visitor’s interests.

The remainder of this paper is organised as follows. In Section 2, we outline related work. Our approaches for predicting visitor locations using collaborative user models are described in Sections 3 and 4. In Section 5, we summarise preliminary findings of our evaluation, followed by a discussion of future work in Section 6.

2. Related Research

Our work lies at the intersection of statistical user modelling and personalised guide systems for physical museum spaces.

Personalised guide systems in physical domains have often employed adaptable user models, which require visitors to explicitly state their interests in some form. For example, the *GUIDE* project [8] developed a handheld tourist guide for visitors to the city of Lancaster, UK. It employed a user model obtained from explicit user input to generate a dynamic and user-adapted city tour, where the order of the visited items could be varied. In the museum domain, the *CHIP* project [2] investigates how Semantic Web techniques can be used to provide personalised access to digital museum collections both online and in the physical museum, based on explicitly initialised user models.

Less attention has been paid to predicting preferences from non-intrusive observations, and to utilising adaptive user models that do not require explicit user input. In the museum domain, adaptive user models have usually been updated from the user’s interactions with the system, with a focus on adapting content presentation, rather than predicting or recommending exhibits to be viewed. For example, *HyperAudio* [17] dynamically adapted the presented content and hyperlinks to stereotypical assumptions about the user, and

to what the user has already accessed and seems interested in. The augmented audio reality system for museums *ec(h)o* [10] treated user interests in a dynamic manner, and adapted its user models on the basis of the users’ interactions with the system. The collected user modelling data were used to deliver personalised information associated with exhibits via audio display. The *PEACH* project [19] developed a multimedia handheld guide, which adapted its user models both from explicit visitor feedback and implicit observations of visitor interactions with the device, and used the information stored in these user models to generate personalised multimedia presentations.

These systems, like most systems in the museum domain, primarily rely on knowledge-based user models, which require an explicit, a-priori engineered representation of the domain knowledge. In contrast, this work investigates non-intrusive statistical user modelling techniques that do not require an explicit representation of the domain knowledge, and takes into account spatial constraints — a factor that has not been considered to date. As far as we are aware, the only instance of the application of a statistical technique [1,20] for predicting a visitor’s behaviour in a museum is described in [9].

3. Using Collaborative Models based on Spatio-Temporal Information to Predict Location Probabilities

We consider two collaborative models for estimating the probability of a visitor viewing a particular exhibit given his/her previous visit trajectory: interest-based (Section 3.1) and transitional (Section 3.2). The interest-based approach predicts a visitor’s next location on the basis of his/her interest in unseen exhibits, which in turn is estimated from the time the visitor spent at the exhibits s/he has seen. The transitional approach predicts a visitor’s next location from the trajectories of other visitors. In Section 3.3 we propose an ensemble approach to combine the predictions generated by these models [14,18]. The utilisation of the estimated location probabilities to predict a set or sequence of next items is described in Section 4.

Recent developments in the area of positioning technology have made possible the non-intrusive indoor tracking of visitors equipped with a positioning device. The availability of such technology to infer a visitor’s high-level activities from sensing data, e. g., [15], is crucial to this work, i. e., to perform non-intrusive, adaptive user modelling. In this research, we assume

access to a visitor’s pathway in the form of a time-annotated sequence of visited items. That is, for each visitor u , we have an ordered sequence of viewing durations $t_{ui_1}, t_{ui_2}, \dots$ for items i_1, i_2, \dots respectively. This information was obtained manually for this study, by tracking 44 visitors to the Marine Life Exhibition in Melbourne Museum, but is of the same type as information inferable from sensing data in a real-world setting [15].¹ In total, our dataset comprises 317 data points (Section 5). The small size of this dataset is due to the difficulties associated with collecting and processing data in physical settings (Section 6).

3.1. Interest Model

In an information-seeking context, users are expected to spend more time on relevant than on irrelevant information, as viewing time correlates positively with preference and interest [16]. That is, the time spent at a given exhibit can be used as a measure of interest. However, viewing time is also positively correlated with item complexity. Additionally, viewing times vary over different visitors depending on their available time.² Hence, in order to infer the interests of visitors in different items, the observed viewing times cannot be used directly, but must be transformed into a measure that takes these factors into account. To this effect we have devised the *relative interest* measure, which reflects the interest of a visitor in an exhibit in the context of the time s/he has spent on previously seen exhibits, and the time spent by other visitors on this exhibit. This measure implicitly takes into account item complexity, as complex items are likely to be viewed for a longer time than simpler items.

Definition 1 (Relative Interest (RI))

The *relative interest* of visitor u in a seen exhibit i is defined as follows:

$$RI_{ui} = \frac{t_{ui}}{\bar{t}_u} - \frac{1}{n_i} \sum_{v \in U} n_{vi} \frac{t_{vi}}{\bar{t}_v}, \quad (1)$$

where t_{ui} is the time visitor u spent at exhibit i , \bar{t}_u is the average viewing time of visitor u , n_i is the number of visitors that viewed exhibit i , U is the set of visitors, and $n_{vi} = 1$ if visitor v viewed exhibit i and 0 otherwise.

¹The consideration of the impact of instrumentation accuracy on user models is outside the scope of this work.

²Viewing time was also found to be negatively correlated with familiarity, positively correlated with novelty, and decreases from beginning to end within a sequence of stops [16]. However, these factors are not yet considered in our models.

Algorithm 1 Estimating the relative interests of the active visitor in unseen exhibits

- 1: Estimate from the observed viewing times the relative interests of all visitors — including the active visitor a — in the exhibits viewed during their visit (Equation 1).
 - 2: **for all** i such that i is an unvisited exhibit **do**
 - 3: Find a set of *item mentors*, who have viewed item i , and have the highest similarity with the active visitor. To calculate a visitor-to-mentor similarity, use Pearson’s correlation coefficient on the vectors of their relative interests.
 - 4: Estimate the active visitor’s relative interest in item i as the weighted mean of the relative interests of his/her item mentors in i , where the weights are the visitor-to-mentor similarities.
 - 5: **end for**
-

The first term in Equation 1 reflects visitor u ’s viewing time of item i relative to his/her average viewing time, and the second term represents the average relative viewing time spent at item i (over all the visitors that viewed this item). Hence, RI_{ui} measures whether visitor u is (relative to his/her average viewing time) more or less interested in item i than the average interest in item i .³

The collaborative *Interest Model (IM)* is built by calculating RI_{ui} , the relative interest of visitor u in exhibit i , for all visitors $u = 1, \dots, |U|$ and all items $i = 1, \dots, |I|$, where $|U|$ is the number of visitors and $|I|$ is the number of exhibits. This yields a relative interest matrix \mathcal{RI} of size $|U| \times |I|$, which contains defined values for all combinations of visitors u and items i that occurred, i. e., combinations referring to an item i viewed by a visitor u . These values, which may be viewed as implicit ratings given by visitors to exhibits, do not take into account the order in which the exhibits were visited. In Section 6, we consider the incorporation of spatial information into our *Interest Model*.

Following a collaborative approach as described in [11], we use Algorithm 1 to predict the missing relative interest values of the active visitor a from the values in \mathcal{RI} .⁴ These values are mapped into the $[0,1]$ range to estimate the probability of visiting an unseen exhibit. Formally, given a visit where a visitor a has

³Other measures of interest are possible. For instance, Bohnert and Zukerman [4] explored a different variant of relative interest, which was slightly outperformed by the measure presented here.

⁴Although visitors sometimes return to previously viewed exhibits, our observations indicate that this rarely happens. Hence, we focus on unseen exhibits.

viewed k items so far, the probability of the $(k + 1)$ -th item being item i is represented by the expression $\Pr(X_{k+1} = i | \mathbf{v}_a^k)$, where \mathbf{v}_a^k is the user's visit history so far. Thus, approximating this expression by a probability estimated using our *Interest Model* yields the following formula:

$$\Pr(X_{k+1} = i | \mathbf{v}_a^k) \approx \Pr_{IM}(X_{k+1} = i | \mathbf{t}_a^k),$$

where \mathbf{t}_a^k is the time component of the visit history \mathbf{v}_a^k (the *Interest Model* depends on viewing times, rather than transitions between locations).⁵

3.2. Transition Model

Our *Interest Model* considers only a visitor's relative interests, and does not take into account the order in which the exhibits were visited. Here we describe an alternative model, denoted *Transition Model (TM)*, which considers the visit order.

The *Transition Model* is represented by a stationary 1-stage Markov model, where the transition matrix \mathcal{TM} approximates the probabilities of moving between exhibits. Specifically, the element $\mathcal{TM}(i, j)$ approximates the probability of a visitor going from exhibit i to exhibit j , where $i, j = 1, \dots, |I|$, and $|I|$ is the number of exhibits. This probability is estimated on the basis of the frequency of transitions between i and j . In order to overcome the data sparseness problem (which is exacerbated by our small dataset) and to smooth out outliers, we added a flattening constant ε ($= 1/|I|$) to each frequency count.

When we employ the *Transition Model* to approximate the probability that the $(k + 1)$ -th exhibit viewed by the active visitor a is item i , we obtain the formula

$$\Pr(X_{k+1} = i | \mathbf{v}_a^k) \approx \Pr_{TM}(X_{k+1} = i | I_a^k),$$

where I_a^k are the exhibits visited by the active visitor.

Since our *Transition Model* is a 1-stage Markov model, the probability of the next exhibit being item i is further approximated by

$$\begin{aligned} \Pr_{TM}(X_{k+1} = i | I_a^k) &\approx \Pr_{TM}(X_{k+1} = i | X_k = i_k) \\ &= \mathcal{TM}(i_k, i), \end{aligned}$$

where i_k is the current item. As mentioned before, our observations indicate that visitors rarely return to previously viewed exhibits. Hence, prior to calculating these probabilities, we set to 0 the entries of \mathcal{TM} that

⁵The subscript $k + 1$ of X in $\Pr_{IM}(X_{k+1} = i | \mathbf{t}_a^k)$ could be replaced by $k + m$, as the *Interest Model* predicts the probability of visiting an unseen exhibit at any point in the future.

correspond to the visited items (items in I_a^k), and appropriately renormalise the rows.

The *Transition Model* implicitly captures the physical layout of the museum space, i. e., the physical proximity of items, on the basis of the assumption that transitions to spatially close items occur more frequently than movements to items that are further away. However, in the future, we will also experiment with spatial models that represent more directly the distance between exhibits (Section 6).

3.3. Combining Interest Model and Transition Model

As indicated above, the probabilities computed by the *Interest Model* are based on temporal information, while the predictions made by the *Transition Model* implicitly capture spatial information. Additionally, while the *Interest Model* adapts to the behaviour of a visitor, the *Transition Model* is not personalised. In this section, we propose a *Hybrid Model (HM)* that combines the predictions made by these models [14,18], thereby jointly taking into account transitional and temporal information.

Formally, we use the probability $\Pr_{HM}(X_{k+1} = i | \mathbf{v}_a^k)$ generated by our ensemble model to approximate $\Pr(X_{k+1} = i | \mathbf{v}_a^k)$:

$$\Pr(X_{k+1} = i | \mathbf{v}_a^k) \approx \Pr_{HM}(X_{k+1} = i | \mathbf{v}_a^k).$$

This probability in turn is calculated by means of a weighted average of the predictions generated by our *Interest Model* and *Transition Model*, i. e.,

$$\begin{aligned} \Pr_{HM}(X_{k+1} = i | \mathbf{v}_a^k) \\ = \omega \Pr_{IM}(X_{k+1} = i | \mathbf{t}_a^k) \\ + (1 - \omega) \Pr_{TM}(X_{k+1} = i | I_a^k), \end{aligned}$$

where $0 \leq \omega \leq 1$. We experimented with different values for ω , with the assignment $\omega = \beta / (\alpha + \beta)$ yielding the best performance,⁶ where

$$\alpha = \min_{i \in I \setminus I_a^k} \Pr_{IM}(X_{k+1} = i | \mathbf{t}_a^k), \text{ and}$$

$$\beta = \min_{i \in I \setminus I_a^k} \Pr_{TM}(X_{k+1} = i | I_a^k),$$

and $I \setminus I_a^k$ is the set of exhibits not yet visited. This choice of ω assigns more weight to the model with the lower minimum prediction, which may be viewed as the more discriminating model.

⁶In the future, we intend to apply machine learning techniques to learn the optimal value for ω .

4. Building Models to Predict a Visitor's Next K Locations

In this section, we describe two approaches for using the probabilities estimated in Section 3 to predict the next K exhibits to be viewed by a visitor: *TopK*, which predicts the next K items as a set and ranks them in descending order of estimated probability, and *SequenceK/N*, which predicts the next K items as the initial portion of a sequence of N items.

4.1. TopK Prediction

The *TopK* approach assumes that the current visit history suffices to predict a visitor's future behaviour, and that it is unnecessary to consider the impact of future transitions on the visitor's subsequent behaviour. Hence, in order to predict the next K items that will be visited (having visited k items), we find the set of K unvisited items i_{k+1}, \dots, i_{k+K} that maximises the product of their visit probabilities. This is done by solving

$$\arg \max_{i_{k+1}, \dots, i_{k+K} \in I \setminus I_a^k} \prod_{m=1}^K \Pr(X_{k+m} = i_{k+m} \mid \mathbf{v}_a^k).$$

This approach is equivalent to computing the probabilities $\Pr(X_{k+1} = i \mid \mathbf{v}_a^k)$ for all unvisited exhibits $i \in I \setminus I_a^k$ (pretending that each of these exhibits is the next one — hence the subscript $k + 1$), then sorting the exhibits in descending order of estimated visit probabilities, and selecting the top K items.

4.2. SequenceK/N Prediction

In contrast to the *TopK* approach, the *SequenceK/N* approach assumes that future transitions influence a visitor's subsequent behaviour. Hence, in order to predict the next K items that will be visited (having visited k items), we first find the maximum-probability sequence of N unvisited items i_{k+1}, \dots, i_{k+N} by solving

$$\arg \max_{i_{k+1}, \dots, i_{k+N} \in I \setminus I_a^k} \Pr(X_{k+1} = i_{k+1}, \dots, X_{k+N} = i_{k+N} \mid \mathbf{v}_a^k),$$

and then select the first K items i_{k+1}, \dots, i_{k+K} within this sequence.

Assuming that X_{k+m} depends only on the past, this probability is decomposed as follows:

$$\begin{aligned} & \Pr(X_{k+1} = i_{k+1}, \dots, X_{k+N} = i_{k+N} \mid \mathbf{v}_a^k) \\ &= \prod_{m=1}^N \Pr(X_{k+m} = i_{k+m} \mid \mathbf{v}_a^{k+m-1}). \end{aligned} \quad (2)$$

Due to this decomposition, the joint probability in Equation 2 can be maximised by recursively spanning a search tree of depth $N - 1$, and performing an exhaus-

sive search for a maximising path from its root to one of the leaves.

The probability $\Pr(X_{k+m} = i_{k+m} \mid \mathbf{v}_a^{k+m-1})$ in Equation 2 depends on the active visitor's visit history up to exhibit i_{k+m-1} , but in practice this history is available only up to item i_k . Future exhibits are incorporated into a "potential history" for the *Transition Model* by iteratively adding predicted unseen exhibits to construct different potential sequences. In order to incorporate such a potential history into the *Interest Model* (and hence the *Hybrid Model*), we require viewing time estimates. The calculation of predicted viewing times is similar to that performed for the estimation of relative interests, and is described in [6].

5. Preliminary Findings

The data used to obtain the results reported in this paper was collected at the Marine Life Exhibition of Melbourne Museum in 2006. This exhibition consists of 56 exhibits in four sections, comprising a rather homogeneous selection of marine-related topics. With the help of curators, we transformed the original set of 56 exhibits into a set of 22 grouped exhibits by unifying logically related exhibits. To collect the dataset, 44 visitors were manually tracked at the exhibition, which allowed us to construct time-annotated visitor pathways. In summary, we obtained 317 observations, with an average visitor path length of 7.20 grouped exhibits, and the shortest and longest visits comprising 3 and 16 exhibits respectively.

In our experiments, we evaluated the performance of our two approaches for predicting a visitor's next K exhibits, *TopK* and *SequenceK/N*, for two values of K ($K = 1$ and $K = 3$), and a fixed value for N ($N = 3$), yielding the four variants *Top1*, *Sequence1/3*, *Top3* and *Sequence3/3*. For every combination of prediction mode (set and sequence) and value of K ($K = 1$ and $K = 3$), we considered the three prediction models defined in Section 3: *Interest Model*, *Transition Model* and *Hybrid Model*. We evaluated the performance of these 12 variants by measuring their *Classification Accuracy* and *Ranking Accuracy* [12]. Here we focus on the results obtained for Classification Accuracy (the results obtained for Ranking Accuracy are detailed in [5,6]). Classification Accuracy was calculated using Precision, i. e., $\text{Pre} = |\mathcal{K} \cap \mathcal{M}|/|\mathcal{K}|$, where \mathcal{K} (with cardinality K) is the set of predicted exhibits, and \mathcal{M} is the set of actually viewed exhibits. We considered two alternatives for \mathcal{M} : (1) the set of K exhibits viewed next ($|\mathcal{M}| = K$), and (2) the set of ex-

hibits viewed during the remainder of the visit (which may be of different size for each visitor). These alternatives for \mathcal{M} correspond to evaluating the accuracy of our predictions in the immediate future and in the eventual future respectively.

Our results indicate that, as expected, our predictions are less accurate when evaluated for the immediate future than when evaluated for the eventual future (on average, immediate accuracy is lower than eventual accuracy by 20%). Additionally, the spatial structure of the museum space dominates in our constrained, homogeneous setup, i. e., the *Transition Model* significantly outperforms the *Interest Model*. For the eventual evaluation setting, we obtained the following Classification Accuracy results for the *Hybrid Model* (analogous insights apply to the immediate setting). When predicting the next exhibit ($K = 1$) to be viewed by a visitor, (1) the accuracy of *Top1* predictions is comparable to the accuracy of the more complex *Sequence1/3* predictions (66% vs. 68%), hence there is no need to apply a complex sequence prediction mechanism; and (2) hybridisation yields only a slight improvement in predictive accuracy over the *Transition Model*. In contrast, when predicting a sequence of $K = 3$ exhibits, (1) *Sequence3/3* is superior to simple *Top3* (59% vs. 49%), meaning that sequence information aids prediction; and (2) the *Hybrid Model* significantly outperforms the individual models. Consequently, hybridising the *Interest Model* and the *Transition Model* is generally beneficial.

6. Discussion and Future Work

Our results raise several issues with respect to modelling users in physical spaces in general, and museums in particular. These issues pertain to exhibit diversity, amount and quality of data, user modelling strategies, and recommendations.

Exhibit diversity. Our experiments were conducted in the rather homogeneous Marine Life Exhibition. This means that the visitors who decided to enter the exhibition were already interested in marine life. Thus, as indicated by our results, a key factor influencing visitor behaviour in a homogeneous exhibition is the physical layout of the space. However, this conclusion may not be valid in a space of heterogeneous exhibits, such as the entire Melbourne Museum, which has exhibits relating to flora, fauna, Australian history and modern life. These observations motivate future experiments at different levels of granularity, e. g., inter-exhibition

versus intra-exhibition, while considering the link between granularity and topic diversity. We expect these experiments to shed light on the influence of exhibition size and exhibit diversity on the applicability of our models.

Amount and quality of data. Our dataset comprises visit trajectories of 44 museum visitors. This is a rather small dataset. However, in contrast to web-based data collections, the collection of visit traces in a physical space is an expensive and time-consuming process, when done by either human trackers or electronic devices. Additionally, cost-efficient tracking technology is still relatively inaccurate, which may affect the accuracy of the derived user models and the quality of the personalisation provided to the visitor [7]. This problem should be addressed prior to deploying such devices. At the same time, this problem obfuscates basic user modelling issues, and should be avoided during initial model development. In contrast, human tracking is precise, and hence ideal for initial model development, but clearly cannot be used during model deployment. Using human trackers, we have just completed the collection of additional visit traces at the level of the entire Melbourne Museum. However, the amount of data we collected remained relatively small.

User modelling strategies. Our current approach for combining user models belongs to the ensemble category, where the predictions made are combined in a weighted manner [18]. However, the models themselves are built separately; the *Transition Model* from trajectory information, and the *Interest Model* from temporal information. In the future, we propose to combine these information sources and conduct hybridisation at the model acquisition stage. For example, this can be done by considering the distance from a current exhibit when computing a visitor's *Interest Model*. That is, the farther a newly visited exhibit is from the last visited exhibit, the higher the interest in the new exhibit. We also plan to investigate an ensemble combination of collaborative user models with content-based models, and intend to take into account the cold-start problem [11,20] by applying machine learning techniques to determine the point in a visit at which personalised models can be deployed. These techniques will also be applied to find the optimal weight of the individual models in ensemble models.

Recommendations. In a physical domain, the transition from interest and location prediction to recommending interesting items (locations; or in this particular work, exhibits) is not trivial. We suggest the fol-

lowing approach. First, generate the set of unvisited items, and rank them using predictions from an interest model. Second, use a transition (or hybrid) model to predict the locations that a visitor is most likely to visit next. Implement a strategy for merging these lists, e. g., for whether locations that the visitor is likely to visit anyway should be included (to help build trust in the system) or excluded (to avoid over-communication). Finally, deliver the list of recommended locations to the visitor allowing the visitor to act on the recommendations as they see fit. The modality of the presentation, e. g., visualised on a site map, or provided more explicitly in textual or audio form, should be taken into account in constructing the list.

In the future, we plan to investigate utility-based recommendation generation strategies which balance factors such as those above, e. g., Markov Decision Processes, which were recently proposed for decision-theoretic and user-adaptable planning in the shopping guide domain [3].

Acknowledgements

This research was supported in part by Discovery grant DP0770931 from the Australian Research Council. The authors thank Enes Makalic for his assistance with ensemble models. Thanks also go to Carolyn Meehan and her team from Museum Victoria for fruitful discussions and the dataset.

References

- [1] David W. Albrecht and Ingrid Zukerman, editors. Special issue on statistical and probabilistic methods for user modeling. *User Modeling and User-Adapted Interaction*, 17(1-2), 2007.
- [2] Lora Aroyo, Natalia Stash, Yiwen Wang, Peter Gorgels, and Lloyd Rutledge. CHIP demonstrator: Semantics-driven recommendations and museum tour generation. In *Proc. of the Sixth Intl. Semantic Web Conf. (ISWC-07)*, pages 879–886, 2007.
- [3] Thorsten Bohnenberger, Oliver Jacobs, Anthony Jameson, and Ilhan Aslan. Decision-theoretic planning meets user requirements: Enhancements and studies of an intelligent shopping guide. In *Proc. of the Third Intl. Conf. on Pervasive Computing (Pervasive-05)*, pages 279–296, 2005.
- [4] Fabian Bohnert and Ingrid Zukerman. Using viewing time for theme prediction in cultural heritage spaces. In *Proc. of the 20th Australian Joint Conf. on Artificial Intelligence (AI-07)*, pages 367–376, 2007.
- [5] Fabian Bohnert, Ingrid Zukerman, Shlomo Berkovsky, Timothy Baldwin, and Liz Sonenberg. Using collaborative models to adaptively predict visitor locations in museums. In *Proc. of the Fifth Intl. Conf. on Adaptive Hypermedia and Adaptive Web-Based Systems (AH-08)*, pages 42–51, 2008.
- [6] Fabian Bohnert, Ingrid Zukerman, Shlomo Berkovsky, Timothy Baldwin, and Liz Sonenberg. Using interest and transition models to predict visitor locations in museums. Technical Report 2008/219, Faculty of Information Technology, Monash University, Clayton, VIC 3800, Australia, 2008.
- [7] David J. Carmichael, Judy Kay, and Bob Kummerfeld. Consistent modelling of users, devices and sensors in a ubiquitous computing environment. *User Modeling and User-Adapted Interaction*, 15(3-4):197–234, 2005.
- [8] Keith Cheverst, Keith Mitchell, and Nigel Davies. The role of adaptive hypermedia in a context-aware tourist guide. *Communications of the ACM*, 45(5):47–51, 2002.
- [9] Karl Grieser, Timothy Baldwin, and Steven Bird. Dynamic path prediction and recommendation in a museum environment. In *Proc. of the ACL Workshop on Language Technology for Cultural Heritage Data (LaTeCH-07)*, in conjunction with ACL-07, pages 49–56, 2007.
- [10] Marek Hatala and Ron Wakkary. Ontology-based user modeling in an augmented audio reality system for museums. *User Modeling and User-Adapted Interaction*, 15(3-4):339–380, 2005.
- [11] Jonathan L. Herlocker, Joseph A. Konstan, Al Borchers, and John T. Riedl. An algorithmic framework for performing collaborative filtering. In *Proc. of the 22nd Annual Intl. ACM SIGIR Conf. on Research and Development in Information Retrieval (SIGIR-99)*, pages 230–237, 1999.
- [12] Jonathan L. Herlocker, Joseph A. Konstan, Loren G. Terveen, and John T. Riedl. Evaluating collaborative filtering recommender systems. *ACM Transactions on Information Systems*, 22(1):5–53, 2004.
- [13] Anthony Jameson and Antonio Krüger, editors. Special issue on user modeling in ubiquitous computing. *User Modeling and User-Adapted Interaction*, 15(3-4), 2005.
- [14] George Lekakos and George M. Giaglis. A hybrid approach for improving predictive accuracy of collaborative filtering algorithms. *User Modeling and User-Adapted Interaction*, 17(1-2):5–40, 2007.
- [15] Lin Liao, Dieter Fox, and Henry Kautz. Extracting places and activities from GPS traces using hierarchical conditional random fields. *Intl. Journal of Robotics Research*, 26(1):119–134, 2007.
- [16] Jeffrey Parsons, Paul Ralph, and Katherine Gallager. Using viewing time to infer user preference in recommender systems. In *Proc. of the AAAI Workshop on Semantic Web Personalization (SWP-04)*, in conjunction with AAAI-04, pages 52–64, 2004.
- [17] Daniela Petrelli and Elena Not. User-centred design of flexible hypermedia for a mobile guide: Reflections on the HyperAudio experience. *User Modeling and User-Adapted Interaction*, 15(3-4):303–338, 2005.
- [18] Robi Polikar. Ensemble based systems in decision making. *IEEE Circuits and Systems Magazine*, 6(3):21–45, 2006.
- [19] Oliviero Stock, Massimo Zancanaro, Paolo Busetta, Charles Callaway, Antonio Krüger, Michael Kruppa, Tsvi Kuflik, Elena Not, and Cesare Rocchi. Adaptive, intelligent presentation of information for the museum visitor in PEACH. *User Modeling and User-Adapted Interaction*, 18(3):257–304, 2007.
- [20] Ingrid Zukerman and David W. Albrecht. Predictive statistical models for user modeling. *User Modeling and User-Adapted Interaction*, 11(1-2):5–18, 2001.