Using Collaborative Models to Adaptively Predict Visitor Locations in Museums

Fabian Bohnert[1], Ingrid Zukerman[1],
Shlomo Berkovsky[2], Timothy Baldwin[2], and Liz Sonenberg[2]

[1] Monash University
Clayton, Victoria 3800, Australia
{Fabian.Bohnert,Ingrid.Zukerman}@infotech.monash.edu.au

[2] The University of Melbourne
Carlton, Victoria 3010, Australia
{shlomo,tim}@csse.unimelb.edu.au
l.sonenberg@unimelb.edu.au

Abstract. The vast amounts of information presented in museums can be overwhelming to a visitor, whose receptivity and time are typically limited. Hence, s/he might have difficulties selecting interesting exhibits to view within the available time. Mobile, context-aware guides offer the opportunity to improve a visitor's experience by recommending exhibits of interest, and personalising the delivered content. The first step in this recommendation process is the accurate prediction of a visitor's activities and preferences. In this paper, we present two adaptive collaborative models for predicting a visitor's next locations in a museum, and an ensemble model that combines their predictions. Our experimental results from a study using a small dataset of museum visits are encouraging, with the ensemble model yielding the best performance overall.

1 Introduction

Museums offer vast amounts of information, but since a visitor's receptivity and time are typically limited, s/he is confronted with the challenge of selecting interesting exhibits to view during a visit. A personal human guide who is knowledgeable about the museum's exhibits and aware of the visitor's interests and time limitations could support the visitor in this selection process, but the provision of personal guides is generally impractical. Advances in mobile, context-aware computing and user modelling point towards an alternative solution: electronic handheld guides. These guides have the potential to (1) make recommendations about items of interest, and (2) personalise the content delivered for these items; based on predictions of a visitor's activities and interests estimated from non-intrusive observations of his/her behaviour.

In this paper, we describe the first step in this process. We consider two collaborative predictive models of visitor behaviour, *Interest* and *Transition*, and an ensemble model that combines their predictions. These models are employed to predict the next K (= 3) exhibits to be viewed by a visitor, using two prediction approaches: *set*, which predicts a set of exhibits, and *sequence*, which predicts a sequence. Accurately predicting a visitor's next locations will enable us to deliver useful recommendations about

Bohnert, Fabian, Ingrid Zukerman, Shlomo Berkovsky, Timothy Baldwin and Elizabeth Sonenberg (2008) Using Collaborative Models to Adaptively Predict Visitor Locations in Museums, In Proceedings of the 5th International Conference on Adaptive Hypermedia and Adaptive Web-Based Systems, Hanover, Germany, pp. 42-51.

exhibits to visit, e. g., by excluding from the set of potential recommendations the exhibits that a visitor is likely to see anyway in the near future. We trained and tested our models on a small dataset 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 dominating factor influencing visitor behaviour. However, the ensemble model yielded the best performance overall with an average accuracy of 59%, demonstrating the importance of considering also a visitor's interests. Additionally, we found that our sequence-based prediction model has a significantly higher accuracy than our set-based prediction model (59% vs. 49%).

The rest of this paper is organised as follows. In Section 2 we outline related work, and in Section 3 we introduce the domain. Our predictive approaches are described in Sections 4 and 5. In Section 6 we present the results of our evaluation, followed by our conclusions in Section 7.

2 Related Research

Our work lies at the intersection of statistical user modelling [1] 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, e. g., [2, 3]. 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 [4–6] rather than predicting or recommending exhibits to be viewed. These systems, like most systems in the museum domain, rely on *knowledge-based user models*, which require an explicit and a-priori built 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.

3 Domain and Dataset

The data used in the experiments reported in this paper was obtained by manually tracking visitors to the Marine Life Exhibition of Melbourne Museum in 2006. This exhibition consists of 56 exhibits in four sections, displaying 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, such as a visual display and its accompanying explanatory panel. Figure 1 depicts the layout of the exhibition space and the exhibition highlight "Whale meets Squid". In the initial stage of their visit, visitors pass through a highly constrained entrance area where they behave similarly. This area leads to a space with several open sections, where visitor behaviour is less prescribed. However, at around the 55%–60% point of their visit, most visitors enter the area from which the "Whale meets Squid" exhibit is visible, and gravitate towards it.

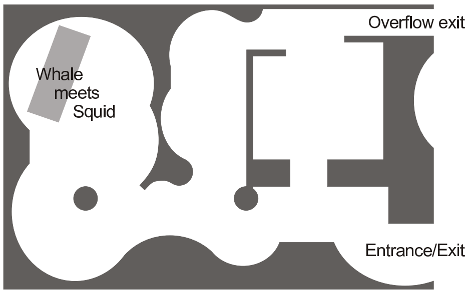

(a) Exhibition layout

(b) “Whale meets Squid” exhibit

Fig. 1: The Marine Life Exhibition

Owing to the difficulties associated with collecting data in a physical space, our dataset consists of only 44 visitor pathways, which comprise a total of 317 stops at grouped exhibits. On average, visitors viewed 7.20 exhibits, with the shortest and longest pathways being 3 and 16 exhibits respectively.

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

In this work, we consider two collaborative approaches for estimating the probability of a visitor viewing a particular exhibit given his/her previous visit trajectory: interest-based (Section 4.1) and transitional (Section 4.2). 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 saw. The transitional approach predicts a visitor's next location on the basis of the pathways followed by other visitors in the museum. In Section 4.3 we propose an ensemble approach that combines the predictions generated by these models [7, 8]. The utilisation of the estimated location probabilities to predict a set or sequence of next items is described in Section 5.

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 as a basis for inferring a visitor's high-level activities from sensing data 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 $t_{ui_1}, t_{ui_2}, \ldots$ for items $i_1, i_2, \ldots$ respectively. As stated above, this information was obtained by tracking people manually, but is of the same type as information inferable from sensing data in a real-world setting.[3]

4.1 Interest Model

In an information-seeking context, users are expected to spend more time on relevant information than on irrelevant information, as viewing time correlates positively with

[3] The consideration of the impact of instrument accuracy on user models is outside the scope of this work.

preference and interest [9]. Hence, viewing time 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 the time available for their visit.[4] In order to infer the interests of visitors in different items while taking into account these factors, we have devised the *relative interest* measure below. It 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.

$$RI_{ui} = \frac{t_{ui}}{\bar{t}_{u\cdot}} - \frac{1}{n_{\cdot i}} \sum_{v \in U} n_{vi} \frac{t_{vi}}{\bar{t}_{v\cdot}} \tag{1}$$

where t_{ui} is the time visitor u spent at exhibit i, $\bar{t}_{u\cdot}$ is the average viewing time of visitor u, $n_{\cdot i}$ 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.

The first term in Equation 1 reflects visitor u's viewing time for 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.[5]

The collaborative *Interest Model (IM)* is built by calculating RI_{ui}, the relative interest of visitor u in exhibit i, for all visitors $u = 1, \ldots, |U|$ and all items $i = 1, \ldots, |I|$, 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 regarded as implicit ratings given by visitors to exhibits, do not take into account the order in which the exhibits were visited.

Following the collaborative approach described in [10], we use Algorithm 1 to predict the missing relative interest values of the active user a from the values in $\mathcal{RI}$. These values are mapped into the $[0, 1]$ range to estimate the probability of visiting an unseen exhibit [11]. Formally, given a visit where a user a has viewed k items so far, the probability of the $(k+1)$-th item being item i is represented by the expression $\Pr(X_{k+1} = i \mid \boldsymbol{v}_a^k)$, where $\boldsymbol{v}_a^k$ is the user's visit history so far. Approximating this expression by a probability estimated using our *Interest Model* yields the following formula.

$$\Pr\left(X_{k+1} = i \mid \boldsymbol{v}_a^k\right) \approx \Pr_{IM}\left(X_{k+1} = i \mid \boldsymbol{t}_a^k\right)$$

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

[4] Viewing 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 [9]. However, these factors are not yet considered in our model.

[5] Clearly, other measures of interest are possible. The measure proposed here outperformed other variants of relative interest we have explored.

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 items 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 whose relative interests are most similar to those of the active visitor. To calculate a visitor-to-mentor similarity, use Pearson's correlation coefficient.
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**

4.2 Transition Model

In contrast to the *Interest Model*, the *Transition Model (TM)* considers the order in which exhibits were visited. 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 $i, j = 1, \ldots, |I|$ and $|I|$ is the number of exhibits. This probability is estimated on the basis of the frequency count 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 before normalising each row of $\mathcal{TM}$ to 1.

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

$$\Pr\left(X_{k+1} = i \mid \boldsymbol{v}_a^k\right) \approx \Pr_{TM}\left(X_{k+1} = i \mid I_a^k\right)$$

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

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

$$\Pr_{TM}(X_{k+1} = i \mid I_a^k) \approx \Pr_{TM}(X_{k+1} = i \mid X_k = i_k) = \mathcal{TM}(i_k, i)$$

where i_k is the current item. Although visitors sometimes return to previously viewed exhibits, our observations indicate that this rarely happens. Hence, we focus on unseen exhibits. That is, prior to calculating the transition probabilities, we set to 0 the entries of $\mathcal{TM}$ that correspond to the visited items, i. e., the 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 intend to experiment with spatial models that represent more directly the spatial proximity between exhibits.

4.3 Combining Interest Model and Transition Model

As outlined 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 an ensemble *Hybrid Model (HM)* that combines the predictions made by these models [7, 8], thereby jointly taking into account transitional and temporal information.

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

$$\Pr\left(X_{k+1} = i \mid \boldsymbol{v}_a^k\right) \approx \Pr_{HM}\left(X_{k+1} = i \mid \boldsymbol{v}_a^k\right)$$

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.,

$$\Pr_{HM}\left(X_{k+1} = i \mid \boldsymbol{v}_a^k\right) = \omega \Pr_{IM}\left(X_{k+1} = i \mid \boldsymbol{t}_a^k\right) + (1-\omega)\Pr_{TM}\left(X_{k+1} = i \mid I_a^k\right)$$

where the weight ω is chosen from the range $[0, 1]$. We experimented with different values for ω, with the assignment $\omega = \beta / (\alpha + \beta)$ yielding the best performance,[6] where

$$\alpha = \min_{i \in I \backslash I_a^k} \Pr_{IM}\left(X_{k+1} = i \mid \boldsymbol{t}_a^k\right) \quad \text{and} \quad \beta = \min_{i \in I \backslash I_a^k} \Pr_{TM}\left(X_{k+1} = i \mid I_a^k\right)$$

and $I \backslash I_a^k$ 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.

5 Building Models to Predict the Next *K* Exhibits

In this section, we describe two approaches for using the probabilities estimated in Section 4 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 *SeqK/N*, which predicts the next K items as the initial portion of a sequence of N items.

5.1 *TopK* Prediction

The *TopK* approach assumes that the current history of the active visitor a is sufficient to predict his/her 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 to be visited (having visited k items), we find the set of K unvisited items $i_{k+1}, \ldots, i_{k+K}$ which maximises the product of their visit probabilities by solving

$$\operatorname*{arg\,max}_{i_{k+1},\ldots,i_{k+K} \in I \backslash I_a^k} \prod_{m=1}^{K} \Pr\left(X_{k+1} = i_{k+m} \mid \boldsymbol{v}_a^k\right)$$

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

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

5.2 *SeqK/N* Prediction

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

$$\underset{i_{k+1},\ldots,i_{k+N} \in I \setminus I_a^k}{\arg\max} \Pr\left(X_{k+1} = i_{k+1}, \ldots, X_{k+N} = i_{k+N} \mid \boldsymbol{v}_a^k\right)$$

and then select the first K items $i_{k+1}, \ldots, i_{k+K}$ within this sequence. Assuming that X_{k+m} depends only on the past, this probability is decomposed as follows.

$$\Pr\left(X_{k+1} = i_{k+1}, \ldots, X_{k+N} = i_{k+N} \mid \boldsymbol{v}_a^k\right) = \prod_{m=1}^{N} \Pr\left(X_{k+m} = i_{k+m} \mid \boldsymbol{v}_a^{k+m-1}\right) \quad (2)$$

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

The probability $\Pr(X_{k+m} = i_{k+m} \mid \boldsymbol{v}_a^{k+m-1})$ in Equation 2 depends on the active user'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 also need to predict viewing times. The calculation of the estimated viewing times is similar to that performed for the estimation of relative interests, and is described in detail in [11].

6 Evaluation

In our experiments, we evaluated the performance of our two approaches for predicting the next K exhibits to be viewed by a visitor, *TopK* and *SeqK/N*, for $K = 3$ and $N = 3$, yielding the two variants *Top*3 and *Seq*3/3. For both prediction modes, we considered the three prediction models defined in Section 4 — *Interest Model (IM)*, *Transition Model (TM)* and *Hybrid Model (HM)* — yielding a total of six variants. Due to the small size of our dataset (Section 3), we used leave-one-out validation, i. e., we trained our prediction models on 43 of the 44 visitors in our dataset, and tested them on the remaining visitor (the active visitor). Additionally, we considered only the portion of a museum visit for which a collaborative *Interest Model* could be constructed (i. e., for which the active visitor's similarity with the other visitors could be computed). Hence, we report on the results obtained only after at least three observations have been made for the active visitor. Also, to be able to evaluate the predictions of the final three items viewed in a visit, we stopped simulating the visit history of the active visitor at that point. To obtain statistically valid results, we considered only visit percentages where at least 10 visitors were observed. Owing to these considerations, the results presented in this paper pertain to the middle part of a museum visit, spanning between 25% and 70% of a visit.

For each visit percentage, we averaged the values obtained for the following evaluation measures for all the active visitors in the test set (we considered visit percentage, rather than number of viewed exhibits, because this number varies across visitors).[7]

- ***Precision (Pre)*** – $Pre = |\mathcal{K} \cap \mathcal{M}|/|\mathcal{K}|$, the proportion of the $|\mathcal{K}|$ (= 3) predicted exhibits in $\mathcal{K}$ that appear in the set $\mathcal{M}$ of exhibits viewed during the remainder of the visit; and
- ***Modified Spearman (mSP)*** – a modified version of Spearman's rank correlation [13], measuring the fit between the predicted exhibit sequence and the sequence of actually visited exhibits (a modified version is required because the sequences being compared may be of different lengths [11]).

The results of our experiments are summarised in Figure 2. For both measures and both prediction modes, the overall performance of *HM* is at least as good as the performance of *TM*, and both of these methods perform considerably better than *IM*. Specifically, for the *Pre* measure (Figures 2a and 2b), the difference between the performance of *HM* and that of *IM* is statistically significant ($p < 0.05$) for most of the visit for both *Top*3 and *Seq*3/3.[8] The difference between *HM* and *TM* is statistically significant ($p < 0.1$) for *Seq*3/3 for up to 50% of a visit, but it is not significant for *Top*3. For the *mSP* measure (Figures 2c and 2d), *HM* and *TM* perform similarly in the *Top*3 mode (the difference is not statistically significant), while *HM* significantly outperforms *IM* for the initial stages of a visit and for visit percentages larger than 45% ($p < 0.05$). For the *Seq*3/3 mode, *HM* outperforms *IM* ($p < 0.05$) throughout a visit, and *TM* ($p < 0.05$) for the first 50% of a visit. Comparing the prediction modes *Top*3 and *Seq*3/3, *Top*3 *IM* and *Seq*3/3 *IM* perform similarly, as do *Top*3 *TM* and *Seq*3/3 *TM*. However, *Seq*3/3 *HM* yields a higher precision than *Top*3 *HM* for most of a visit ($p < 0.1$), and a higher value for *mSP* for 30%–50% of a visit ($p < 0.1$). On average, *Seq*3/3 *HM* yields 59% for *Pre* and 46% for *mSP*, whereas *Top*3 *HM* performs at 49% with respect to *Pre* and at 40% with respect to *mSP*.

The results in Figure 2 highlight the relationship between the exhibition layout and the relative performance of our predictive models. Figures 2b–2d show a divergence in the performance of *TM* and *IM* during the initial stages of a visit (the accuracy of *TM* is relatively high, while the accuracy of *IM* is relatively low), and Figures 2a–2c show such a divergence around the 55%–60% point of a visit. These regions of divergence coincide with those visit percentages where a visitor's behaviour is constrained by the physical layout (the entrance area and the point where the highlight exhibit becomes visible, Section 3). Additionally, our results show a performance decrease for the *IM* variants as the visit percentage increases (Figures 2a–2c). This may be due to our *Interest Model* disregarding the fact that viewing time decreases within a sequence of stops (Section 4.1).

[7] In agreement with Herlocker *et al.*'s observations regarding the impracticality of using recall for recommender systems [12], we eschew the calculation of recall. That is, due to the large number of exhibits left to be viewed at most stages of a visit (i. e., $|\mathcal{M}| \gg 3$), our setup would yield low recall values, which are not comparable to the values obtained for precision.

[8] Throughout this paper, the statistical tests performed are paired two-tailed t-tests. Also, we consider $p > 0.1$ to indicate a lack of statistical significance.

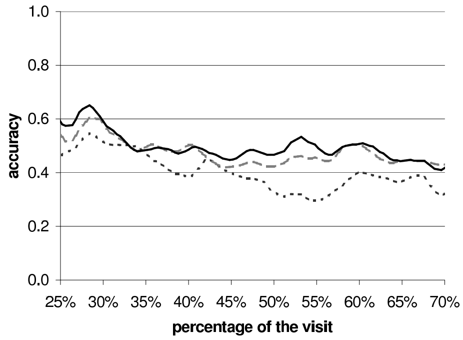

(a) *Pre* for *Top*3

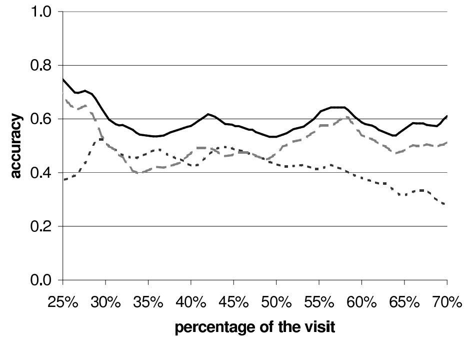

(b) *Pre* for *Seq*3/3

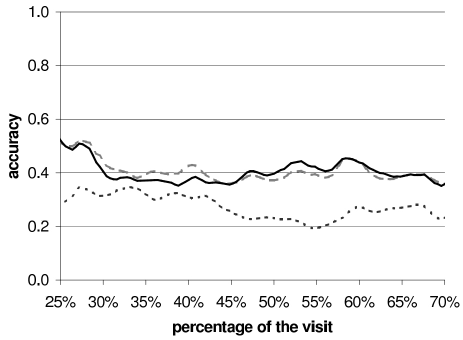

(c) *mSP* for *Top*3

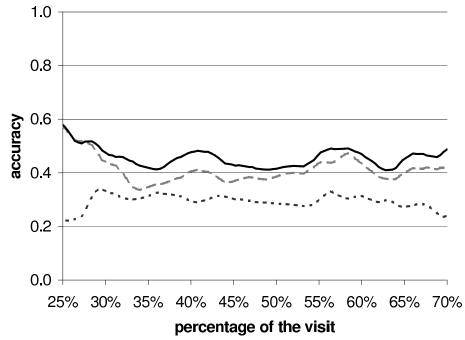

(d) *mSP* for *Seq*3/3

Fig. 2: Performance of predictive models

In summary, when predicting a sequence of $K = 3$ exhibits, (1) *Seq*3/3 is superior to *Top*3, meaning that sequence information aids prediction, and (2) *TM* and *IM* should be hybridised, as their combined predictive accuracy surpasses that of the individual methods.

7 Conclusions and Future Work

We have offered two models for predicting visitor locations in a museum — a *Transition Model* implicitly capturing spatial information, and an *Interest Model* based on viewing times — and we have combined these models into a hybrid ensemble model. The performance of these models was tested on a small dataset collected from visitors to 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 is a dominating factor influencing visitor behaviour. Nevertheless, the *Hybrid Model* yielded the best performance overall, which demonstrates the importance of also considering a visitor's interests. Additionally, our results show that when predicting the next three exhibits to be viewed, a model that predicts a sequence of items has a higher accuracy than a model that predicts a ranked set.

In this work, our experiments were conducted using a small dataset obtained from a single exhibition comprising a homogeneous set of exhibits. The small size of the dataset affects the applicability of probabilistic models. Additionally, its homogeneity reduces the impact of a visitor's interests on his/her behaviour, and consequently the usefulness of a predictive model of interest. In the near future, we intend to address these problems by collecting additional traces of visit trajectories over areas of the museum with more heterogeneous content.

Acknowledgements. This research was supported in part by 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, their support of this research, and the dataset.

References

1. Zukerman, I., Albrecht, D.W.: Predictive statistical models for user modeling. User Modeling and User-Adapted Interaction **11**(1-2) (2001) 5–18
2. Cheverst, K., Mitchell, K., Davies, N.: The role of adaptive hypermedia in a context-aware tourist guide. Communications of the ACM **45**(5) (2002) 47–51
3. Aroyo, L., Stash, N., Wang, Y., Gorgels, P., Rutledge, L.: CHIP demonstrator: Semantics-driven recommendations and museum tour generation. In: Proceedings of the Sixth International Semantic Web Conference (ISWC-07). (2007) 879–886
4. Petrelli, D., Not, E.: User-centred design of flexible hypermedia for a mobile guide: Reflections on the HyperAudio experience. User Modeling and User-Adapted Interaction **15**(3-4) (2005) 303–338
5. Hatala, M., Wakkary, R.: Ontology-based user modeling in an augmented audio reality system for museums. User Modeling and User-Adapted Interaction **15**(3-4) (2005) 339–380
6. Stock, O., Zancanaro, M., Busetta, P., Callaway, C., Krüger, A., Kruppa, M., Kuflik, T., Not, E., Rocchi, C.: Adaptive, intelligent presentation of information for the museum visitor in PEACH. User Modeling and User-Adapted Interaction **18**(3) (2007) 257–304
7. Lekakos, G., Giaglis, G.M.: A hybrid approach for improving predictive accuracy of collaborative filtering algorithms. User Modeling and User-Adapted Interaction **17**(1-2) (2007) 5–40
8. Polikar, R.: Ensemble based systems in decision making. IEEE Circuits and Systems Magazine **6**(3) (2006) 21–45
9. Parsons, J., Ralph, P., Gallager, K.: Using viewing time to infer user preference in recommender systems. In: Proceedings of the AAAI Workshop on Semantic Web Personalization (SWP-04). (2004) 52–64
10. Herlocker, J.L., Konstan, J.A., Borchers, A., Riedl, J.: An algorithmic framework for performing collaborative filtering. In: Proceedings of the 22nd Annual International ACM Conference on Research and Development in Information Retrieval (SIGIR-99). (1999) 230–237
11. Bohnert, F., Zukerman, I., Berkovsky, S., Baldwin, T., Sonenberg, L.: Using interest and transition models to predict visitor locations in museums. Technical Report 2008/219, Faculty of Information Technology, Monash University, Clayton, Victoria 3800, Australia (2008)
12. Herlocker, J.L., Konstan, J.A., Terveen, L.G., Riedl, J.T.: Evaluating collaborative filtering recommender systems. ACM Transactions on Information Systems **22**(1) (2004) 5–53
13. Siegel, S., Castellan, N.J.: Non-Parametric Statistics for the Behavioral Sciences. 2nd edn. McGraw–Hill, Inc. (1988)