

PathRec: Visual Analysis of Travel Route Recommendations

Dawei Chen^{*†}, Dongwoo Kim^{*}, Lexing Xie^{*†}, Minjeong Shin^{*}, Aditya Krishna Menon^{†*},
Cheng Soon Ong^{†*}, Iman Avazpour[‡], John Grundy[‡]

^{*}The Australian National University, [†]Data61/CSIRO, [‡]Deakin University
{u5708856,dongwoo.kim,lexing.xie,minjeong.shin,aditya.menon,chengsoon.ong}@anu.edu.au
{iman.avazpour,j.grundy}@deakin.edu.au

ABSTRACT

We present an interactive visualisation tool for recommending travel trajectories. This system is based on new machine learning formulations and algorithms for the sequence recommendation problem. The system starts from a map-based overview, taking an interactive query as starting point. It then breaks down contributions from different geographical and user behavior features, and those from individual points-of-interest versus pairs of consecutive points on a route. The system also supports detailed quantitative interrogation by comparing a large number of features for multiple points. Effective trajectory visualisations can potentially benefit a large cohort of online map users and assist their decision-making. More broadly, the design of this system can inform visualisations of other structured prediction tasks, such as for sequences or trees.

KEYWORDS

Route Visualisation, Travel Recommendation, Learning to rank

1 INTRODUCTION

Sequence recommendation has emerged as an important framework for modelling diverse problems such as travel route and music playlist recommendation [3]. Unlike classical ranking algorithms where items are considered independently [7], a sequence recommendation algorithm requires modelling a structure between items and suggests a set of items as a whole. For example, consider recommending a trajectory of *points-of-interest* (POIs) in a city to a visitor. While a classical ranking algorithm can learn a user's preference for each individual location, it may ignore the distances between them and could suggest a longer trajectory than is optimal. Several sequence recommendation algorithms have been proposed to solve this problem and demonstrated superior performance compared to classical ranking algorithms [3, 9]. Nonetheless, recommendation algorithms for sequences and trajectories [1, 3] have many components and can be difficult for a user to understand. This is part of the general challenge of introducing transparency and accountability for machine learning algorithms [4].

In this paper, we tackle the problem of sequence visualisation, specifically focussing on travel routes recommendation. A travel route is a sequence of POIs, and the sequence recommendation problem can be formulated as a structured prediction problem [3]. Based on a diverse set of features for individual and pairs of POIs, we train the prediction model with trajectory data extracted from geo-tagged photos taken in Melbourne [1]. To visualise the suggested routes, we develop a novel tool that efficiently displays multiple suggested routes, which helps users understand the process behind

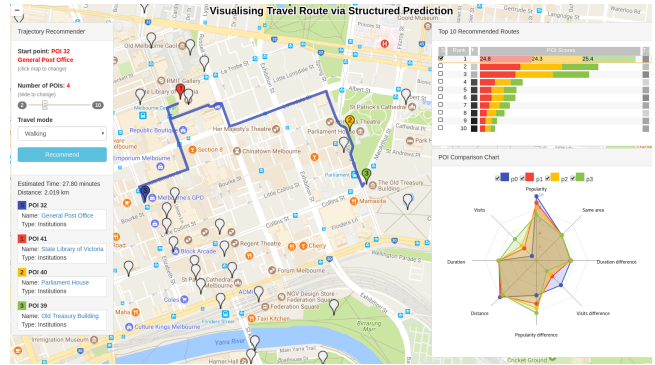


Figure 1: Travel route visualisation system¹. Given a starting POI and the number of POIs to be visited, the system recommends multiple routes from travel history of tourists. Shown above: recommendation in central Melbourne.

the recommendations. Specifically, our system decomposes a total score of each route into a set of features and their corresponding scores, and shows the total score as a stacked bar plot of the features. The system also visualises the differences between POIs in a single route to show how POIs in that route can exhibit vast diversity.

This visualisation helps tourists who want diverse experiences by choosing the best route among multiple recommendations. Generalising to a broader class of routes, such a visualisation could also help users of online mapping apps to make decisions on suggested travel routes, such as by trading off distance, traffic, and scenery.

2 TRAVEL ROUTE RECOMMENDATION

The travel route recommendation problem involves a set of POIs in a city. Given a trajectory query $\mathbf{x} = (s, l)$, comprising a start POI s and trip length l , the goal is to suggest one or more sequences of POIs that maximise some notion of utility.

Following [3], we first cast travel recommendation as a structured prediction problem, which allows us to leverage the well-studied literature of structured SVMs (SSVM) [6]. From a visualisation perspective, an advantage of the SSVM is the explicit representation of feature scores in its final decision process. Specifically, we can disassemble the final score of a route into feature scores of each POI and the transition between two adjacent POIs. We use hand-crafted POI features such as the category, popularity, and average visit duration of previous tourists. We also crafted transition features such as the distance and neighbourhood of two POIs to maximise the interpretability of the outcome.

¹<http://www.pathrec.ml>




Figure 2: Visualisation of POI and transition scores for top 10 recommended routes. Each bar from left to right represents a relative score of each POI or transition along the route. The length of stacked bars represents the total score of the suggested route.

3 VISUALISATION

Our goal is to design an interactive visualisation system on top of the structured prediction framework. Figure 1 shows the overview of a live demo system, which consists of five major components: a map to display the suggested routes, an input box for user query (upper left), a stacked score of routes (upper right), a POI list box (lower left), and a radar chart to compare features of multiple POIs (lower right). The role and the construction of the four major components, besides the main map, are as follows:

Query input: A query consists of a starting POI and a trip length. Users can choose the starting POI by clicking icons on the map and can adjust the slide to set the trip length. In addition, three different travelling modes (e.g. bicycling, driving and walking) are supported, and we optimise the suggested routes for each mode.

Route score visualisation: The SSVM evaluates relevance scores of POIs and transitions in a candidate route to the given query and uses the sum of the relevance scores to determine the ranks of the routes. To visualise the POI and transition scores, we adopt a stacked bar representation [5], designed to support the visualisation of multi-attribute ranking. In Figure 2, the system decompose the scores of top 10 recommended routes into POI and transition scores via the stacked bar representation, where the size of each bar is proportional to the relevance score of the corresponding POI and transition in the route. Note that the POI and transition scores are scaled differently to support better visual discrimination². For a seamless match between a route on the map and the corresponding POI scores in the bar plot, we use the same colour for both POI score and POI icon on the map. We also allow users to select multiple rows to visualise the corresponding routes on the map.

POI list: The POI list box provides the list of POI names and categories along the recommended route. The list is sorted according to the suggested visiting order, and again, the same POI colour is used to match the corresponding POI on the map. On top of the list, the system also provides an estimated travel time and total distance of the route. The POI list box is updated whenever a user selects a different route or the system makes a new recommendation. If more than one route is selected, the system displays the information of the most recent chosen route.

POI feature visualisation: We further provide a radar chart to analyse the variation between POIs in a single route. For example, in Figure 3, we compare two POIs (*Melbourne Aquarium* and *Queen Victoria Market*) in terms of POI features and their importance in




Figure 3: POI feature comparison between *Melbourne Aquarium* and *Queen Victoria Market*: the former scores higher on *Popularity* and *Visits difference* features whereas the latter scores higher on *Visits* and *Popularity difference* features.

the suggested route. The radar chart shows the corresponding POI feature scores when a user selects a route. In particular, the user can check/uncheck any POI in the selected route, and the feature scores of all checked POIs will be shown in the chart.

4 CONCLUSION

In this demonstration, we detail an interactive route analyser which helps the interaction between users and a route recommendation system. The system benefits from the explicit feature construction of the structured prediction model, and visualises recommended routes in terms of information on both the routes and the POIs.

Acknowledgements. We thank the National Computational Infrastructure (NCI), supported by the Australian Government, for computational resources. This work is supported in part by the Australian Research Council via Discovery Project program DP140102185.

REFERENCES

- [1] Dawei Chen, Cheng Soon Ong, and Lexing Xie. 2016. Learning Points and Routes to Recommend Trajectories. In *Proceedings of the 25th ACM International on Conference on Information and Knowledge Management*. ACM, 2227–2232.
- [2] Dawei Chen, Cheng Soon Ong, and Lexing Xie. 2016. Learning Points and Routes to Recommend Trajectories (*CIKM '16*).
- [3] Dawei Chen, Lexing Xie, Aditya Krishna Menon, and Cheng Soon Ong. 2017. Structured Recommendation. *CoRR* abs/1706.09067 (2017). <https://arxiv.org/abs/1706.09067>
- [4] FAT/ML. 2017. Fairness, Accountability, and Transparency in Machine Learning. (2017). <http://www.fatml.org>, retrieved May 2017.
- [5] Samuel Gratzl, Alexander Lex, Nils Gehlenborg, Hanspeter Pfister, and Marc Streit. 2013. Lineup: Visual analysis of multi-attribute rankings. *IEEE transactions on visualization and computer graphics* 19, 12 (2013), 2277–2286.
- [6] Thorsten Joachims, Thomas Hofmann, Yisong Yue, and Chun-Nam Yu. 2009. Predicting structured objects with support vector machines. *Commun. ACM* 52, 11 (2009), 97–104.
- [7] Yehuda Koren, Robert Bell, and Chris Volinsky. 2009. Matrix factorization techniques for recommender systems. *Computer* 42, 8 (2009).
- [8] Ching-Pei Lee and Chuan-bi Lin. 2014. Large-scale linear rankSVM. *Neural computation* 26, 4 (2014), 781–817.
- [9] Kwan Hui Lim, Jeffrey Chan, Christopher Leckie, and Shanika Karunasekera. 2015. Personalized tour recommendation based on user interests and points of interest visit durations (*IJCAI '15*).
- [10] N.B. Robbins. 2005. *Creating more effective graphs*. Wiley-Interscience.
- [11] Ben Shneiderman. 1996. The eyes have it: A task by data type taxonomy for information visualizations. In *Visual Languages, 1996. Proceedings., IEEE Symposium on*. IEEE, 336–343.
- [12] Wikipedia. 2017. Parallel coordinates. https://en.wikipedia.org/wiki/Parallel_coordinates. (2017). Accessed: July 2017.

²See Appendix for details at <http://arxiv.org/abs/1707.01627>

A ALTERNATIVE APPROACHES TO TRAJECTORY RECOMMENDATION

A number of approaches have been proposed to solve the trajectory recommendation problem. Lim et al. [9] formulated an optimisation problem inspired by the travelling salesman problem, and Chen et al. [2] proposed to learn a RankSVM [8] model to rank POIs with respect to a query, in particular, the training objective is

$$\min_{\mathbf{w}} \frac{1}{2} \mathbf{w}^\top \mathbf{w} + C \cdot \sum_{i=1}^n \sum_{(p,p') \in \mathcal{R}(\mathbf{x}^{(i)})} \max\left(0, 1 - \mathbf{w}^\top \left(\Phi(\mathbf{x}^{(i)}, p) - \Phi(\mathbf{x}^{(i)}, p')\right)\right)^2, \quad (1)$$

where \mathbf{w} denotes the model parameters, Φ is a query-POI feature mapping, $C > 0$ is a regularisation constant, and $\mathcal{R}(\mathbf{x})$ is the set of POI pairs (p, p') such that p is ranked above p' , e.g. POI p is observed more often than POI p' , with respect to query \mathbf{x} . Lastly, the top-ranked POIs with respect to the given query were taken to form a trajectory.

Instead of ranking POIs, a Markov chain was learned from routes in the travel history [2], and recommending a trajectory is achieved with either the classic Viterbi algorithm or an integer linear programming when repeated POI visits are discouraged. Further, Chen et al. [2] proposed to combine the POI ranks learned by RankSVM (1) and the transition preferences learned by a Markov chain using a heuristic which traded off between point affinity and transition preference.

B SCORE NORMALISATION FOR VISUALISATION

We visualise multiple suggested trip by ranking them according to their scores from SSVM, in addition, we leverage the linear form of SSVM that score of a trajectory is the sum of point and transition scores in the trip. As shown in Figure 2, scores of both POIs and transitions in a trip are visualised in a stacked bar plot, which helps users of the system to better understand what contributes to a suggested trip.

To support better visual discrimination, we perform linear scaling on scores of suggested trajectories, scores of POIs and transitions.

B.1 Scaling overall trajectory scores

Specifically, we scale the trajectory scores for the top 10 suggested routes such that:

- The first trajectory scores 100,
- and the last (i.e. the 10-th) trajectory scores 10.

B.2 Displaying unary and pairwise scores

We further scale the POI scores using the same scaling parameters as that of trajectory scores, however, to visualise the transition scores in the stacked bar plot, we perform another linear scaling as the scores of POIs and transitions are in quite different range, in particular, the transition scores are scaled linearly into the range [0.1, 1].

B.3 Displaying POI features

Besides scores of POIs and transitions of a suggested trip, we also visualise the scores for individual POI features (e.g. popularity, visit duration, category) in a radar chart, as shown in Figure 3. We perform another linear scaling such that scores of POI features are in the range [1, 10].

C CHOICES ON VISUALISATION PARADIGMS

We want to visualise a ranked list of suggested trips instead of a single choice, further, we would like to break down the score of a trip into the contributions of its unary and pairwise components. Following the Visual Information-Seeking Mantra [11], we made the following design choices for this visualisation system:

Overview first. For a given query, as shown in Figure 1, the visualisation system shows the top-ranked trip on a map. The other components of the system, as described in Section 3, can display more details of these recommendations after a user zooms in, as described below.

Zoom and filter. The system shows the scores of top-10 suggested trajectories and the scores of the corresponding POIs and transitions in a stacked bar plot, as shown in Figure 2. Further, a POI list box at the lower left side of the system interface displays the information of POIs including names and categories of the selected trip, as described in Section 3. Lastly, at the lower right side, contributions of individual features of all POIs in the selected trip are compared, and Figure 3 shows an example which compares individual feature scores between two POIs.

Details-on-demand. The system also support display and compare multiple suggested trajectories on map, this is achieved by select the trips at the left side of the stacked bar plot (Figure 2). All selected trips will be displayed on the map, which helps user choose the best trip among the multiple recommendations. Moreover, comparison of individual features between multiple POIs in a trip is also supported, in the radar chart, users can check one or more POIs in the most recent choose trip, the feature scores of all checked POIs will then be displayed and compared, using the same colour as the POI icon on the map.

We note that to visualise POI feature contributions, instead of radar chart, parallel coordinates [12] is another viable option, although both approaches have their own advantages and drawbacks [10].