Utilizing Language Models for Tour Itinerary Recommendation

Ngai Lam Ho¹, Kwan Hui Lim¹

¹Information Systems Technology and Design Pillar, Singapore University of Technology and Design

Abstract

Tour itinerary recommendation involves planning a sequence of relevant Point-of-Interest (POIs), which combines challenges from the fields of both Operations Research (OR) and Recommendation Systems (RS). As an OR problem, there is the need to maximize a certain utility (e.g., popularity of POIs in the tour) while adhering to some constraints (e.g., maximum time for the tour). As a RS problem, it is heavily related to problem or filtering or ranking a subset of POIs that are relevant to a user and recommending it as part of an itinerary. In this paper, we explore the use of language models for the task of tour itinerary recommendation and planning. This task has the unique requirement of recommending personalized POIs relevant to users and planning these POIs as an itinerary that satisfies various constraints. We discuss some approaches in this area, such as using word embedding techniques like Word2Vec and GloVe for learning POI embeddings and transformer-based techniques like BERT for generating itineraries.

Keywords

Tour Recommendation, Itinerary Planning, Language Models, Natural Language Processing

1. Introduction

The tour itinerary recommendation is a popular and challenging problem, with significant impact for tourism and other domains such as transportation and logistics [1]. The tour itinerary recommendation problem has garnered immense interest in both academia and industry. This problem contains both aspects of an recommendation problem as well as a planning problem. From the recommendation perspective, there are elements of top-K item recommendation and learning to rank, where we aim to recommend a subset of most relevant POIs to a user in the form of an itinerary. From the operation research perspective, it is akin to a constrained optimization problem, where we need to maximize the utility that a user obtains from the planned itinerary while ensuring that the itinerary adheres to certain time and location constraints.

In this paper, we discuss the tour itinerary recommendation problem from both perspectives of a recommendation and operation research problem, particularly on how recent advances on natural language processing have advanced research in this area. In particular, we discuss how language models have been adapted for the task of tour itinerary recommendation such as how word embedding techniques like Word2Vec and GloVe are used for POI representation learning and transformer-based models like BERT are used for next POI and itinerary recommendation. Figure 1 provides an overview of potential applications of such a framework.

^{© 0000-0003-4768-2208 (}N. L. Ho); 0000-0002-4569-0901 (K. H. Lim) © 2022 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)

PMAI'23: 2nd International Workshop on Process Management in the AI Era, August 29, 2023, Macau, S.A.R.

ngailam ho@mymail.sutd.edu.sg (N.L. Ho); kwanhui lim@sutd.edu.sg (K.H. Lim)



Figure 1: Application of Language Models for Tour Recommendation, including (a) Representation of POIs and Itineraries as Words and Sentences; (b) Feeding of inputs in the form of (a) into various language models; (c) Potential outputs and downstream applications.

2. Tour Itinerary Recommendation Problem

There are various variants of the tour itinerary recommendation and a popular formulation is based on the Orienteering problem [2, 3], which we briefly discuss.

Given a set of POIs P, our main objective is to plan for a tour itinerary $T = (p_1, ..., p_N)$ that:

$$Max \sum_{p_i}^{P} \sum_{p_j}^{P} x_{p_i, p_j} S_u(p_i) \tag{1}$$

where $x_{p_i,p_j} = 1$ is the itinerary includes a travelling path from POI p_i to p_j , and $x_{p_i,p_j} = 0$ otherwise. $S_u(p_i)$ is the utility score that a user u benefits from based on this tour itinerary T.

Variants of this problem differ in various aspects, mainly in terms of the definition of the utility score used and the types of constraints being implemented. For example, the utility score is typically a global score such as POI popularity for the OR community, while a more personalized interest or relevance score is used for the RS community. Similarly, the two communities might differ in terms of the type and range of constraints from simple ones such as having a fixed number of POIs in a tour itinerary to more diverse constraints such as the need to start/end the tour itinerary at specific locations, complete the tour itinerary within specific time limits or budget constraints, and other unique considerations.

3. Language Models and Tour Recommendation

Natural Language Processing (NLP) and particularly large language models have seen rapid progress in recent years with developments ranging from word embedding techniques like Word2Vec [4] and GloVe [5] to the more recent foundation models like BERT [6], GPT [7] and their variants.

Word embedding techniques aim to find appropriate dense vector representation at the word level, by using techniques such as the Continuous Bag of Words and Skip-gram Models that are used in Word2Vec. Both model uses a simple neural network with a single hidden layer and differ in terms of the task. Continuous Bag of Words is used to predict a single target word given the context of its neighbouring words, where Skip-gram does the opposite and tries to predict the neighbouring words that surround a given input word. Thereafter, the weights learned by the hidden layer are used as the word embedding for the specific word, where the number of neurons/nodes in the hidden layer corresponds to the dimension size of the word embedding.

Foundation models like BERT and GPT utilizes the transformer architecture with self-attention and are now commonly used for many downstream NLP tasks from standard text classification to text generation. Models like BERT and its variants uses a self-supervised approach to training and typically use the Mask Language Modelling and Next Sentence Prediction tasks on top of the Transformer architecture. Given a sentence, the Mask Language Modelling task involves hiding a subset of words in the sentence and training the model to predict those words. For Next Sentence Prediction, it is to predict the following sentence given the earlier sentence.

POI Embedding Models. Word embedding models have increasingly been used for generating POI embedding. To adapt language models for POI representation learning, the set of POIs in a city can be treated as the vocabulary of words where each POI is akin to an individual word or token. Similar to sentences in NLP, past itineraries or sequences of POI visits are used as a proxy of sentences made up of a series of words. Thereafter, models such as Continuous Bag of Words and Skip-gram are used to learn the vector representation of POIs. Similar models and variants that are adapted for learning latent attributes relating to geographic and temporal factors have been proposed in recent years [8, 9, 10].

Transformer-based Models. More recently, Transformer-based models have gained popularity for not just NLP tasks but increasingly for various types of next POI prediction and tour itinerary recommendation tasks. Similarly, past itineraries comprising sequences of POI visits can be modelled in the same context as sentences that are made up of a series of words. For example, [11] used the BERT model trained on past POI visit sequences, coupled with an iterative process for generating intermediate POIs for recommending tour itineraries. Others like [12] have used Transformers for the POI recommendation problem with joint training on the task of next POI prediction and queuing time prediction. More generally, Transformer-based architectures have seen numerous applications for various sequence-related recommendation tasks [13, 14, 15].

4. Conclusion

In this paper, we discussed the problem of tour itinerary recommendation and highlighted its relation to both fields of OR and RS in its consideration of trip constraints and user-relevant recommendation. Following which, we provided an overview of popular techniques used in a variety of NLP tasks and discussed how these NLP techniques, such as word embedding and Transformers, have been adapted to the tour itinerary recommendation task.

Acknowledgments. This research is funded in part by the Singapore University of Technology and Design under grant RS-MEFAI-00005-R0201.

References

- K. H. Lim, J. Chan, S. Karunasekera, C. Leckie, Tour recommendation and trip planning using location-based social media: A survey, Knowledge and Information Systems 60 (2019) 1247–1275.
- [2] B. L. Golden, L. Levy, R. Vohra, The orienteering problem, Naval Research Logistics (NRL) 34 (1987) 307–318.
- [3] A. Gunawan, H. C. Lau, P. Vansteenwegen, Orienteering problem: A survey of recent variants, solution approaches and applications, European Journal of Operational Research 255 (2016) 315–332.
- [4] T. Mikolov, I. Sutskever, K. Chen, G. S. Corrado, J. Dean, Distributed representations of words and phrases and their compositionality, Proceedings of NeurIPS'13 26 (2013).
- [5] J. Pennington, R. Socher, C. D. Manning, Glove: Global vectors for word representation, in: Proceedings of EMNLP'14, 2014, pp. 1532–1543.
- [6] J. Devlin, M.-W. Chang, K. Lee, K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, in: Proceedings of NAACL'19, 2019, pp. 4171– 4186.
- [7] T. Brown, B. Mann, N. Ryder, M. Subbiah, J. D. Kaplan, P. Dhariwal, A. Neelakantan, P. Shyam, G. Sastry, A. Askell, et al., Language models are few-shot learners, Proceedings of NeurIPS'20 33 (2020) 1877–1901.
- [8] S. Feng, G. Cong, B. An, Y. M. Chee, Poi2vec: Geographical latent representation for predicting future visitors, in: Proceedings of AAAI'17, volume 31, 2017.
- [9] B. Chang, Y. Park, D. Park, S. Kim, J. Kang, Content-aware hierarchical point-of-interest embedding model for successive poi recommendation., in: Proceedings of IJCAI'18, volume 2018, 2018, p. 27th.
- [10] N. L. Ho, K. H. Lim, User preferential tour recommendation based on poi-embedding methods, in: Proceedings of IUI'21 Companion, 2021, pp. 46–48.
- [11] N. L. Ho, K. H. Lim, Poibert: A transformer-based model for the tour recommendation problem, Proceedings of BigData'22 (2022) 5925–5933.
- [12] S. Halder, K. H. Lim, J. Chan, X. Zhang, Poi recommendation with queuing time and user interest awareness, Data Mining and Knowledge Discovery 36 (2022) 2379–2409.
- [13] F. Sun, J. Liu, J. Wu, C. Pei, X. Lin, W. Ou, P. Jiang, Bert4rec: Sequential recommendation with bidirectional encoder representations from transformer, in: Proceedings of CIKM'19, 2019, pp. 1441–1450.
- [14] L. Wu, S. Li, C.-J. Hsieh, J. Sharpnack, Sse-pt: Sequential recommendation via personalized transformer, in: Proceedings of RecSys'20, 2020, pp. 328–337.
- [15] Q. Chen, H. Zhao, W. Li, P. Huang, W. Ou, Behavior sequence transformer for e-commerce recommendation in alibaba, in: Proceedings of the 1st International Workshop on Deep Learning Practice for High-Dimensional Sparse Data, 2019, pp. 1–4.