

Out-of-Town Recommendation with Travel Intention Modeling

Haoran Xin,^{1, 2} Xinjiang Lu,^{2, 3*} Tong Xu,¹ Hao Liu,^{2, 3} Jingjing Gu,⁴ Dejing Dou,^{2, 3} Hui Xiong⁵

¹University of Science and Technology of China

²Business Intelligence Lab, Baidu Research

³National Engineering Laboratory of Deep Learning Technology and Application, China

⁴Nanjing University of Aeronautics and Astronautics

⁵Rutgers University

xinhaoran@mail.ustc.edu.cn, {luxinjiang, liuhao30, doudejing}@baidu.com, tongxu@ustc.edu.cn, gujingjing@nuaa.edu.cn, hxiong@rutgers.edu

Abstract

Out-of-town recommendation is designed for those users who leave their home-town areas and visit the areas they have never been to before. It is challenging to recommend Point-of-Interests (POIs) for out-of-town users since the out-of-town check-in behavior is determined by not only the user's home-town preference but also the user's travel intention. Besides, the user's travel intentions are complex and dynamic, which leads to big difficulties in understanding such intentions precisely. In this paper, we propose a TRAVEL-INTENTION-aware Out-of-town Recommendation framework, named TRAINOR. The proposed TRAINOR framework distinguishes itself from existing out-of-town recommenders in three aspects. First, graph neural networks are explored to represent users' home-town check-in preference and geographical constraints in out-of-town check-in behaviors. Second, a user-specific travel intention is formulated as an aggregation combining home-town preference and generic travel intention together, where the generic travel intention is regarded as a mixture of inherent intentions that can be learned by Neural Topic Model (NTM). Third, a non-linear mapping function, as well as a matrix factorization method, are employed to transfer users' home-town preference and estimate out-of-town POI's representation, respectively. Extensive experiments on real-world data sets validate the effectiveness of the TRAINOR framework. Moreover, the learned travel intention can deliver meaningful explanations for understanding a user's travel purposes.

Introduction

Point-of-Interest (POI) recommendation is an important task in location-based services (LBS), which tends to act as a more pivotal part in people's daily life. Recently, since the POI check-in data having accumulated rapidly over time, a more refined recommendation problem, out-of-town recommendation, is coming into focus. To be specific, out-of-town recommendations are designed for those users who travel from their home-town areas to out-of-town areas they have seldom been to before.

*Corresponding author. This work was done when Haoran Xin was an intern at the Baidu Research.
Copyright © 2021, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

Out-of-town recommendation problem suffers from the cold-start issue a lot due to the insufficiency of out-of-town check-ins (FERENCE, YE, and LEE 2013). Traditional POI recommender systems (POI RSs) fail to make appropriate recommendations to tackle such severe cold start issues. The reasons are: 1) Individual's home-town preferences cannot be used for out-of-town recommendations directly due to the gap between home-town preferences and out-of-town behaviors (i.e. interest drifts); and 2) The travel intention, which tends to affect the out-of-town check-in behaviors, is often ignored in these POI RSs.

In the literature, some research efforts have been made to attack the out-of-town recommendation problem. For instance, (PHAM, LI, and CONG 2017) recommends out-of-town region of POIs instead of individual POIs by exploiting the proximity of human mobility. (FERENCE, YE, and LEE 2013) proposes a recommender for out-of-town users by taking into account user preference, social influence and geographical proximity. Besides, some researchers have also paid attentions to interest drifts when addressing the out-of-town recommendation problem (YIN ET AL. 2014, 2016; WANG ET AL. 2017). However, none of these approaches comprehensively integrate users' preferences, interest drifts and complex travel intentions as a whole.

To this end, in this paper, we propose a TRAVEL-INTENTION-aware Out-of-town Recommendation framework, named TRAINOR. Specifically, we first devise a user's preference representation module based on Gated Graph Neural Network (G-GNN) to explore the underlying structural information encoded in user's home-town check-ins. After being aggregated via an attention network, the user's home-town preference is further transferred into out-of-town preference through a non-linear mapping function, i.e. multi-layer perceptron (MLP). In this way, the interest drifts from home-town to out-of-town can be captured directly. Besides, we devise a travel intention discovery module by developing a Neural Topic Model (NTM) followed by user-specific travel intention aggregation. In particular, we assume that each out-of-town check-in activity can be drawn from a latent topic mixture which can be further generated by Gaussian Softmax construction, then we adopt variational inference to uncover users' generic travel intention without extra su-

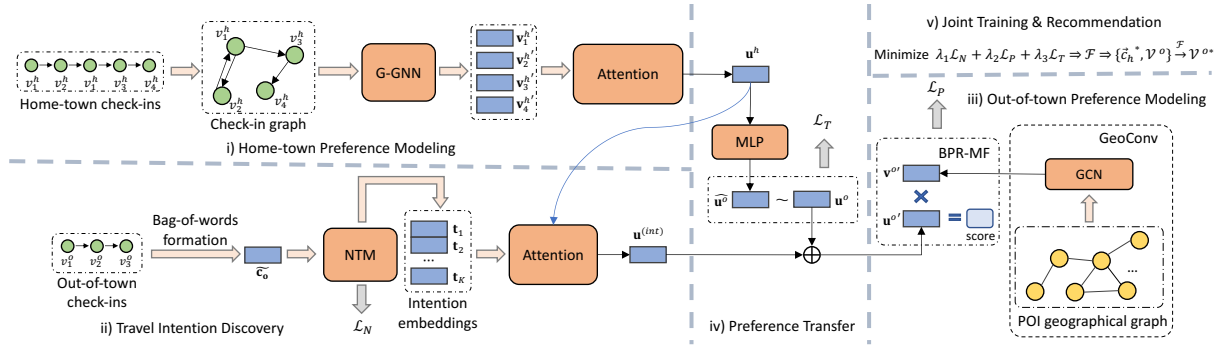


Figure 1: The overview of TRAINOR framework.

pervision. Moreover, the aforementioned user’s home-town preference is integrated into the disclosed generic travel intention to generate user-specific travel intention via another attention network. In addition, we represent user’s out-of-town preference by exploiting a matrix factorization (MF) approach and enrich such out-of-town preference by taking into account the geographical proximity among out-of-town POIs. Finally, a joint learning method is employed in an end-to-end manner to yield the trained recommender. To sum up, our major contributions are as follows:

- We study the out-of-town recommendation problem by modeling user’s complex travel intention.
- We devise a framework TRAINOR which is able to capture the user’s home-town preference, user’s interest drift from home-town to out-of-town, out-of-town geographical influence and user’s travel intention comprehensively.
- We demonstrate the effectiveness of TRAINOR quantitatively and qualitatively through extensive experiments.

Problem Definition

In this section, we formally define the out-of-town recommendation problem. We start by defining several concepts.

Definition 1 (POI) A POI is a spatial item related to a geographical location. We use v to represent a POI identifier.

Definition 2 (Check-in) A user’s check-in activity c is represented by a three-tuple (u, t, v) which indicates that a user u visits POI v at timestamp t .

Definition 3 (User Home-Town) Given a user u , we denote a region \tilde{r}_u as the user’s home-town where the user lives for a period of time, say, 6 months.

Definition 4 (Travel Behavior) Given a user u , his/her travel behavior is represented by a five-tuple $\tau = (u, \vec{c}_h, \vec{c}_o, \tilde{r}_u, r_o)$ which indicates that the user u travels from his/her home-town \tilde{r}_u to out-of-town r_o and leaves check-in records in both home-town and out-of-town, which are represented by \vec{c}_h and \vec{c}_o , respectively.

When a user u travels from his/her home-town \tilde{r}_u to an out-of-town r_o , we take u as an out-of-town user and aim to recommend a list of POIs located at r_o that u may be interested in. Formally, we have the following problem statement:

Problem 1 (Out-of-town Recommendation) Given a set of users \mathcal{U} who live in \tilde{r} , a target region r_o , a set of out-of-town POIs \mathcal{V}^o in r_o , and the travel behavior records \mathcal{T} generated by \mathcal{U} when traveling from \tilde{r} to r_o , learn a function $\mathcal{F}(\cdot)$ by exploring \mathcal{T} and \mathcal{V}^o . Then, recommend a list of POIs $\mathcal{V}^{o*} \subset \mathcal{V}^o$ to a new coming user $u^* \notin \mathcal{U}$ given his/her home-town check-ins \vec{c}_h^* observed in \tilde{r} : $\{\vec{c}_h^*, \mathcal{V}^o\} \xrightarrow{\mathcal{F}} \mathcal{V}^{o*}$.

The Proposed Approach

Framework Overview

We first present the overview of TRAINOR framework which is illustrated in Fig. 1. The TRAINOR framework consists of five components:

- **Home-town preference modeling** takes user’s home-town check-ins as input and assigns a d -dimensional embedding to each of the visited POIs. Then the user’s home-town preference is encoded and aggregated by adopting G-GNN model and attention network.
- **Travel intention discovery** takes the user’s visited POIs in out-of-town as input in bag-of-words, and then an NTM model takes such input to discover the generic travel intention. Afterward, another attention network is adopted to summarize user-specific intention by integrating discovered intention and user’s home-town preference.
- **Out-of-town preference modeling** assigns another two d -dimensional embeddings to each user and out-of-town POI, and utilizes MF to learn the latent representations of users and POIs. Moreover, to model the geographical influence of POIs, a GeoConv is explored to process the geo-information bundled with POIs.
- **Preference transfer** receives the home-town preference embedding and captures the non-linear relationship from home-town to out-of-town via an MLP.
- **Model learner** jointly minimizes the intention inference loss, preference estimation loss, and preference transfer loss to output the trained recommender \mathcal{F} .

Home-town Preference Modeling

To encode users’ home-town preference, we represent the structural information with the G-GNN model (Wu et al. 2019; Li et al. 2015).

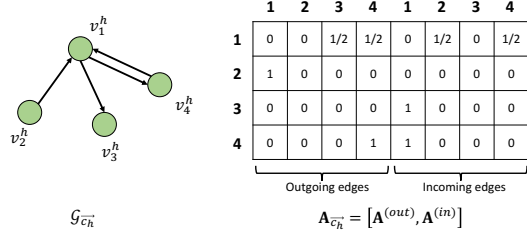


Figure 2: An illustration of check-in graph $\mathcal{G}_{\vec{c}_h}$ and the construction of corresponding $\mathbf{A}_{\vec{c}_h} = [\mathbf{A}^{(out)}, \mathbf{A}^{(in)}]$.

Given a user u and his/her home-town check-ins \vec{c}_h , we first build a directed graph $\mathcal{G}_{\vec{c}_h} = (\mathcal{V}_{\vec{c}}, \mathcal{E}_{\vec{c}})$, where $\mathcal{V}_{\vec{c}}$ denotes the set of home-town check-ins and each pair of adjacent check-ins is represented by $(v_{i-1}^h, v_i^h) \in \mathcal{E}_{\vec{c}} (v_i^h \in \vec{c}_h)$. Notably, duplicated pairs of spatial items may exist in \vec{c}_h , we normalize all the weights of the edges in $\mathcal{G}_{\vec{c}_h}$. Then, we construct the adjacent matrix $\mathbf{A}_{\vec{c}_h}$ (refer to Fig. 2). The matrix $\mathbf{A}_{\vec{c}_h} \in \mathbb{R}^{D_1 \times 2D_1}$ (NOTE: $D_1 = |\mathcal{V}_{\vec{c}}|$) determines how spatial items communicate with each other via user’s check-ins. Next, we assign a d -dimensional embedding \mathbf{v}_i^h to each vertex v_i^h in $\mathcal{G}_{\vec{c}_h}$ and feed the corresponding embeddings $\mathbf{V}^h = (\mathbf{v}_1^h, \mathbf{v}_2^h, \dots, \mathbf{v}_{D_1}^h)$ into the G-GNN. $\forall v \in \mathcal{V}_{\vec{c}}$, the network propagates as follows:

$$\mathbf{a}_v^{(t)} = \mathbf{A}_{v:}^T \left[\mathbf{v}_1^{h(t-1)}, \mathbf{v}_2^{h(t-1)}, \dots, \mathbf{v}_{D_1}^{h(t-1)} \right]^T + \mathbf{b}^g, \quad (1)$$

$$\mathbf{z}_v^{(t)} = \zeta \left(\mathbf{W}^z \mathbf{a}_v^{(t)} + \mathbf{U}^z \mathbf{v}_v^{h(t-1)} \right), \quad (2)$$

$$\mathbf{r}_v^{(t)} = \zeta \left(\mathbf{W}^r \mathbf{a}_v^{(t)} + \mathbf{U}^r \mathbf{v}_v^{h(t-1)} \right), \quad (3)$$

$$\widetilde{\mathbf{v}}_v^{h(t)} = \tanh \left[\overline{\mathbf{W}} \mathbf{a}_v^{(t)} + \overline{\mathbf{U}} \left(\mathbf{r}_v^{(t)} \odot \mathbf{v}_v^{h(t-1)} \right) \right], \quad (4)$$

$$\mathbf{v}_v^{h(t)} = (1 - \mathbf{z}_v^{(t)}) \odot \mathbf{v}_v^{h(t-1)} + \mathbf{z}_v^{(t)} \odot \widetilde{\mathbf{v}}_v^{h(t)}, \quad (5)$$

where $\mathbf{A}_{v:}$ are the two columns of blocks in $\mathbf{A}^{(out)}$ and $\mathbf{A}^{(in)}$ corresponding to v , and $\zeta(\cdot)$ is the sigmoid function. In particular, Eq. (1) is the step that passes information between different POIs based on $\mathcal{G}_{\vec{c}_h}$. Eqs. (2) to (5) are the update steps similar to GRU (Cho et al. 2014). The updated embeddings learned by G-GNN are denoted as $\mathbf{V}^{h'} = [\mathbf{v}_1^{h'}, \mathbf{v}_2^{h'}, \dots, \mathbf{v}_{D_1}^{h'}]$.

Furthermore, to summarize user’s home-town preference, we adopt an attention network as follows:

$$\alpha_i = \mathbf{q}^T \zeta \left(\mathbf{W}^p \mathbf{v}_i^{h'} + \mathbf{b}^p \right), \quad (6)$$

$$\mathbf{u}^h = \sum_{i=1}^{D_1} \alpha_i \mathbf{v}_i^{h'},$$

where $\mathbf{q} \in \mathbb{R}^d$ and $\mathbf{W}^p \in \mathbb{R}^{d \times d}$ weigh the home-town POIs, and \mathbf{u}^h is the user’s home-town preference embedding.

Travel Intention Discovery

Understanding travel intentions plays an important role in out-of-town recommendation. Inspired by (Miao, Grefenstette, and Blunsom 2017; Srivastava and Sutton 2017), we develop a Neural Topic Model (NTM) to uncover the inherent travel intentions without extra supervision.

Uncovering Generic Travel Intentions. Assume that each out-of-town check-in is generated by a latent topic mixture $\Theta \in \mathbb{R}^K$, which can be regarded as the generic travel intention mixture of users, where K denotes the number of generic intentions. Then, $\forall i (1 \leq i \leq K)$, we adopt an embedding $\mathbf{t}_i \in \mathbb{R}^d$ to represent the i -th travel intention. Afterward, given the out-of-town POI embedding matrix $\mathbf{E} \in \mathbb{R}^{|\mathcal{V}^o| \times d}$, the i -th generic out-of-town travel intention distribution over the out-of-town POIs, denoted as Φ_i , can be determined as follows:

$$\Phi_i = \text{softmax}(\mathbf{E} \mathbf{t}_i), \quad (7)$$

where $\Phi_i \in \mathbb{R}^{|\mathcal{V}^o|}$. Then we denote the whole out-of-town intention-POI distribution as $\Phi = (\Phi_1, \Phi_2, \dots, \Phi_K)^T$.

We assume that the distribution Θ can be generated by Gaussian Softmax construction. Let $\tilde{\mathbf{c}}_o \in \mathbb{R}^{|\mathcal{V}^o|}$ be the bag-of-words vector to represent the user’s out-of-town check-ins, then the generation of $\tilde{\mathbf{c}}_o$ can be conducted as follows:

- Draw a latent variable \mathbf{z} from a standard Gaussian distribution: $\mathbf{z} \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.
- Generate the out-of-town intention distribution Θ : $\Theta = \text{softmax}(F_\Theta(\mathbf{z}))$, where F_Θ is a fully connected layer.
- For the i -th POI in $\tilde{\mathbf{c}}_o$, draw a POI $v_i \sim \Phi^T \Theta$.

As shown above, we can find that $p(\mathbf{z}) = \mathcal{N}(\mathbf{0}, \mathbf{I})$. In order to make \mathbf{z} traceable, a variational posterior distribution is introduced as below:

$$q(\mathbf{z} | \tilde{\mathbf{c}}_o) = \mathcal{N}(\boldsymbol{\mu}, \boldsymbol{\sigma}^2), \quad (8)$$

where $\boldsymbol{\mu}, \boldsymbol{\sigma}^2$ are two prior parameters determined by the input bag-of-words vectors:

$$\boldsymbol{\mu} = F_\mu(F_{enc}(\tilde{\mathbf{c}}_o)), \quad (9)$$

$$\boldsymbol{\sigma}^2 = F_\sigma(F_{enc}(\tilde{\mathbf{c}}_o)),$$

where F_μ, F_σ are two multi-layer perceptrons (MLP) and F_{enc} is an encoder layer which accepts bag-of-words inputs extracted from out-of-town check-ins.

As the neural variational inference instructs, we would like to maximize the variational lower bound. Thus, the intention inference loss is defined as follows:

$$\mathcal{L}_N = - \sum_{u \in \mathcal{U}} \left[\mathbb{E}_{q(\mathbf{z} | \tilde{\mathbf{c}}_o)} \left(\tilde{\mathbf{c}}_o^T \log(\Phi^T \Theta) \right) + \mathbb{D}_{\text{KL}}(q(\mathbf{z} | \tilde{\mathbf{c}}_o) || p(\mathbf{z})) \right], \quad (10)$$

where \mathbb{D}_{KL} is the Kullback-Leibler divergence.

By optimizing the above loss, the generic travel intentions can be discovered without extra supervision.

Summarizing User-Specific Travel Intention. Previous works (Zeng et al. 2018; Wei and Mao 2019) have paid attention to integrating the topic knowledge with downstream tasks. Inspired by these, we further design an attention network to probe the dynamic travel intentions of users, which can explore intention knowledge according to user’s home-town preference.

Specifically, after the generic out-of-town intention $\mathbf{T} = (\mathbf{t}_1, \mathbf{t}_2, \dots, \mathbf{t}_K)^\top$ being acquired with NTM, we implement the attention network as follows:

$$\begin{aligned} \beta_i &= \text{softmax}(\mathbf{t}_i^\top \mathbf{W}^t \mathbf{u}^h), \\ \mathbf{u}^{(int)} &= \sum_{i=1}^K \beta_i \mathbf{t}_i, \end{aligned} \quad (11)$$

where $\mathbf{W}^t \in \mathbb{R}^{d \times d}$ is a trainable transition matrix. By fitting the user’s preference, the user-specific intention embedding $\mathbf{u}^{(int)}$ can be aggregated adaptively.

Out-of-town Preference Modeling

Geographical influence underlying out-of-town POIs is helpful in understanding users’ out-of-town check-in behaviors. On the other hand, with the logged travel records \mathcal{T} , we can further enrich the representations of out-of-town POIs via exploiting the interactions between POIs and users.

Specifically, we first assign another d -dimensional embedding to each out-of-town POI denoted as $\mathbf{v}^o \in \mathbb{R}^d$, and we have $\mathbf{V}^o = (\mathbf{v}_1^o, \mathbf{v}_2^o, \dots, \mathbf{v}_{D_2}^o)^\top$ where $\mathbf{V}^o \in \mathbb{R}^{D_2 \times d}$ (NOTE: $D_2 = |\mathcal{V}^o|$).

Then, we build an undirected graph $\mathcal{G}_{geo} = (\mathcal{V}^o, \mathcal{E}^o)$ based on the geographical relations among POIs, and the edge $e_{i,j}^o \in \mathcal{E}^o$ is defined as:

$$e_{i,j}^o = \exp(-\text{dist}(i, j)), \quad (12)$$

where $\text{dist}(\cdot, \cdot)$ denotes the distance between POI i and j . The adjacent matrix \mathbf{A}_{geo} can be constructed based on the edge constraints between each pair of out-of-town POIs.

Recently, GNN has been proved to be effective in modeling spatial data (Zhang et al. 2020; Li et al. 2020; Geng et al. 2019). To capture the relations among POIs in a spatial perspective, we employ the graph neural network (Kipf and Welling 2016) as below:

$$\mathbf{V}^{o'} = \text{ReLU}(\mathbf{A}_{geo} \mathbf{V}^o \mathbf{W}^c + \mathbf{b}^c), \quad (13)$$

where $\mathbf{W}^c \in \mathbb{R}^{d \times d}$ is a transition matrix and $\mathbf{b}^c \in \mathbb{R}^d$ is a bias vector. $\mathbf{V}^{o'} = (\mathbf{v}_1^{o'}, \mathbf{v}_2^{o'}, \dots, \mathbf{v}_{D_2}^{o'})^\top$ is the updated out-of-town POI embedding matrix, which encodes geographical influence of POIs.

Moreover, from the users’ point of view, we adopt the matrix factorization (MF) method to explore the interactions between users and POIs in out-of-town. In particular, we first assign a d -dimensional embedding, denoted by $\mathbf{u}^o \in \mathbb{R}^d$, to each of the users who left out-of-town check-ins.

Then, we aggregate the user’s out-of-town preference and travel intention:

$$\mathbf{u}^{o'} = \text{ReLU}(\mathbf{W}^f \text{concat}(\mathbf{u}^o, \mathbf{u}^{(int)}) + \mathbf{b}^f), \quad (14)$$

where $\mathbf{W}^f \in \mathbb{R}^{d \times 2d}$ is a transition matrix, $\mathbf{b} \in \mathbb{R}^d$ is a bias vector, and $\text{concat}(\cdot, \cdot)$ is a function concatenating its two input vectors.

Afterward, following the idea of MF that a user’s scores over POIs can be regarded as the inner product of the user’s latent embedding and the POIs’, we define the score of user i over out-of-town POI j as follows:

$$s(i, j) = (\mathbf{u}_i^{o'})^\top \mathbf{v}_j^{o'}. \quad (15)$$

At last, following the assumption of BPR (Rendle et al. 2012) that the observed items should be ranked higher than those unobserved, for each user u , we randomly select a fixed size of positive samples visited by u and their counterparts not checked in by u . Based on pairwise comparisons, the out-of-town preference loss is given by:

$$\mathcal{L}_P = - \sum_{u \in \mathcal{U}} \sum_{j \in \mathcal{E}_o} \sum_{k \notin \mathcal{E}_o} \log \zeta(s(i, j) - s(i, k)), \quad (16)$$

where \mathcal{E}_o comprises u ’s out-of-town check-ins.

Preference Transfer

Inspired by (Man et al. 2017), we adopt an MLP as the non-linear mapping function to transfer user’s home-town preference to out-of-town check-in behavior. We define the preference transfer loss as follows:

$$\mathcal{L}_T = \sum_{i \in \mathcal{U}} \|F_{tr}(\mathbf{u}_i^h) - \mathbf{u}_i^{o'}\|^2, \quad (17)$$

where F_{tr} is the MLP-based mapping function.

Joint Training and Recommendation

By combining the intention inference loss in Eq. (10), the preference loss in Eq. (16) and the transfer loss in Eq. (17), we can minimize the following composite loss function to jointly train our model in an end-to-end fashion:

$$\mathcal{L} = \lambda_1 \mathcal{L}_N + \lambda_2 \mathcal{L}_P + \lambda_3 \mathcal{L}_T, \quad (18)$$

where λ_1 , λ_2 and λ_3 are three hyper-parameters that control the respective contributions to the composite loss function.

After the parameters in our model are optimized, we can make recommendations for out-of-town users. Specifically, given a user $u^* \notin \mathcal{U}$ and his/her home-town check-ins, we first generate his/her affine out-of-town user preference by using the trained preference transfer:

$$\widehat{\mathbf{u}}_*^{o'} = F_{tr}(\mathbf{u}_*^h), \quad (19)$$

where \mathbf{u}_*^h is u^* ’s home-town preference embedding obtained from Eq. (6). Meanwhile, we can obtain his/her intention embedding $\mathbf{u}_*^{(int)}$ by using Eq. (11). Similar to Eq. (14), the travel intention embedding can be calculated as:

$$\widehat{\mathbf{u}}_*^{o'} = \text{ReLU}(\mathbf{W}^f \text{concat}(\widehat{\mathbf{u}}_*^{o'}, \mathbf{u}_*^{(int)}) + \mathbf{b}^f). \quad (20)$$

Then, with $\widehat{\mathbf{u}}_*^{o'}$ and $\mathbf{V}^{o'}$, we can estimate the score of user u^* over out-of-town POI j :

$$\widehat{s}(*, j) = (\widehat{\mathbf{u}}_*^{o'})^\top \mathbf{v}_j^{o'}. \quad (21)$$

Finally, we can pick the top-k out-of-town POIs based on the estimated scores as the recommendations for the out-of-town user u^* .

Dataset		# Users	# POIs	# Check-ins
BJ→SH	Beijing	10,776	2,111	127,528
	Shanghai		1,140	70,794
SH→HZ	Shanghai	19,997	3,415	263,158
	Hangzhou		1,203	116,475
GZ→FS	Guangzhou	12,788	4,228	220,006
	Foshan		1,225	57,229

Table 1: Basic description of datasets.

Experiments

Experimental Setups

Dataset. We chose three real-world travel behavior datasets including BJ→SH, SH→HZ and GZ→FS, to evaluate our approach. BJ→SH stands for traveling from Beijing to Shanghai, SH→HZ for Shanghai to Hangzhou and GZ→FS for Guangzhou to Foshan. The travel records of the above three datasets were generated between 07/01/2019 and 12/31/2019. To ensure the data quality, in each dataset, we filtered out the POIs that is visited less than 5 times. Besides, the users, whose home-town check-ins are less than 5 or out-of-town check-ins are less than 3, were eliminated. Then, we randomly split users following the proportions: 80%, 10%, and 10% to form a training set, a test set, and a validation set. The statistics of our dataset are given in Table 1. Notably, in our datasets, each user has only one travel record, which guarantees the fairness of our evaluations for out-of-town recommendation.

Evaluation Metrics. Since there are more than one out-of-town check-ins (i.e. multiple ground-truths) for each user in our dataset, we apply *Recall@k* (Rec@k) and *mean average precision* (MAP) to evaluate the performance of different recommender systems. The larger the values of the above metrics are, the better the models perform.

Baselines. We compared our approach with various baselines that could be used for out-of-town recommendation.

- **TOP** is a naive method which recommends the top-N frequently visited POIs in the target city.
- **UCF** is a user-based collaborative filtering method which recommends POIs for a target user in accordance with POI check-in behaviors of similar users.
- **BPR-MF** (Rendle et al. 2012) takes MF as the underlying predictor, which aims to factorize the user-POI matrix into the latent factors, and optimizes the MF by Bayesian Personalized Ranking (BPR). Recommendations are implemented based on the reconstruction of the matrix.
- **GRU4Rec** (Hidasi et al. 2015) utilizes RNNs to model users’ sequential check-ins. To make this method capable of our problem, we take the home-town check-ins as RNNs’ input, predict the out-of-town check-ins by utilizing the hidden state, and train the model by BPR.
- **SR-GNN** (Wu et al. 2019) utilizes GNNs to model the complex transitions of items. Similar to GRU4Rec, we regard each user’s home-town check-ins as a directed graph,

predict the out-of-town check-ins and train the model using BPR.

- **LA-LDA** (Yin et al. 2014) is a location-aware recommendation model which is suitable for out-of-town recommendation scenario. It takes personal interests and geographical gaps into consideration by exploiting POI co-visiting patterns.
- **EMCDR** (Man et al. 2017) is a cross-domain recommendation approach, which uses a multi-layer perceptron to capture the nonlinear mapping function across domains.

Moreover, to explore the respective contributions of different modules in our approach, we further come up with three variants of TRAINOR as follows:

- **TRAINOR-I**: this variant removes travel intention discovery module. As a result, it recommends only based on users’ preference.
- **TRAINOR-C**: this variant removes the GeoConv, such that the geographical influence of out-of-town POIs is neglected.
- **TRAINOR-IC**: this variant removes both travel intention discovery module and the geographical influence.

Implementations. The number d (i.e. the hidden size) was fixed to 128 for all latent representations. And, the number of layers in G-GNN was set to 1. In the travel intention discovery module, we set the topic number K as 15 for better explanation. In the joint training stage, we set $\lambda_1 = \lambda_2 = \lambda_3 = 1$ in Eq. (18). We used Adam optimizer to train our approach with an initial learning rate as 0.001 and an L2 regularization with weight 10^{-5} . When the quantity measures were evaluated, the test was repeated over 5 times using different data splits and the average was reported.

Experimental Results

Recommendation Performance. The performances of TRAINOR as well as its variants and the baselines are illustrated in ?? . Basically, TRAINOR consistently outperforms the baselines w.r.t. Rec@k. Regarding MAP, TRAINOR achieves best performance on GZ→FS dataset, and second best on BJ→SH and SH→HZ datasets.

In particular, UCF, LA-LDA and BPR-MF perform relatively worse. UCF and BPR-MF are two collaborative filtering algorithms for item recommendation, which cannot be directly applied to out-of-town recommendation due to the data scarcity issues. Though LA-LDA considers the geographical gaps, it is still insufficient to model the fine-grained personal interest drifts when the difference is big (e.g. cross city), which makes it less competitive for out-of-town recommendation.

TOP is not a personalized method and makes recommendations only based on the popularity of POIs according to history logs, yet has surprisingly better performances than some personalized approaches. The probable reason is that the out-of-town travel behaviors are usually dominated by tourism, which makes some hot attractions (e.g. famous landmarks) be frequently visited by the out-of-town users.

Methods	BJ→SH				SH→HZ				GZ→FS			
	Rec@10	Rec@20	Rec@30	MAP	Rec@10	Rec@20	Rec@30	MAP	Rec@10	Rec@20	Rec@30	MAP
LA-LDA	0.0160	0.0335	0.0417	0.0151	0.0008	0.0021	0.0028	0.0019	0.0020	0.0036	0.0057	0.0021
UCF	0.0443	0.0700	0.0935	0.1133	0.0628	0.0874	0.0981	0.2577	0.0386	0.0661	0.0800	0.1071
SR-GNN	0.1168	0.1807	0.2627	0.1071	0.2287	0.4550	0.5661	0.2013	0.0933	0.1670	0.2541	0.0566
BPR-MF	0.1768	0.2379	0.2844	0.0901	0.2812	0.3588	0.4116	0.1910	0.1642	0.2545	0.3173	0.0947
TOP	0.2062	0.3103	0.3818	0.1494	0.3713	0.4620	0.5176	0.2896	0.1964	0.2838	0.3483	0.1202
GRU4Rec	0.2091	0.3011	0.3763	0.1438	0.3619	0.4650	0.5150	0.2807	0.1789	0.2742	0.3422	0.1034
EMCDR	0.2163	0.3008	0.3649	0.1553	0.3772	0.4358	0.4732	0.3260	0.1928	0.2770	0.3368	0.1246
TRAINOR-IC	0.2029	0.2880	0.3513	0.1497	0.3679	0.4406	0.4963	0.3020	0.1937	0.2609	0.3178	0.1245
TRAINOR-I	0.2177	0.3084	0.3825	0.1543	0.3825	0.4624	0.5177	0.3016	0.2028	0.2841	0.3449	0.1266
TRAINOR-C	0.2233	0.3194	0.3955	0.1538	0.3914	0.4757	0.5300	0.2950	0.2032	0.2918	0.3569	0.1246
TRAINOR	0.2226	0.3198	0.3938	0.1541	0.3914	0.4768	0.5295	0.2955	0.2039	0.2922	0.3551	0.1246

Table 2: The overall performance of TRAINOR and baselines.

GRU4Rec and SR-GNN are two session-based deep recommender systems that take sequential and structural information into account, respectively. However, they also neglect the users’ interest drifts and context differences between home-town areas and out-of-town areas.

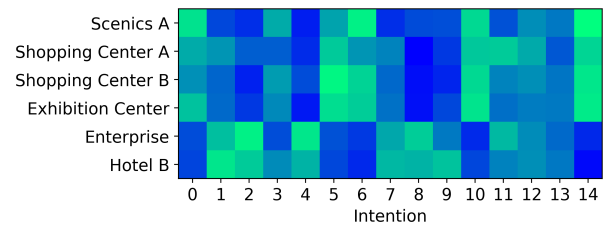
EMCDR is the state-of-the-art cross-domain recommendation framework. Because of its capability of non-linear mapping function that transfers features from the source domain to the target domain, EMCDR achieves the almost best ranking performance when making out-of-town recommendations. One possible reason is that in our TRAINOR framework, negative sampling strategy is not adopted in home-town preference modeling compared to EMCDR, which may lead to higher ranking of some negative items in home-town area and may have a negative impact on ranking performance, i.e. MAP. However, TRAINOR outperforms EMCDR with satisfactory margins in terms of Rec@k, which indicates that TRAINOR is more effective in retrieving relevant out-of-town POIs and more beneficial for out-of-town recommendation in practice.

Ablation Analysis. As for the variants of TRAINOR, we have the following main observations:

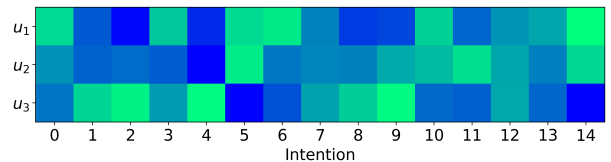
1) TRAINOR outperforms TRAINOR-I w.r.t. Rec@k, which indicates the effectiveness of taking into account users’ out-of-town intentions. Besides, the MAP slightly falls when travel intention discovery module is utilized. The reason might be that global signals such as intentions can become disturbance while the model is trying to put every item in a right ranking.

2) With comparing the results between TRAINOR-IC and TRAINOR-I, the removal of GeoConv decreases the performances on all metrics. However, we also find that with the existence of travel intention discovery module, GeoConv barely contributes to the performances, since GeoConv may result in overfitting as the learned intention embedding contains potential relations between POIs.

Case Study on Intention Discovery. We next present a case study on the discovered intentions to further evaluate TRAINOR framework. We randomly selected 3 recommended cases with promising Rec@30 (e.g. 0.67 for u_1 ,



(a) The distributions of visited POIs (over generic intentions).



(b) The weights of generic intentions for user-specific intentions.

Figure 3: The visualization of the case study.

0.33 for u_2 and 0.5 for u_3) from BJ→SH dataset. The out-of-town check-ins of these cases are illustrated in ?? . Besides, we visualized the POI-intention distributions (i.e. Φ^T) of some POIs visited by u_1 , u_2 and u_3 in Fig. 3(a), and the weights of generic intentions for the user-specific attentions (refer to Eq. (11)) related to these users in Fig. 3(b). The deeper the color, the greater the value.

As depicted in ?? , we can infer that u_1 traveled to Shanghai for vacation, u_2 for shopping and u_3 for business, respectively. Besides, based on the inference of users’ intentions, we can clearly tell that the difference between u_1 and u_2 is small, while, the difference between u_3 and u_1/u_2 is large, regarding the travel intention (refer to Fig. 3(b)).

Moreover, the intention distributions of the representative POIs are also distinguishable (refer to Fig. 3(a)). For example, the distributions of scenics and shopping centers have more weights on intention #2, #4, #8, etc., whereas the distributions of functional facilities (e.g. hotels and enterprises) have more weights on intention #5, #6, #10, #14, etc. Hence, we can conclude that the former set of intentions is more

User	Out-of-town check-ins
u_1	Scenics A, Scenics B, Art Gallery, Shopping Center A
u_2	Shopping Center B, Exhibition Center, Life Plaza, Shopping Center C, Hotel A
u_3	Enterprise, Hotel B, Hotel C

Table 3: Out-of-town check-ins of three selected users from test set.

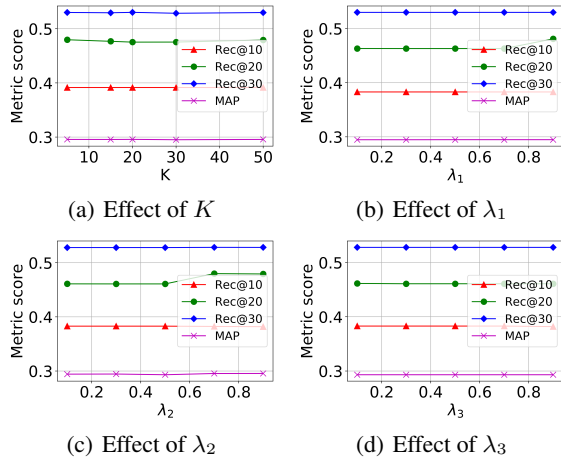


Figure 4: The impacts of inherent intention number (K) and different loss function weights (λ_1 , λ_2 and λ_3) on SH→HZ dataset w.r.t. the recommendation performance.

leisure-related while the latter is more function-related.

Parameter Sensitivity. We report the influence of the number of inherent travel intentions, i.e. K , and the impacts of λ_1 , λ_2 and λ_3 in loss functions, respectively. We only demonstrate the results on SH→HZ dataset. The results are similar on the other two datasets. Besides, to evaluate the impacts of parameters λ_1 , λ_2 and λ_3 on the recommendation performance, we first vary λ_1 while setting $\lambda_2 = \lambda_3 = \frac{1-\lambda_1}{2}$, and we use the similar strategy for evaluating λ_2 and λ_3 , respectively.

As shown in Fig. 4(a), the scores are relatively stable as K increases. The reason behind is that the user-specific intention is a summary of all generic intention embeddings, which makes the overall performance insensitive to the number of generic inherent intentions. Moreover, from Figs. 4(b) to 4(d), we can observe that, with increasing λ_1 , λ_2 and λ_3 from 0.1 to 0.9, both the Recall and MAP metrics are stable in general with an exception that the Rec@20 is slightly better with larger λ_1 and λ_2 .

Related Work

Out-of-town POI recommendation attempts to provide a list of POIs the out-of-town users are interested in (Wang et al. 2017). Comparing with general POI recommendations (Zhou, Mascolo, and Zhao 2019; Wang et al. 2018) that

involve factors like geographical influence, temporal influence, social influence and so forth, the out-of-town recommendation is more intractable and specialized due to the cold start, interest drift, and other domain gap issues. Some researchers have studied the aforementioned informative features to realize out-of-town recommendation, such as user preference (FERENCE, Ye, and Lee 2013), geographical influence (FERENCE, Ye, and Lee 2013; Pham, Li, and Cong 2017), and social influence (FERENCE, Ye, and Lee 2013). Interest drift refers to the phenomenon that user’s out-of-town check-ins are not aligned to user’s home-town preference. Some out-of-town recommenders have paid attentions to user’s interest drift (Yin et al. 2014, 2016; Wang et al. 2017). Most of them take the textual reviews as input via topic models. However, the data sparsity issue is getting worse when utilizing the textual content related to POIs and users. Besides, as a number of POIs in out-of-town check-ins relate to tourism, there also exist researches focusing on out-of-town tourism POI recommendation (Liu et al. 2011; Brilhante et al. 2013; Hu et al. 2017). Our work differentiates itself from previous works by comprehensively considering user’s preference, travel intention, geographical constraints and user interest drifts for out-of-town recommendation.

Topic models have been widely applied as generative models for different tasks (Wang et al. 2017; Xu et al. 2017; Shen et al. 2018; Zhou, Mascolo, and Zhao 2019; Luo et al. 2020). However, as the dimensionality grows, these methods are scant to perform fast and accurate inference. Recently, deep learning techniques and neural variational inference have accelerated the development of latent variable models (Miao, Yu, and Blunsom 2016; Miao, Grefenstette, and Blunsom 2017; Kingma and Welling 2013; Srivastava and Sutton 2017). For example, (Miao, Yu, and Blunsom 2016) developed a neural variational document model (NVDM) for text mining. (Miao, Grefenstette, and Blunsom 2017) proposed Neural Topic Model (NTM) to discover latent topics by variational inference. These methods offer us a new data driven paradigm towards topic discovery problem.

Concluding Remarks

In this paper, we studied the out-of-town recommendation problem via travel intention modeling. We proposed a data-driven framework TRAINOR to learn an out-of-town recommender by comprehensively considering user preference, interest drifts, travel intention and out-of-town geographical influence. To investigate user’s home-town preference, a G-GNN model was exploited. Besides, the user’s out-of-town preference was estimated in a collective manner and enriched through a geographical GCN. Afterward, we devised a preference transfer module to map home-town preference to out-of-town check-in behavior via an MLP. Moreover, to understand the user’s complex travel intention, we developed an NTM based travel intention discovery module. Finally, with jointly minimizing composite loss, the learned recommender was yielded. Through extensive experiments on real-world datasets, we validated the effectiveness of TRAINOR quantitatively. A case study further validated the ability of TRAINOR to understand users’ travel intentions.

Acknowledgments

This work is supported in part by NSFC 71531001, 61725205, 91746301, and 61703386.

References

- Brilhante, I.; Macedo, J. A.; Nardini, F. M.; Perego, R.; and Renso, C. 2013. Where shall we go today? Planning touristic tours with TripBuilder. In *CIKM'13*, 757–762.
- Cho, K.; Van Merriënboer, B.; Bahdanau, D.; and Bengio, Y. 2014. On the properties of neural machine translation: Encoder-decoder approaches. *arXiv preprint arXiv:1409.1259*.
- Ference, G.; Ye, M.; and Lee, W.-C. 2013. Location recommendation for out-of-town users in location-based social networks. In *CIKM'13*, 721–726.
- Geng, X.; Li, Y.; Wang, L.; Zhang, L.; Yang, Q.; Ye, J.; and Liu, Y. 2019. Spatiotemporal multi-graph convolution network for ride-hailing demand forecasting. In *AAAI'19*, volume 33, 3656–3663.
- Hidasi, B.; Karatzoglou, A.; Baltrunas, L.; and Tikk, D. 2015. Session-based recommendations with recurrent neural networks. *arXiv preprint arXiv:1511.06939*.
- Hu, G.; Shao, J.; Shen, F.; Huang, Z.; and Shen, H. T. 2017. Unifying multi-source social media data for personalized travel route planning. In *SIGIR'17*, 893–896.
- Kingma, D. P.; and Welling, M. 2013. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*.
- Kipf, T. N.; and Welling, M. 2016. Semi-supervised classification with graph convolutional networks. *arXiv preprint arXiv:1609.02907*.
- Li, S.; Zhou, J.; Xu, T.; Liu, H.; Lu, X.; and Xiong, H. 2020. Competitive Analysis for Points of Interest. In *KDD'20*, 1265–1274.
- Li, Y.; Tarlow, D.; Brockschmidt, M.; and Zemel, R. 2015. Gated graph sequence neural networks. *arXiv preprint arXiv:1511.05493*.
- Liu, Q.; Ge, Y.; Li, Z.; Chen, E.; and Xiong, H. 2011. Personalized travel package recommendation. In *ICDM'11*, 407–416. IEEE.
- Luo, H.; Zhou, J.; Bao, Z.; Li, S.; Culpepper, J. S.; Ying, H.; Liu, H.; and Xiong, H. 2020. Spatial object recommendation with hints: When spatial granularity matters. In *SIGIR'20*, 781–790.
- Man, T.; Shen, H.; Jin, X.; and Cheng, X. 2017. Cross-Domain Recommendation: An Embedding and Mapping Approach. In *IJCAI'17*, 2464–2470.
- Miao, Y.; Grefenstette, E.; and Blunsom, P. 2017. Discovering discrete latent topics with neural variational inference. *arXiv preprint arXiv:1706.00359*.
- Miao, Y.; Yu, L.; and Blunsom, P. 2016. Neural variational inference for text processing. In *ICML'16*, 1727–1736.
- Pham, T.-A. N.; Li, X.; and Cong, G. 2017. A general model for out-of-town region recommendation. In *WWW'17*, 401–410.
- Rendle, S.; Freudenthaler, C.; Gantner, Z.; and Schmidt-Thieme, L. 2012. BPR: Bayesian personalized ranking from implicit feedback. *arXiv preprint arXiv:1205.2618*.
- Shen, D.; Zhu, H.; Zhu, C.; Xu, T.; Ma, C.; and Xiong, H. 2018. A joint learning approach to intelligent job interview assessment. In *IJCAI'18*, 3542–3548.
- Srivastava, A.; and Sutton, C. 2017. Autoencoding variational inference for topic models. *arXiv preprint arXiv:1703.01488*.
- Wang, H.; Fu, Y.; Wang, Q.; Yin, H.; Du, C.; and Xiong, H. 2017. A location-sentiment-aware recommender system for both home-town and out-of-town users. In *SIGKDD'17*, 1135–1143.
- Wang, J.; Feng, Y.; Naghizade, E.; Rashidi, L.; Lim, K. H.; and Lee, K. 2018. Happiness is a choice: sentiment and activity-aware location recommendation. In *WWW'18*, 1401–1405.
- Wei, P.; and Mao, W. 2019. Modeling Transferable Topics for Cross-Target Stance Detection. In *SIGIR'19*, 1173–1176.
- Wu, S.; Tang, Y.; Zhu, Y.; Wang, L.; Xie, X.; and Tan, T. 2019. Session-based recommendation with graph neural networks. In *AAAI'19*, volume 33, 346–353.
- Xu, T.; Zhu, H.; Zhu, C.; Li, P.; and Xiong, H. 2017. Measuring the popularity of job skills in recruitment market: A multi-criteria approach. *arXiv preprint arXiv:1712.03087*.
- Yin, H.; Cui, B.; Sun, Y.; Hu, Z.; and Chen, L. 2014. LCARS: A spatial item recommender system. *TOIS* 32(3): 1–37.
- Yin, H.; Cui, B.; Zhou, X.; Wang, W.; Huang, Z.; and Sadiq, S. 2016. Joint modeling of user check-in behaviors for real-time point-of-interest recommendation. *TOIS* 35(2): 1–44.
- Zeng, J.; Li, J.; Song, Y.; Gao, C.; Lyu, M. R.; and King, I. 2018. Topic memory networks for short text classification. *arXiv preprint arXiv:1809.03664*.
- Zhang, W.; Liu, H.; Liu, Y.; Zhou, J.; and Xiong, H. 2020. Semi-Supervised Hierarchical Recurrent Graph Neural Network for City-Wide Parking Availability Prediction. In *AAAI'20*, volume 34, 1186–1193.
- Zhou, X.; Mascolo, C.; and Zhao, Z. 2019. Topic-enhanced memory networks for personalised point-of-interest recommendation. In *SIGKDD'17*, 3018–3028. ACM.