Origin-Aware Next Destination Recommendation with Personalized Preference Attention

Nicholas Lim

GrabTaxi Holdings, Singapore
nic.lim@grab.com

Bryan Hooi

Institute of Data Science and School
of Computing, National University of
Singapore, Singapore
dcsbkh@nus.edu.sg

See-Kiong Ng

Institute of Data Science and School
of Computing, National University of
Singapore, Singapore
seekiong@nus.edu.sg

Xueou Wang

Grab-NUS AI Lab, National
University of Singapore, Singapore
idswx@nus.edu.sg

Yong Liang Goh

GrabTaxi Holdings, Singapore
yongliang.goh@grab.com

Renrong Weng

GrabTaxi Holdings, Singapore
renrong.weng@grab.com

Rui Tan

GrabTaxi Holdings, Singapore
rui.tan@grab.com

ABSTRACT

Next destination recommendation is an important task in the transportation domain of taxi and ride-hailing services, where users are recommended with personalized destinations given their current origin location. However, recent recommendation works do not satisfy this origin-awareness property, and only consider learning from historical destination locations, without origin information. Thus, the resulting approaches are unable to learn and predict origin-aware recommendations based on the user's current location, leading to sub-optimal performance and poor real-world practicality. Hence, in this work, we study the origin-aware next destination recommendation task. We propose the Spatial-Temporal Origin-Destination Personalized Preference Attention (STOD-PPA) encoder-decoder model to learn origin-origin (OO), destination-destination (DD), and origin-destination (OD) relationships by first encoding both origin and destination sequences with spatial and temporal factors in local and global views, then decoding them through personalized preference attention to predict the next destination. Experimental results on seven real-world user trajectory taxi datasets show that our model significantly outperforms baseline and state-of-the-art methods.

CCS CONCEPTS

• Information systems → Recommender systems.

KEYWORDS

Recommender System; Recurrent Neural Network; Spatio-Temporal

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

WSDM '21, March 8–12, 2021, Virtual Event, Israel

© 2021 Association for Computing Machinery.

ACM ISBN 978-1-4503-8297-7/21/03...\$15.00

<https://doi.org/10.1145/3437963.3441797>

ACM Reference Format:

Nicholas Lim, Bryan Hooi, See-Kiong Ng, Xueou Wang, Yong Liang Goh, Renrong Weng, and Rui Tan. 2021. Origin-Aware Next Destination Recommendation with Personalized Preference Attention. In *Proceedings of the Fourteenth ACM International Conference on Web Search and Data Mining (WSDM '21)*, March 8–12, 2021, Virtual Event, Israel. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/3437963.3441797>

1 INTRODUCTION

Recent years have seen rapid growth in the popularity of ride-hailing and mobile applications for taxi-booking, such as Uber, Didi, and others, where users book taxi rides from an origin (O) to a destination (D) location. This surge in taxi booking transactions has resulted in massive user trajectory datasets that provide the opportunity to learn to perform personalized recommendations for users, such as predicting their next destination location.

Different from the popular existing taxi trajectory datasets [2, 16] that record the sequence of locations from multiple different customers or users for the same taxi, a user trajectory instead records the taxi riding patterns for the *same user*. Thus, they record the sequence of past taxi trips by a user, in the form of origin-destination (OD) tuples that reflect her transport mobility behaviors and latent preferences. Accordingly, these user trajectory datasets present the opportunity to learn each user's preferences, in order to provide personalized recommendations, and to improve performance on the next destination recommendation task.

Existing works for the next Point-of-Interest (POI) or destination location recommendation task mainly focus on the Location-Based Social Network (LBSN) domain, where they learn from LBSN's user trajectories datasets, containing only users' sequential historical destinations or checked-in locations with *no origin information* to predict the next destination. However, in the transportation domain, the historical taxi trips do not just record where the user has visited (i.e. destination), but also where the user has come from (i.e. origin). The inclusion of such origin information can provide several key benefits to existing next destination recommendation works that currently only rely on destination sequences. First, as

shown in Fig. 1, the sequential transition of destinations from d_{t_1} to d_{t_2} can be used to learn destination-destination (DD) relationships and predict the next destination of d_{t_i} , as studied extensively in existing recommendation works. However, we can also see that the associating origin locations of o_{t_1} , o_{t_2} and current origin o_{t_i} can serve as a valuable source of input to learn both origin-origin (OO) and OD relationships to help predict the next destination d_{t_i} correctly, but such relationships (dashed arrows on Fig. 1) have not been studied by existing works. Second, by conditioning the next destination prediction d_{t_i} on the current origin o_{t_i} or where the user is currently at when booking a taxi ride, this extends the problem to *origin-aware next destination recommendation*, where this conditioning of origin-awareness eliminates the problem whereby the same recommendation set is always suggested regardless of where the user is, which is impractical in a real-world setting (e.g. very far away destinations should not be recommended).

Among the recent existing works which do not consider origin information for their next destination recommendations, RNN-based approaches have been shown to surpass classical methods of Markov Chains (MCs) and Matrix Factorization (MF). For instance, [15] proposed the Spatial Temporal Recurrent Neural Network (ST-RNN) to exploit spatial and temporal intervals between neighbouring destination locations, where a time window is used to consider several destination locations as input. [12] proposed the Hierarchical Spatial-Temporal Long-Short Term Memory (HST-LSTM) to leverage spatial and temporal intervals directly into the existing multiplicative gates of LSTM. [25] proposed the Spatio-Temporal Gated Coupled Network (STGCN) to capture short and long-term user preferences with new distance and time gates from the continuous spatial and temporal intervals. [18] proposed the Long- and Short-Term Preference Modeling (LSTPM) to learn long and short term user preferences through a nonlocal network and a geo-dilated RNN respectively and is the state-of-the-art approach.

In this paper, we study the origin-aware next destination recommendation task by proposing a Spatial-Temporal Origin-Destination Personalized Preference Attention (STOD-PPA) model based on an encoder-decoder framework to learn the OO, DD, and OD relationships as shown in Fig. 1. The STOD-PPA model first encodes both historical origin and destination sequences using a novel Spatial-Temporal LSTM (ST-LSTM) module, an extension of the LSTM to enable OD relationships to be learned by new cell states based on spatial and temporal factors from both local and global views. Unlike recent works that focus on short and long term user preferences [18, 25], we propose a novel Personalized Preference Attention (PPA) decoder module to compute a hidden representation for the prediction task by performing personalized preference attention directly on all encoded hidden states of both OD sequences to best learn the dynamic preferences of the user. To summarise, the contributions of this paper are as follows:

- We propose a novel encoder-decoder STOD-PPA model to learn OO, DD, and OD relationships for the origin-aware next destination recommendation task. To the best of our knowledge, we are the first work to study the origin-aware next destination recommendation task.
- Our proposed ST-LSTM encoder module is the first to consider both local and global views for spatial and temporal factors and

Figure 1: Illustration of a user trajectory containing both origin (blue nodes) and destination (red nodes) sequences, available in the transportation domain to predict the next destination d_{t_i} . Maps © OpenStreetMap contributors, CC BY-SA.

incorporate them into their own spatial and temporal cell states to learn OD relationships.

- Experiments conducted on seven real-world user trajectory taxi datasets of different countries in the transportation domain show that our approach outperforms baseline and state-of-the-art methods significantly.

2 RELATED WORK

The closest line of work to our origin-aware next destination recommendation task is the next Point-of-Interest (POI) or location recommendation task, which does not use origin information, but only sequentially visited locations (i.e. destinations). This problem is mostly studied in the LBSN domain, such as the popular Gowalla, Brightkite, and Foursquare datasets [4, 23]. Early works studied conventional sequential approaches of MC, as well as collaborative filtering via MF. FPMC-LR [3] extended FPMC [17] to incorporate personalized MC and applied a localized region constraint where only nearby new POIs around the users' historical POIs are considered for the recommendation task, reducing noisy information and improves user experience [26]. PRME [9] learns both sequential transitions of POIs and user preferences by modelling them in a latent space. PRME-G is a variant that considers the geographical influence to improve the recommendations. [10] proposed a Bayesian Personalized Ranking (BPR) approach to learn check-in behaviors and fuses the personalized MC with latent patterns. CAPE [1] learns content-aware POI embeddings from user check-in sequences and POI textual information. [21] proposed GE, an approach to learn POI, time slot, region, and word embeddings from various graphs, such as POI-POI graphs [14, 19, 20, 24], then using these embeddings to score each POI for the recommendation task.

Recently, RNN-based approaches have been demonstrated to perform better than these existing methods due to their ability to model sequential dependencies in user location sequences. ST-RNN [15] first proposed the use of spatial and temporal intervals among sequences of POIs by incorporating them into an RNN after performing linear interpolation. [12] proposed HST-LSTM, a hierarchical LSTM model for session data, and similarly incorporates

spatial and temporal intervals into LSTM's existing gates after using linear interpolation. STGN [25] is a LSTM-based model that includes new time and distance gates to incorporate temporal and spatial intervals respectively. The model also includes an additional cell state so that both short and long-term user preferences can be learned. STGCN, a variant of STGN, was also proposed where the forget gate is removed for efficiency. CatDM [22] considers both POI categories and spatial proximity to mitigate data sparsity, but is limited to only predicting the POIs visited in the next window of 24 hours. DeepMove [8] uses a historical attention module and a GRU to learn sequential transitions from visit sequences. LSTPM [18] learns both long and short term user preferences through a nonlocal network architecture where the spatial and temporal correlations between current and past trajectories are considered. The long term preferences exploit the historical trajectories using a nonlocal network, while the short term preferences focus on the most recent trajectory to model users' latest preferences using a geo-dilated RNN. LSTPM is also the state-of-the-art model for the next POI recommendation task. However, these existing works are not designed to work with origin information.

Among the other existing works, approaches that consider the use of both origin and destination information are limited, where the only work of [7] applied a simple frequency approach under their suggestion framework to compute a recommendation list by considering both user's historical destinations and popular city-wide locations; however, this approach is *not origin-aware*. In this paper, we only consider their suggestion framework as their prediction framework does not compute a ranked list for evaluation.

3 PRELIMINARIES

Problem Formulation. Let $U = \{u_1, u_2, \dots, u_M\}$ be the set of M users and $L = \{l_1, l_2, \dots, l_N\}$ be the set of N locations for the users in U to visit, where each location $l_n \in L$ can either have the role of an origin o or a destination d , or both among the dataset of user trajectories. Each user u_m has an OD sequence of pick-up (origin) and drop-off (destination) tuples $s_{u_m} = \{(o_{t_1}, d_{t_1}), (o_{t_2}, d_{t_2}), \dots, (o_{t_i}, d_{t_i})\}$ that record their taxi trips, and $S = \{s_{u_1}, s_{u_2}, \dots, s_{u_M}\}$ is the set of all users' OD sequences. All location visits in S , regardless of origin or destination, have their own location coordinates LC and timestamp $time$. The objective of the origin-aware next destination recommendation task is to consider the user u_m , the current origin o_{t_i} , and her historical sequence of OD tuples $\{(o_{t_1}, d_{t_1}), (o_{t_2}, d_{t_2}), \dots, (o_{t_{i-1}}, d_{t_{i-1}})\}$ to recommend an ordered set of destinations from L , where the next destination d_{t_i} should be highly ranked in the recommendation set. We further denote $s_{u_m}^{train}$ and S^{train} as sets from the training partition with the superscript *train* for clarity.

3.1 LSTM

The LSTM [11] is a variant of RNN that is capable of learning sequential transitions across timesteps of the sequence through gating mechanisms and a memory cell state to retain information. The gating mechanism helps to mitigate the vanishing gradient problem for sequences and has been found effective in various sequential predictive applications. Here, we include the original

LSTM model and its equations:

$$i_{t_i} = \sigma(\mathbf{W}_i x_{t_i} + \mathbf{U}_i h_{t_{i-1}} + \mathbf{b}_i) \quad (1)$$

$$f_{t_i} = \sigma(\mathbf{W}_f x_{t_i} + \mathbf{U}_f h_{t_{i-1}} + \mathbf{b}_f) \quad (2)$$

$$o_{t_i} = \sigma(\mathbf{W}_o x_{t_i} + \mathbf{U}_o h_{t_{i-1}} + \mathbf{b}_o) \quad (3)$$

$$\tilde{c}_{t_i} = \tanh(\mathbf{W}_c x_{t_i} + \mathbf{U}_c h_{t_{i-1}} + \mathbf{b}_c) \quad (4)$$

$$c_{t_i} = f_{t_i} \odot c_{t_{i-1}} + i_{t_i} \odot \tilde{c}_{t_i} \quad (5)$$

$$h_{t_i} = o_{t_i} \odot \tanh(c_{t_i}) \quad (6)$$

where i_{t_i} , f_{t_i} , o_{t_i} are the input, forget and output gates respectively that regulate information flow in the scale of 0 to 1 from the sigmoid activation function. For each timestep t_i , the LSTM extracts information via the input gate i_{t_i} from the cell input $\tilde{c}_{t_i}$ using Hadamard product $\odot$, representing "how much to store in the cell state c_{t_i} ". The forget gate f_{t_i} regulates "how much to forget from the previous cell state $c_{t_{i-1}}$ ". The LSTM outputs c_{t_i} and h_{t_i} for the next timestep accordingly, where h_{t_i} is the hidden state or output hidden representation regulated by the output gate indicating "how much to output for this timestep" from the current cell state c_{t_i} .

4 APPROACH

In this section, we first propose the ST-LSTM module (encoder), an extension of the LSTM to learn OD relationships, followed by the PPA module (decoder) to best learn users' dynamic preferences.

4.1 ST-LSTM

Recent RNN and LSTM based works [12, 15, 25] have incorporated spatial (distance) and temporal (time) intervals of Δd_{t_i} and Δt_{t_i} respectively, between neighbouring pairs of destinations as input to their models. For example, given the previous destination $d_{t_{i-1}}$ and the next destination d_{t_i} to be visited and predicted by the model, the spatial interval $\Delta d_{t_i} = d(LC_{t_i}, LC_{t_{i-1}})$ is computed using a distance function d on the location coordinates LC of both destinations d_{t_i} and $d_{t_{i-1}}$. Similarly, the temporal interval $\Delta t_{t_i} = time_{t_i} - time_{t_{i-1}}$ is computed from the corresponding timestamps $time$ of both d_{t_i} and $d_{t_{i-1}}$ destinations. Intuitively, the spatial and temporal intervals seek to describe "how far is the next destination" and "how long before they visit the next destination" respectively, using these intervals as inputs to their models accordingly to predict the next destination d_{t_i} . Notably, this usage of intervals makes the key assumption that the next destination d_{t_i} and its details (i.e. location coordinates LC_{t_i} and timestamp $time_{t_i}$) are known in advance, making it impractical in a real-world setting where the next destination visit d_{t_i} is not yet known [8]. Apart from this impracticality flaw, the consideration of spatial and temporal factors between only neighbouring destination pairs restricts the model to only learn from a local view and not from a global view where all locations from L are involved.

Inspired from these two flaws, we propose to extend the LSTM with spatial and temporal cell states to learn OD relationships based on the spatial and temporal factors in both local and global views. This extension seeks to allow OD relationships to be learned as the LSTM is already capable of learning OO or DD relationships given an origin or a destination sequence respectively but is unable to learn OD relationships.

Learning Spatial OD Relationships. We propose the use of Geohash embeddings (local view) and spatial intervals (global view)

Figure 2: Learning spatial factors of origins (blue nodes) and destinations (red nodes). Maps © OpenStreetMap contributors, CC BY-SA.

to learn spatial OD relationships among all locations in L , which includes both origins and destinations. First, we partition the map with a Geohash grid. As each location $l_n \in L$ has its own location coordinates LC , we map l_n to its corresponding Geohash cell l_n^{geo} , as well as the Geohash embedding of that cell $\tilde{l}_n^{geo}$. Accordingly, this means that locations in the same region or Geohash would be trained to be similar to each other under a *local view*, such as the origin l_1 and destinations of l_2 and l_3 in Fig. 2(a), allowing region specific semantics such as shopping districts and housing neighbourhood areas to be learned. However, a limitation with this approach is that the relationship between l_1 and l_4 may be deemed dissimilar by the model due to the assignment of different Geohash embeddings based on their respective Geohash cells, even when they are near to each other, due to their different cells or areas.

To mitigate this limitation, we propose to also learn the spatial factor from a *global view*. In Fig. 2(b), we can see that instead of partitioning the same map and locations by a Geohash grid, for the global view, we compute pair-wise spatial or distance intervals Δs between the origin l_1 and all locations available of $\{l_1, l_2, l_3, l_4, l_5\}$ (including l_1 itself for simplicity). The intervals would be able to embed the spatial proximity of the current location l_1 to all the other locations and itself, where the origin l_1 can now learn that it is also similar or nearby to the destination l_4 , which was not achievable under a local view as shown on Fig. 2(a). However, the global view also has a limitation of being unable to learn area or Geohash specific semantics, necessitating the consideration of both local and global views to learn OD relationships.

Now, instead of just considering $\{l_1, l_2, l_3, l_4, l_5\}$ for the example in Fig. 2(b), we perform the pair-wise spatial interval computation between l_1 and all locations $l_n \in L$ using the Haversine distance function, to compute a vector representation of the spatial intervals $\Delta \vec{s}_{l_1} \in \mathbb{R}^{|L|}$ for l_1 to represent the *global view*. Intuitively, this provides a global context on how far or how near are all other locations (origin or destination) in the whole city or country is to l_1 , whereas the local view focuses on region level semantics. Different from the existing works that require the last spatial and temporal intervals of Δd_{t_i} and Δt_{t_i} respectively as input to the model [12, 15, 25] to predict the next destination d_{t_i} , here, we compute intervals between a location l_1 to all locations $l_n \in L$ without the need to know the next destination location d_{t_i} in advance.

Figure 3: ST-LSTM module where green and brown connections are the proposed spatial and temporal cell states.

With an input location embedding $\tilde{l}_{t_i} \in \mathbb{R}^{dim}$ as x_{t_i} for Eq. (1) to (4) of the LSTM, as well as its mapped Geohash embedding $\tilde{l}_{t_i}^{geo} \in \mathbb{R}^{dim}$ (local view) and its spatial interval vector $\Delta \vec{s}_{l_{t_i}} \in \mathbb{R}^{|L|}$ (global view), where dim is the embedding dimension, we propose a new spatial cell state and its associating gates and cell input *in addition* to LSTM's cell state:

$$i_{t_i}^s = \sigma(\mathbf{W}_i^s \tilde{l}_{t_i}^{geo} + \mathbf{V}_i^s \Delta \vec{s}_{l_{t_i}} + \mathbf{U}_i^s h_{t_{i-1}} + \mathbf{b}_i^s) \quad (7)$$

$$f_{t_i}^s = \sigma(\mathbf{W}_f^s \tilde{l}_{t_i}^{geo} + \mathbf{V}_f^s \Delta \vec{s}_{l_{t_i}} + \mathbf{U}_f^s h_{t_{i-1}} + \mathbf{b}_f^s) \quad (8)$$

$$\tilde{c}_{t_i}^s = \tanh(\mathbf{W}_c^s \tilde{l}_{t_i}^{geo} + \mathbf{V}_c^s \Delta \vec{s}_{l_{t_i}} + \mathbf{U}_c^s h_{t_{i-1}} + \mathbf{b}_c^s) \quad (9)$$

$$c_{t_i}^s = f_{t_i}^s \odot c_{t_{i-1}}^s + i_{t_i}^s \odot \tilde{c}_{t_i}^s \quad (10)$$

where $i_{t_i}^s, f_{t_i}^s, \tilde{c}_{t_i}^s, c_{t_i}^s$ are the set of input gate, forget gate, cell input and cell state respectively to learn the sequential spatial transitions with the superscript s . Here, $\mathbf{W}_i^s, \mathbf{W}_f^s, \mathbf{W}_c^s \in \mathbb{R}^{dim \times Hdim}$

learns a projection of the Geohash embedding $\tilde{l}_{t_i}^{geo}$ (local view) for the gates $i_{t_i}^s, f_{t_i}^s$ and cell input $\tilde{c}_{t_i}^s$ respectively, where $Hdim$ is the hidden representation dimension. Similarly, for the global view, $\mathbf{V}_i^s, \mathbf{V}_f^s, \mathbf{V}_c^s \in \mathbb{R}^{|L| \times Hdim}$ learns a representation from the spatial interval vector $\Delta \vec{s}_{l_{t_i}}$. Same as Eq. (1) to (4), we also add $\mathbf{U}_i^s, \mathbf{U}_f^s, \mathbf{U}_c^s \in \mathbb{R}^{Hdim \times Hdim}$ and their biases $\mathbf{b}_i^s, \mathbf{b}_f^s, \mathbf{b}_c^s \in \mathbb{R}^{Hdim}$ for the previous hidden state $h_{t_{i-1}}$, as this enforces the recurrent structure where the output hidden state of the previous timestep is used as input to the current timestep to model the sequential spatial transitions. The spatial cell state $c_{t_i}^s$ is then computed from Eq. (10) from its own set of gates and cell input, encoding the sequential spatial factor for use by the next timestep accordingly.

Learning Temporal OD Relationships. Similar to the spatial cell state, we also propose a *separate* temporal cell state $c_{t_i}^t$ to learn temporal OD relationships from both local and global views:

$$i_{t_i}^t = \sigma(\mathbf{W}_i^t \tilde{l}_{t_i}^{slot} + \mathbf{V}_i^t \Delta \vec{t}_{l_{t_i}} + \mathbf{U}_i^t h_{t_{i-1}} + \mathbf{b}_i^t) \quad (11)$$

$$f_{t_i}^t = \sigma(\mathbf{W}_f^t \tilde{l}_{t_i}^{slot} + \mathbf{V}_f^t \Delta \vec{t}_{l_{t_i}} + \mathbf{U}_f^t h_{t_{i-1}} + \mathbf{b}_f^t) \quad (12)$$

$$\tilde{c}_{t_i}^t = \tanh(\mathbf{W}_c^t \tilde{l}_{t_i}^{slot} + \mathbf{V}_c^t \Delta \vec{t}_{l_{t_i}} + \mathbf{U}_c^t h_{t_{i-1}} + \mathbf{b}_c^t) \quad (13)$$

$$c_{t_i}^t = f_{t_i}^t \odot c_{t_{i-1}}^t + i_{t_i}^t \odot \tilde{c}_{t_i}^t \quad (14)$$

where the main differences with the spatial cell state equations from Eq. (7) to (10), apart from its own temporal weight matrices, are the inclusion of $\tilde{l}_{t_i}^{slot} \in \mathbb{R}^{dim}$ as the corresponding timeslot embedding

of the input location l_{t_i} (local view) and $\Delta \vec{l}_{l_{t_i}} \in \mathbb{R}^{|L|}$ as the temporal interval vector between the input location l_{t_i} and all locations in L (global view). As each location visit (origin or destination) includes timestamp data *time* (taxi pick-up timestamp for origin and drop-off timestamp for destination), we can map each location visit in s^{train} to its corresponding timeslot $l_{t_i}^{slot}$. Similar to learning OD relationships in a Geohash cell for the spatial factor, this would allow origin and destination locations in the same timeslot to learn to be similar under a local view for the temporal factor. We use a total of eight timeslots of three hours each to represent the periodic changes across a day. For the global view of temporal interval vector $\Delta \vec{l}_{l_{t_i}}$, we first identify OD tuples in all users' sequences s^{train} that contain l_{t_i} (origin or destination), and compute the pairwise temporal interval as $\Delta t = time^d - time^o$, where $time^o$ and $time^d$ are the origin and destination timestamps respectively of the OD tuple.

Next, we modify Eq. (6) to combine the spatial, temporal and LSTM's memory cell states as the output hidden state for the current timestep:

$$h_{t_i} = o_{t_i} \odot \tanh(\mathbf{W}_h(c_{t_i} \parallel c_{t_i}^s \parallel c_{t_i}^t)) \quad (15)$$

where $\mathbf{W}_h \in \mathbb{R}^{3 \times Hdim \times Hdim}$ fuses the three cell states after the concatenate operation $\parallel$, followed by the hyperbolic tangent activation function, then applies the LSTM's existing output gate o_{t_i} to learn "how much to output" from the fused sequential location, spatial and temporal cell states. Accordingly, we illustrate the ST-LSTM module on Fig. 3.

4.2 STOD-PPA

We further propose STOD-PPA based on an encoder-decoder framework to first encode OD sequences using our ST-LSTMs, then use a PPA decoder to perform personalized preference attention to all the encoded OD hidden states.

Encoder. As each user's sequence of OD tuples s_{u_m} is partitioned into training and testing partitions, here, we use the training partition $s_{u_m}^{train} = \{(o_{t_1}, d_{t_1}), (o_{t_2}, d_{t_2}), \dots, (o_{t_i}, d_{t_i})\}$ and split them into separate origin and destination sequences of $s_{u_m}^{train^O} = \{o_{t_2}, o_{t_3}, \dots, o_{t_i}\}$ and $s_{u_m}^{train^D} = \{d_{t_1}, d_{t_2}, \dots, d_{t_{i-1}}\}$ respectively. For efficiency, we omitted the first origin o_{t_1} from $s_{u_m}^{train^O}$ and the last destination d_{t_i} from $s_{u_m}^{train^D}$ so that both the encoder and decoder will use the same set of input sequences, allowing batch training to be performed for each user. Then, we propose to encode both $s_{u_m}^{train^O}$ and $s_{u_m}^{train^D}$ separately with a ST-LSTM each, allowing OO and DD relationships to be learned in their own ST-LSTM respectively, as well as OD relationships from the newly proposed spatial and temporal cell states. First, we use a multi-modal embedding layer *Emb* to embed the location l_{t_i} (origin or destination), its corresponding Geohash $l_{t_i}^{geo}$ and timeslot $l_{t_i}^{slot}$:

$$(\vec{l}_{t_i}, \vec{l}_{t_i}^{geo}, \vec{l}_{t_i}^{slot}) = \underset{\mathbf{W} \in \{\mathbf{W}_L, \mathbf{W}_G, \mathbf{W}_T\}}{\text{Emb}} \mathbf{W} (l_{t_i}, l_{t_i}^{geo}, l_{t_i}^{slot}) \quad (16)$$

where $\mathbf{W}_L \in \mathbb{R}^{|L| \times dim}$, $\mathbf{W}_G \in \mathbb{R}^{|G| \times dim}$, $\mathbf{W}_T \in \mathbb{R}^{|T| \times dim}$ are the location, Geohash and timeslot weight matrices respectively, $|G|$ is the total number of corresponding Geohashes after mapping all

Figure 4: Illustration of STOD-PPA.

locations in L , and $|T|$ is the number of timeslots, where we use $|T| = 8$ of three hours each. Accordingly, at each timestep, the input to an ST-LSTM would have an input tuple of the mapped embeddings and the global interval vectors $(\vec{l}_{t_i}, \vec{l}_{t_i}^{geo}, \Delta \vec{l}_{l_{t_i}}, \vec{l}_{t_i}^{slot}, \Delta \vec{l}_{l_{t_i}})$, as shown on Fig. 3. Next, we proceed to encode the origin and destination sequences separately with their own ST-LSTMs, denoted as $\phi^O(\cdot)$ and $\phi^D(\cdot)$ respectively, specifically:

$$h_{u_m}^O = \phi^O(s_{u_m}^{train^O}) \quad (17)$$

$$h_{u_m}^D = \phi^D(s_{u_m}^{train^D}) \quad (18)$$

where $h_{u_m}^O$ and $h_{u_m}^D$ are sets containing all origin and destination encoded hidden states across the timesteps respectively and $|h_{u_m}^O| = |h_{u_m}^D|$. Then, we concatenate both $h_{u_m}^O$ and $h_{u_m}^D$ for a final set of all hidden states $h_{u_m}^{OD}$ for user $u_m \in U$ as the output of the encoder, for use by the decoder in training and testing.

Decoder. After encoding the OD sequences, we propose the Personalized Preference Attention (PPA) decoder module to attend to all the encoded OD hidden states and compute an origin-aware and personalized hidden representation based on the users' dynamic preferences. Different from the existing works that includes user embeddings by addition [15] or concatenation [8] to the hidden representation used for prediction, here, we propose the novel approach of using user embeddings to learn to perform attention on the encoded OD hidden states. First, we update the multi-modal embedding layer *Emb* in Eq. (16) with an additional user weight matrix $\mathbf{W}_U \in \mathbb{R}^{|U| \times dim}$ and use this layer to map user u_m , current origin o_{t_i} and previous destination $d_{t_{i-1}}$ to their corresponding embeddings of $\vec{u}_m, \vec{o}_{t_i}$ and $\vec{d}_{t_{i-1}}$ respectively. Then, from each user's encoded hidden states $\vec{h}_i \in h_{u_m}^{OD}$, we can compute its respective personalized attention score:

$$\vec{\alpha}_i = \frac{\exp\left(\sigma_{LR}\left(\mathbf{W}_A(\vec{u}_m \parallel \vec{o}_{t_i} \parallel \vec{d}_{t_{i-1}} \parallel \vec{h}_i)\right)\right)}{\sum_{\vec{p}_i \in h_{u_m}^{OD}} \exp\left(\sigma_{LR}\left(\mathbf{W}_A(\vec{u}_m \parallel \vec{o}_{t_i} \parallel \vec{d}_{t_{i-1}} \parallel \vec{p}_i)\right)\right)} \quad (19)$$

where $\vec{\alpha}_i \in \mathbb{R}^{Hdim}$, $\mathbf{W}_A \in \mathbb{R}^{(3 \cdot dim + Hdim) \times Hdim}$, and σ_{LR} is the Leaky ReLU activation function. At every timestep, the decoder takes the input of $\vec{u}_m, \vec{o}_{t_i}, \vec{d}_{t_{i-1}}$, and each hidden state $\vec{h}_i \in h_{u_m}^{OD}$ to predict a personalized attention score $\vec{\alpha}_i \in \alpha_{t_i}$ after softmax

normalization, where $|\alpha_{t_i}| = |h_{u_m}^{OD}|$ accordingly. Effectively, the computed personalized attention scores α_{t_i} indicates “how much to attend to each encoded transition or hidden state” among all the encoded OD hidden states $h_{u_m}^{OD}$, that can best predict the next destination d_{t_i} , by optimizing the trainable weight matrix $\mathbf{W}_A$.

Unlike recent works [18, 25] that focus on learning long and short term preferences, our approach considers all hidden states directly for the personalized preference attention process to best learn dynamic user preferences. With the personalized attention scores, we then perform a weighted sum on all of the corresponding hidden states:

$$\tilde{y}_{t_i} = \sum_{\tilde{\alpha}_i \in \alpha_{t_i}, \tilde{h}_i \in h_{u_m}^{OD}} \tilde{\alpha}_i \odot \tilde{h}_i \quad (20)$$

$$P(d_{t_i} | o_{t_i}, d_{t_{i-1}}, u_m) = \text{softmax}(\mathbf{W}_{loc}(\tilde{y}_{t_i})) \quad (21)$$

where $\tilde{y}_{t_i} \in \mathbb{R}^{Hdim}$ is the output hidden representation for timestep t_i . Lastly, we project $\tilde{y}_{t_i}$ to the number of locations or $|L|$ where $\mathbf{W}_{loc} \in \mathbb{R}^{Hdim \times |L|}$, followed by a softmax function to derive a probability distribution of all locations by learning $P(d_{t_i} | o_{t_i}, d_{t_{i-1}}, u_m)$ as a multi-classification problem. Accordingly, we can sort the distribution in descending order to achieve the final ranked recommendation set where the next destination location d_{t_i} should be highly ranked. Fig. 4 illustrates the STOD-PPA model.

Prediction. After training STOD-PPA, at only prediction or test time, we *deactivate the encoder* and only use the decoder for prediction as the encoded hidden states $h_{u_m}^{OD}$ from each users’ historical OD sequences are stored for efficiency after the training phase. The stored or pre-computed $h_{u_m}^{OD}$ are then actively retrieved by the decoder at test time for the respective user to perform attention on them and compute a ranked set prediction based on the different test case or input tuple of the user u_m , current origin o_{t_i} and previous destination $d_{t_{i-1}}$, as shown on Fig. 4, in order to compute Eq. (19) to (21).

5 EXPERIMENTS

5.1 Datasets

We use seven real-world user trajectory taxi datasets of different Southeast Asia (SE) countries in the recent year of 2019 from the ride-hailing company Grab for evaluation, where users would book taxis from the mobile application. Table 1 shows the details of the datasets (country names omitted to not reflect business insights) where the total number of locations is computed from the total number of origins and destinations, and each trip is an OD tuple accordingly from pick-up (origin) to drop-off (destination). For preprocessing, we use the same settings as [25], where users with less than 10 OD tuples or trips are removed and locations (i.e. origin or destination) visited by less than 10 users are removed. Lastly, we sort each user’s visit records by timestamps in ascending order, taking the first 70% as the training set and the remaining 30% as the testing set in order to predict future visited destinations.

5.2 Baseline Methods and Evaluation Metrics

- **TOP:** This rank locations based on the global frequencies in s_{train} and **U-TOP** rank locations in $s_{u_m}^{train}$ based on each users’ individually most frequent locations.

Table 1: Statistics of the seven datasets (after preprocessing).

Datasets	#users	#locations	#origins	#destinations	#trips
SE-1	2,662	1,694	1,131	563	23,730
SE-2	2,595	1,523	1,008	515	22,867
SE-3	2,677	1,469	972	497	21,164
SE-4	3,083	1,625	1,003	622	24,344
SE-5	2,452	1,397	891	506	22,256
SE-6	1,363	1,001	642	359	21,156
SE-7	3,301	2,044	1,315	729	23,019

- **TAXI** [7]: This frequency method considers both city-wide location frequencies and users’ historical destination frequencies.
 - **MF** [13]: MF is a traditional approach to many recommendation problems and can be used to learn the user-destination matrix.
 - **RNN** [6]: RNN takes advantage of sequential dependencies in location visit sequences with a basic recurrent structure. Its variants of **LSTM** [11] and **GRU** [5] includes the use of different multiplicative gates to regulate information flow. For fair comparison, we extend the RNN, GRU and LSTM baselines to $f(\vec{o}_{t_i} || \vec{d}_{t_{i-1}})$ where $f(\cdot)$ is the model of choice with the concatenated inputs of the current origin $\vec{o}_{t_i}$ and previous destination $\vec{d}_{t_{i-1}}$ to predict the next destination d_{t_i} . This ensures the use of both origin and destination information instead of just the latter, making it the same as the OD input to our STOD-PPA’s decoder in Eq. (19).
 - **HST-LSTM** [12]: This LSTM based model includes spatial and temporal intervals into LSTM existing gates. Following [25], we apply their proposed ST-LSTM variant here as the data does not include session information.
 - **STGN** [25]: A LSTM variant that learns both long and short term location visit preferences using new distance and time gates to model spatial and temporal factors, as well as a new cell state. Their variant **STGCN** removes the forget gate for better performance and efficiency.
- For HST-LSTM, STGN and STGCN, these methods use the spatial and temporal intervals between $d_{t_{i-1}}$ and d_{t_i} to predict d_{t_i} , which, as explained in Section 4.1, cannot be used in practice as this requires knowing the details of d_{t_i} in advance. Here, we use $d_{t_{i-2}}$ and $d_{t_{i-1}}$ instead to compute the intervals, so as to leverage the most recent available historical destination visits to predict d_{t_i} .
- **LSTPM** [18]: A LSTM-based variant that models long term preferences through the use of a nonlocal network, and short term preferences via a geo-dilated network. LSTPM is the state-of-the-art approach for the next destination recommendation task.
 - **LSTPM-OD**: Same as the RNN, GRU and LSTM baselines, for fair comparison, we extend LSTPM to LSTPM-OD to use $f(\vec{o}_{t_i} || \vec{d}_{t_{i-1}})$ for each timestep instead of just the previous destination $f(\vec{d}_{t_{i-1}})$ as described in their work [18].
 - **STOD-PPA** Our proposed approach, as shown in Fig. 4, using an encoder-decoder framework with our ST-LSTM as the encoder and the PPA as the decoder.

Same as [25] and other existing works, we use the standard metrics of Acc@K where $K \in \{1, 5, 10\}$ and Mean Average Precision (MAP) for evaluation. Acc@K measures the performance of the recommendation set up to K, where the smaller K is, the more challenging it is to perform well, such as Acc@1 where a score of 1 is awarded only if the ground truth next destination is in the first

Table 2: Performance in Acc@ K and MAP on seven user trajectory datasets from the transportation domain.

		TOP	U-TOP	TAXI	MF	RNN	GRU	LSTM	HST-LSTM	STGN	STGCN	LSTPM	LSTPM-OD	STOD-PPA
SE-1	Acc@1	0.0000	0.1344	0.0000	0.0155	0.1694	0.1621	0.2039	0.1103	0.0294	0.0530	0.3402±0.001	0.3683±0.002	0.4173±0.002
	Acc@5	0.0127	0.5275	0.0139	0.0232	0.2340	0.2348	0.2978	0.2181	0.1012	0.1407	0.5001±0.002	0.5357±0.002	0.5615±0.001
	Acc@10	0.0246	0.5791	0.0280	0.0380	0.2642	0.2694	0.3388	0.2737	0.1501	0.2019	0.5380±0.003	0.5689±0.002	0.5846±0.002
	MAP	0.0147	0.2956	0.0170	0.0294	0.2039	0.2006	0.2517	0.1666	0.0724	0.1040	0.4162±0.001	0.4482±0.001	0.4865±0.002
SE-2	Acc@1	0.0000	0.1250	0.0000	0.0164	0.1651	0.1636	0.2055	0.1051	0.0380	0.0525	0.3260±0.002	0.3617±0.002	0.4108±0.002
	Acc@5	0.0101	0.5233	0.0115	0.0267	0.2555	0.2648	0.3300	0.2324	0.1266	0.1588	0.5064±0.001	0.5434±0.002	0.5761±0.002
	Acc@10	0.0469	0.5814	0.0471	0.0453	0.2917	0.3104	0.3833	0.2937	0.1818	0.2270	0.5582±0.002	0.5889±0.002	0.6077±0.002
	MAP	0.0199	0.2906	0.0209	0.0316	0.2119	0.2150	0.2678	0.1702	0.0900	0.1119	0.4119±0.001	0.4477±0.001	0.4895±0.001
SE-3	Acc@1	0.0000	0.1285	0.0000	0.0123	0.1129	0.1003	0.1344	0.0652	0.0285	0.0384	0.2848±0.001	0.3163±0.002	0.3544±0.003
	Acc@5	0.0156	0.4542	0.0209	0.0211	0.1771	0.1757	0.2291	0.1541	0.0741	0.0908	0.4344±0.003	0.4762±0.004	0.4985±0.001
	Acc@10	0.0844	0.5083	0.0899	0.0381	0.2146	0.2109	0.2746	0.2048	0.1101	0.1362	0.4794±0.002	0.5176±0.004	0.5267±0.002
	MAP	0.0274	0.2608	0.0285	0.0268	0.1489	0.1416	0.1843	0.1156	0.0621	0.0758	0.3594±0.001	0.3948±0.002	0.4249±0.002
SE-4	Acc@1	0.0000	0.0856	0.0000	0.0140	0.1423	0.1341	0.1652	0.0830	0.0359	0.0460	0.2798±0.001	0.3047±0.001	0.3329±0.002
	Acc@5	0.0274	0.4405	0.0274	0.0209	0.2027	0.2035	0.2558	0.1830	0.1022	0.1329	0.4222±0.002	0.4565±0.002	0.4757±0.002
	Acc@10	0.0439	0.4912	0.0439	0.0341	0.2317	0.2363	0.2923	0.2326	0.1414	0.1870	0.4634±0.002	0.4955±0.002	0.5064±0.002
	MAP	0.0192	0.2323	0.0200	0.0264	0.1743	0.1709	0.2116	0.1355	0.0743	0.0948	0.3492±0.001	0.3780±0.001	0.4018±0.002
SE-5	Acc@1	0.0000	0.1063	0.0000	0.0123	0.1065	0.0940	0.1249	0.0546	0.0242	0.0286	0.2507±0.001	0.2739±0.001	0.3049±0.002
	Acc@5	0.0679	0.3985	0.0699	0.0219	0.1603	0.1537	0.2087	0.1361	0.0858	0.0972	0.3847±0.002	0.4184±0.003	0.4433±0.003
	Acc@10	0.1050	0.4583	0.1055	0.0365	0.1910	0.1890	0.2498	0.1818	0.1323	0.1405	0.4333±0.003	0.4617±0.003	0.4732±0.003
	MAP	0.0392	0.2297	0.0409	0.0265	0.1385	0.1278	0.1705	0.1010	0.0638	0.0714	0.3185±0.001	0.3453±0.001	0.3726±0.002
SE-6	Acc@1	0.1278	0.1093	0.1278	0.0078	0.0981	0.0830	0.1291	0.0563	0.0865	0.0998	0.2369±0.001	0.2568±0.001	0.2863±0.003
	Acc@5	0.1769	0.3599	0.1775	0.0153	0.1760	0.1721	0.2405	0.1539	0.1795	0.1854	0.3942±0.004	0.4234±0.003	0.4647±0.004
	Acc@10	0.1769	0.4411	0.1782	0.0331	0.2256	0.2130	0.2933	0.2077	0.2262	0.2366	0.4588±0.003	0.4814±0.004	0.5127±0.004
	MAP	0.1567	0.2161	0.1574	0.0233	0.1418	0.1283	0.1876	0.1171	0.1368	0.1485	0.3161±0.001	0.3395±0.002	0.3706±0.002
SE-7	Acc@1	0.0257	0.0719	0.0257	0.0124	0.0948	0.0869	0.1078	0.0586	0.0206	0.0313	0.2231±0.001	0.2407±0.001	0.2709±0.001
	Acc@5	0.0838	0.3638	0.0838	0.0202	0.1318	0.1304	0.1639	0.1245	0.0563	0.0833	0.3297±0.002	0.3663±0.002	0.3895±0.001
	Acc@10	0.1026	0.4106	0.1096	0.0305	0.1544	0.1577	0.1937	0.1614	0.0814	0.1209	0.3667±0.003	0.4012±0.002	0.4148±0.002
	MAP	0.0521	0.1902	0.0539	0.0234	0.1166	0.1120	0.1398	0.0954	0.0455	0.0643	0.2774±0.001	0.3022±0.001	0.3289±0.001

Table 3: Performance for the cold start problem in Acc@1.

	TOP	U-TOP	TAXI	MF	RNN	GRU	LSTM	HST-LSTM	STGN	STGCN	LSTPM	LSTPM-OD	STOD-PPA
SE-1	0.0098	0.0476	0.0098	0.0129	0.1124	0.1162	0.1377	0.0904	0.0664	0.0610	0.2416	0.2469	0.2563
SE-2	0.0152	0.0394	0.0152	0.0113	0.1285	0.1394	0.1599	0.1059	0.0877	0.1020	0.2235	0.2326	0.2324
SE-3	0.0000	0.0572	0.0000	0.0130	0.0952	0.1045	0.1246	0.0965	0.0749	0.0761	0.2319	0.2400	0.2402
SE-4	0.0102	0.0435	0.0102	0.0159	0.1060	0.1103	0.1320	0.0990	0.0737	0.0840	0.2045	0.2114	0.2136
SE-5	0.0000	0.0584	0.0000	0.0178	0.0955	0.1007	0.1219	0.0880	0.0619	0.0599	0.2260	0.2267	0.2360
SE-6	0.3574	0.0534	0.3574	0.0205	0.1491	0.1605	0.1648	0.1201	0.0622	0.0945	0.1836	0.1922	0.2154
SE-7	0.0087	0.0431	0.0090	0.0080	0.0815	0.0822	0.1029	0.0601	0.0369	0.0483	0.1760	0.1838	0.1946

position ($K = 1$) of the predicted ranked set, i.e. given the highest probability, and 0 score otherwise. Unlike Acc@ K that focuses on top K , MAP evaluates the quality of the entire recommendation set and measures the overall performance of the model.

5.3 Experimental Settings

We apply the Adam optimizer with a batch size of 1 user using cross entropy loss for the multi-classification problem, and used 15 epochs and a learning rate of 0.0001 for training. We set our location, user, timeslot, Geohash embedding dimension dim , and the hidden representation dimension $Hdim$ to both be 256. For the map partitioning using Geohash grid, we use Geohash precision of 5, which corresponds to Geohash cells of approximately 4.9km×4.9km. For RNN, GRU, LSTM, and MF, we use the same settings of our model where possible for fair comparison. For all other works, we use their stated recommended settings accordingly.

5.4 Results

We report the evaluation results of our proposed model STOD-PPA and the baselines in Table 2. For our STOD-PPA, as well as the LSTPM-OD and LSTPM baselines, we show the averaged results of 10 runs on different random seeds and with their respective standard deviations. We observe that:

- STOD-PPA outperforms all the baselines, including the state-of-the-art LSTPM and its OD extension of LSTPM-OD significantly for all metrics on all seven datasets of different countries for the origin-aware next destination recommendation task.
- LSTPM-OD is mostly the second best among the results, even when it uses both origin and destination information, the same as our best performing STOD-PPA model. This implies that trivial inclusion of origin information by concatenation is unable to exploit the underlying semantics to best learn OD relationships and perform well for the predictive task.
- Comparing LSTPM and LSTPM-OD, where the only difference is the additional inclusion of origin information in LSTPM-OD, we can see a notable increase of performance for all metrics on

all datasets consistently for LSTPM-OD. This demonstrates that origin information is indeed important and serves as a valuable source of information for the task.

- For the other LSTM based works of STGCN, STGN, and HST-LSTM, these methods do not perform as well. We believe that these methods were unable to learn from the available historical neighbouring intervals of $d_{t_{i-2}}$ and $d_{t_{i-1}}$ to predict d_{t_i} .
- For RNN, LSTM, and GRU baselines, although these models also use both origin and destination information in the same way of concatenation as LSTPM-OD, these baselines, however, do not consider the spatial and temporal factors when learning location-location relationships, whereas LSTPM, LSTPM-OD, and our STOD-PPA do, but in different ways.
- For the baselines of TOP, U-TOP, TAXI and MF, these methods do not learn the sequential transitions between locations, and hence, they do not perform as well.

5.5 Analysis of Cold Start Performance

Same as [25], we evaluate the robustness of our model for the cold start problem by first reversing our existing preprocessing, where instead of removing users with less than 10 trips on the whole dataset, we now only consider these cold start users as they have less than 10 trips. Then, on this test set, we evaluate the methods for the challenging Acc@1 metric as shown in Table 3. For most of the datasets, our STOD-PPA has the best results, followed by mostly LSTPM-OD being the second best. Our STOD-PPA model performed slightly poorer than LSTPM-OD by 0.0002 for the SE-2 dataset and is second best for the SE-6 dataset where TAXI and TOP have the best results. We believe that TOP and TAXI methods only performed well for the SE-6 dataset because the cold start users for this country tend to be tourists taking taxi rides to popular tourist destinations on an ad-hoc basis, where user preferences are easier to predict by city-wide location frequencies.

5.6 Ablation Study

In this section, we perform an ablation study to evaluate the effectiveness of our proposed STOD-PPA model on the same challenging Acc@1 metric for all datasets on their respective test sets.

Spatial and Temporal Factors. We evaluate the effectiveness of our proposed ST-LSTM by replacing all ST-LSTM modules with LSTMs, denoting this variant as OD-PPA (without considering ST factors). Notably, in Fig. 5, for all datasets, our STOD-PPA has an increase of performance over OD-PPA, demonstrating the importance of spatial and temporal factors to learn OD relationships for the recommendation task.

Encoder-Decoder. In Fig. 6, we evaluate the performance of STOD-PPA (encoder-decoder), encoder-only, and decoder-only, where STOD-PPA has the best results by a large margin:

- For decoder-only, we replace Eq. (17) and (18) with $h_{u_m}^O = s_{u_m}^{train^O}$ and $h_{u_m}^D = s_{u_m}^{train^D}$ respectively, effectively forcing the decoder to learn from the input sequences directly where the sequential transitions and its underlying spatial-temporal factors are not considered, hence the poor results.
- For encoder-only, we replace Eq. (20) with $\tilde{y}_{t_i} = \tilde{h}_{t_i}^O \parallel \tilde{h}_{t_{i-1}}^D$ where $\tilde{y}_{t_i}$ is the concatenation of the hidden state from the current

Figure 5: Effectiveness of proposed ST-LSTM.

Figure 6: Performance of encoder-decoder architecture.

Figure 7: Comparison of personalization inclusions.

(a) Size of hidden units $Hdim$. (b) Number of training epochs.

Figure 8: Sensitivity analysis of STOD-PPA in Acc@1.

origin location $\tilde{h}_{t_i}^O$ and the hidden state of the previous destination location $\tilde{h}_{t_{i-1}}^D$ from their own ST-LSTM of $\phi^O(\cdot)$ and $\phi^D(\cdot)$ respectively, before Eq. (21) to predict the next destination d_{t_i} .

Personalization. From Fig. 7, we can see that for all datasets, our STOD-PPA, where user embedding is included in the computation of attention scores for each hidden state on Eq. (19), has the best result as compared to simple addition [15] or concatenation [8] of user embeddings to the hidden representations, specifically, replacing $\tilde{y}_{t_i}$ in Eq. (21) to $\tilde{y}_{t_i} + \tilde{u}_m$ and $\tilde{y}_{t_i} \parallel \tilde{u}_m$ respectively.

Sensitivity Analysis. From Fig. 8(a), we can observe that the model generally converges at $Hdim = 64$ and maintains very similar performances towards $Hdim = 256$ with a clear insensitivity trend. Similarly, in Fig. 8(b), we can see a clear insensitivity trend from epoch 10 onward, indicating that the model has converged.

5.7 Case Study: Interpretable User Preferences

In Fig. 9, we can see a test input tuple of the user ID 3250, previous destination ID 1321 and current origin ID 6, as input to the PPA decoder where it applies the personalized preference attention on all the encoded OD hidden states from the user's historical OD sequences. In the encoder, we can see the corresponding origin and destination ID sequences, as well as the attention weights computed for each hidden state (in percentages for clarity) by the PPA decoder, where a notable difference of weights computed can be observed to best perform the predictive task and can support interpretability (e.g. transition 1671 $\rightarrow$ 1331 has the highest weight and origin ID

Figure 9: Interpretable user preferences with the PPA decoder on a test case from the SE-7 dataset.

79 has the lowest weight). With the ground truth destination ID of 1671, our STOD-PPA approach was able to correctly predict the destination ID 1671 with the highest probability score of 0.93.

6 CONCLUSION

This paper proposed a novel STOD-PPA encoder-decoder model for the origin-aware next destination recommendation task by learning OO, DD, and OD relationships. Specifically, we developed a ST-LSTM encoder module to allow OD relationships to be learned, as well as a PPA decoder module to learn personalized preferences among the encoded OD hidden states. Experimental results on seven real-world datasets from different countries demonstrate the effectiveness of the proposed approach with substantial improvements over existing works. For future work, we plan to study how users' side information can be used to improve the performance of this recommendation task.

ACKNOWLEDGMENT

This work was funded by the Grab-NUS AI Lab, a joint collaboration between GrabTaxi Holdings Pte. Ltd. and National University of Singapore, and the Industrial Postgraduate Program (Grant: S18-1198-IPP-II) funded by the Economic Development Board of Singapore.

REFERENCES

- [1] Buru Chang, Yonggyu Park, Donghyeon Park, Seongsoon Kim, and Jaewoo Kang. 2018. Content-Aware Hierarchical Point-of-Interest Embedding Model for Successive POI Recommendation. In *IJCAI*. 3301–3307.
- [2] Meng Chen, Xiaohui Yu, and Yang Liu. 2019. MPE: A Mobility Pattern Embedding Model for Predicting Next Locations. *WWW* 22, 6 (2019), 2901–2920.
- [3] Chen Cheng, Haiqin Yang, Michael R. Lyu, and Irwin King. 2013. Where You Like to Go Next: Successive Point-of-Interest Recommendation. In *IJCAI*. 2605–2611.
- [4] Eunjoon Cho, Seth A. Myers, and Jure Leskovec. 2011. Friendship and mobility: User movement in location-based social networks. In *KDD*. 1082–1090.
- [5] Kyunghyun Cho, Dzmitry Bahdanau, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *EMNLP*. 1724–1734.
- [6] Jeffrey L. Elman. 1990. Finding Structure in Time. In *COGNITIVE SCIENCE* 14. 179–211.
- [7] Q. Fan, L. Jiao, C. Dai, Z. Deng, and R. Zhang. 2019. Golang-Based POI Discovery and Recommendation in Real Time. In *MDM*. 527–532.
- [8] Jie Feng, Yong Li, Chao Zhang, Funing Sun, Fanchao Meng, Ang Guo, and Depeng Jin. 2018. DeepMove: Predicting Human Mobility with Attentional Recurrent Networks. In *WWW*. 1459–1468.
- [9] Shanshan Feng, Xutao Li, Yifeng Zeng, Gao Cong, Yeow Meng Chee, and Quan Yuan. 2015. Personalized Ranking Metric Embedding for Next New POI Recommendation. In *IJCAI*. 2069–2075.
- [10] Jing He, Xin Li, Lejian Liao, Dandan Song, and William K. Cheung. 2016. Inferring a Personalized Next Point-of-Interest Recommendation Model with Latent Behavior Patterns. In *AAAI*. 137–143.
- [11] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short Term Memory. In *Neural Computation* 9(8). 1735–1780.
- [12] Dejiang Kong and Fei Wu. 2018. HST-LSTM: A Hierarchical Spatial-Temporal Long-Short Term Memory Network for Location Prediction. In *IJCAI*. 2341–2347.
- [13] Y. Koren, R. Bell, and C. Volinsky. 2009. Matrix Factorization Techniques for Recommender Systems. *Computer* 42, 8 (2009), 30–37.
- [14] Nicholas Lim, Bryan Hooi, See-Kiong Ng, Xueou Wang, Yong Liang Goh, Renrong Weng, and Jagannadan Varadarajan. 2020. STP-UDGAT: Spatial-Temporal Preference User Dimensional Graph Attention Network for Next POI Recommendation. In *CIKM*. 845–854.
- [15] Qiang Liu, Shu Wu, Liang Wang, and Tieniu Tan. 2016. Predicting the Next Location: A Recurrent Model with Spatial and Temporal Contexts. In *AAAI*. 194–200.
- [16] Qingjie Liu, Yixuan Zuo, Xiaohui Yu, and Meng Chen. 2019. TTDM: A Travel Time Difference Model for Next Location Prediction. In *MDM*. 216–225.
- [17] Steffen Rendle, Christoph Freudenthaler, and Lars Schmidt-Thieme. 2010. Factorizing Personalized Markov Chains for Next-Basket Recommendation. In *WWW*. 811–820.
- [18] Ke Sun, Tiejun Qian, Tong Chen, Yile Liang, Quoc Viet Hung Nguyen, and Hongzhi Yin. 2020. Where to Go Next: Modeling Long-and-Short-Term User Preferences for Point-of-Interest Recommendation. In *AAAI*. 214–221.
- [19] Pengyang Wang, Yanjie Fu, Hui Xiong, and Xiaolin Li. 2019. Adversarial Sub-structured Representation Learning for Mobile User Profiling. In *KDD*. 130–138.
- [20] Pengyang Wang, Jiawei Zhang, Guannan Liu, Yanjie Fu, and Charu Aggarwal. 2018. Ensemble-Spotting: Ranking Urban Vibrancy via POI Embedding with Multi-view Spatial Graphs. In *SIAM*. 351–359.
- [21] Min Xie, Hongzhi Yin, Hao Wang, Fanjiang Xu, Weitong Chen, and Sen Wang. 2016. Learning Graph-based POI Embedding for Location-based Recommendation. In *CIKM*. 15–24.
- [22] Fuqiang Yu, Lizhen Cui, Wei Guo, Xudong Lu, Qingzhong Li, and Hua Lu. 2020. A Category-Aware Deep Model for Successive POI Recommendation on Sparse Check-in Data. In *WWW*. 1264–1274.
- [23] Quan Yuan, Gao Cong, Zongyang Ma, Aixin Sun, and Nadia Magnenat-Thalmann. 2013. Time-aware Point-of-interest Recommendation. In *SIGIR*. 363–372.
- [24] Yunchao Zhang, Yanjie Fu, Pengyang Wang, Xiaolin Li, and Yu Zheng. 2019. Unifying Inter-region Autocorrelation and Intra-region Structures for Spatial Embedding via Collective Adversarial Learning. In *KDD*. 1700–1708.
- [25] Pengpeng Zhao, Haifeng Zhu, Yanchi Liu, Jiajie Xu, Zhixu Li, Fuzhen Zhuang, Victor S. Sheng, and Xiaofang Zhou. 2019. Where to Go Next: A Spatio-Temporal Gated Network for Next POI Recommendation. In *AAAI*. 5877–5884.
- [26] Han Zhu, Xiang Li, Pengye Zhang, Guozheng Li, Jie He, Han Li, and Kun Gai. 2018. Learning Tree-based Deep Model for Recommender Systems. In *KDD*. 1079–1088.