Itinerary-aware Personalized Deep Matching at Fliggy

Jia Xu
xujia@gxu.edu.cn
College of Computer, Electronics and
Information, Guangxi University
Nanning, China

Ziyi Wang[*]
Alibaba Group
Hangzhou, China
jianghu.wzy@alibaba-inc.com

Zulong Chen
Alibaba Group
Hangzhou, China
zulong.cz@alibaba-inc.com

Detao Lv
Alibaba Group
Hangzhou, China
detao.ldt@alibaba-inc.com

Yao Yu
Alibaba Group
Hangzhou, China
sichen.yy@alibaba-inc.com

Chuanfei Xu
Concordia University
Montreal, Quebec, Canada
chuanfeixu@126.com

ABSTRACT

Matching items for a user from a travel item pool of large cardinality have been the most important technology for increasing the business at Fliggy, one of the most popular online travel platforms (OTPs) in China. There are three major challenges facing OTPs: sparsity, diversity, and implicitness. In this paper, we present a novel Fliggy ITinerary-aware deep matching NETwork (**FitNET**) to address these three challenges. FitNET is designed based on the popular deep matching network, which has been successfully employed in many industrial recommendation systems, due to its effectiveness. The concept *itinerary* is firstly proposed under the context of recommendation systems for OTPs, which is defined as the list of unconsumed orders of a user. All orders in a user itinerary are learned as a whole, based on which the implicit travel intention of each user can be more accurately inferred. To alleviate the sparsity problem, users' profiles are incorporated into FitNET. Meanwhile, a series of itinerary-aware attention mechanisms that capture the vital interactions between user's itinerary and other input categories are carefully designed. These mechanisms are very helpful in inferring a user's travel intention or preference, and handling the diversity in a user's need. Further, two training objectives, i.e., prediction accuracy of user's travel intention and prediction accuracy of user's click behavior, are utilized by FitNET, so that these two objectives can be optimized simultaneously. An offline experiment on Fliggy production dataset with over 0.27 million users and 1.55 million travel items, and an online A/B test both show that **FitNET** effectively learns users' travel intentions, preferences, and diverse needs, based on their itineraries and gains superior performance compared with state-of-the-art methods. **FitNET** now has been successfully deployed at Fliggy, serving major online traffic.

[*]Ziyi Wang is the corresponding author.

CCS CONCEPTS

• **Information systems** → **Recommender systems**; **Personalization**; • **Computing methodologies** → **Neural networks**.

KEYWORDS

personalized Recommendation, neural networks, deep matching, Itinerary-aware

1 INTRODUCTION

Fliggy from Alibaba group[1], is one of the most popular online travel platforms (OTPs) in China, which serves tens of millions of online users by providing million-scale travel-related items (such as air ticket, hotel, and package tour). With the continuous growth in the amount of users and items at Fliggy, it is increasingly critical to help users find items that they might need and prefer to. Therefore, in recent years, great efforts have been paid to develop the Fliggy recommendation system (FRS), which has made significant contribution to the optimization of user experience and the elevation of business value. For example, several important interfaces on Mobile Fliggy App (see Figure 1) use FRS to automatically generate personalized items to satisfy users' needs, and serve major traffic.

With large-scale users and items, the recommendation task in FRS follows the popular two-phase model, which is composed of the matching phase and the ranking phase, used by many mainstream industrial recommendation systems [6, 17], to make sure the recommended items for users are generated in an efficient way. In the matching phase, thousands of candidate items are retrieved from an item pool of million-scale items for each user. Then, the ranking phase is responsible for predicting probabilities of users interacting with these candidate items, and finally generates recommendations. The quality of candidate items retrieved in the matching phase plays a key role in the whole model [19]. Currently, inspired by the success of deep learning from computer vision domain [10] and natural language processing (NLP) [3], many deep

[1]https://www.fliggy.com/

Interface of 'Payment Successful'   Interface of 'My Orders'   Interface of 'Order Details'
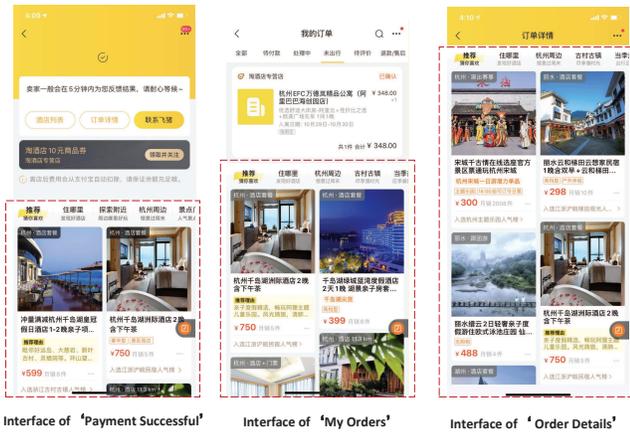
Figure 1: Areas highlighted with dashed rectangles are personalized for tens millions of users in Mobile Fliggy App.

matching networks (DMNs) that match items for users on the basis of deep neural networks have been successfully deployed in many industrial recommendation systems, due to their effectiveness and efficiency [6, 17, 19, 32, 34]. However, without considering the challenges determined by business characteristics of OTPs, these networks become ineffective when the matching objects are travel-related items, which makes the personalized deep matching problem at Fliggy extremely challenging. Major challenges facing OTPs are:

- **Sparsity.** Compared with other e-commerce platforms (e.g., Taobao in China), users' behaviors collected at OTPs are more sparse, for travel is usually an infrequent need of users.
- **Diversity.** When a user plans a travel, the user may have cross-domain diverse needs, including buying air ticket, booking hotel, calling taxi service, choosing a package tour, etc. The diversity of a user's needs must be considered in recommendation process.
- **Implicitness.** According to statistics reported by Fliggy from Aug. 10 to Sep. 10 of the year 2020, over 53% users create more than one order for a travel, and 16% of users even have more than two orders (see Figure 2 for details). The reported results reveal that a user's travel intention, which is of importance in matching items for the user, is generally implicitly delivered by multiple orders rather a single order. For example, Figure 3 displays a user's orders created at Fliggy for a travel. By considering those orders as a whole, we can accurately infer the user's travel intention of the third order is for sightseeing, while the intention of the remaining orders is for transfer. Under such circumstance, tourist items of Nanning (i.e., the transit city) should not be presented to the user.

In this paper, we focus on the above challenges in the matching phase of OTPs and address the above challenges through the creation of a novel Fliggy ITinerary-aware deep matching NETwork called FitNET. On one hand, to cope with the **implicitness** problem, the proposed network uses all unconsumed orders of a user
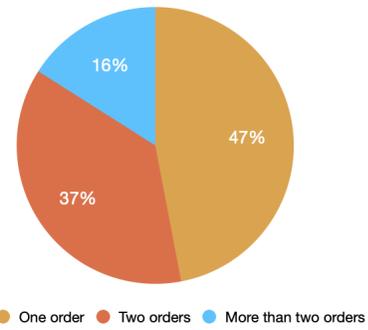

Figure 2: Statistics of orders created by Fliggy users in terms of a travel plan (September 10, 2020 to October 10, 2020).

Order Id	Trajectory	Category	Start Date&Time	End Date&Time
1	Guangzhou→Nanning	Train ticket	2020/1/7 11:47	2020/1/7 15:33
2	Nanning→Bangkok	Air ticket	2020/1/7 19:25	2020/1/7 20:45
3	Bangkok	Package tour	2020/1/8 00:00	2020/1/11 12:00
4	Bangkok→Nanning	Air ticket	2020/1/11 19:15	2020/1/11 22:00
5	Nanning→Guangzhou	Train ticket	2020/1/12 15:40	2020/1/12 19:24

Figure 3: A user's orders created at Fliggy for a travel.

which are deemed to correspond to a travel as a whole input category, defined as *itinerary*, to enhance the understanding of user's intention in each unconsumed order. Besides user's itinerary information, user's profiles which consists of abundant user's side information are incorporated into the framework FitNET, so as to alleviate the **sparsity** and cold start problems. On the other hand, inspired by the recent successes of attention mechanism [14, 30, 33], a series of itinerary-aware attention mechanisms are carefully designed in FitNET to capture abundant interactions between the embedding vectors of users' itineraries and the embedding vectors of other input categories. The proposed attention mechanisms offer great help in improving the accuracy of inferring users' travel intentions, identifying users' preferences to an item's features (such as destination city and category), and understanding users' **diverse** needs. Moreover, the network loss is defined by considering both of the prediction accuracy to users' travel intentions (e.g., sightseeing or others) and the prediction accuracy towards the click behavior of a user to an impression item. These two aspects of loss are seamlessly integrated together under the FitNET framework, so that they can get feedback from each other. The main contributions of this work are summarized below:

- To the best of our knowledge, this is the first work that proposes to leverage a user's *itinerary* (i.e., the list of unconsumed orders of the user) rather than a single order of the user to optimize the deep matching phase of a large-scale recommendation system at OTPs.
- We design a novel itinerary-aware deep matching network (FitNET for short) for Fliggy, one of the most popular online travel platforms (OTPs) in China, to improve the effectiveness of deep neural network in matching personalized travel-related items from a item pool with large cardinality.

In FitNET, unconsumed orders in a user's itinerary is considered as a whole, and a series of carefully designed itinerary-aware attention mechanisms are proposed, which enables FitNET to cope with the three tough challenges facing OTPs effectively.

- We verify FitNET on Fliggy production dataset with over 0.27 million users and 1.55 million travel items. Experimental results show that FitNET effectively learns users' travel intentions and preferences based on their itinerary information and beats state-of-the-art methods. In particular, FitNET improves the average CTR by 1.3% compared with the next best method in an online A/B test, which is an apparent progress made by an industrial recommendation system. Recently, FitNET has been deployed at Fliggy successfully, serving most online traffic.

The remainder of this paper is organized as follows: related works are reviewed in Section 2; Section 3 elaborates the technical details of the FitNET model; In Section 4, we detail the experiments for comparing FitNET with existing methods; The last section gives conclusions and points out future directions of this paper.

2 RELATED WORKS

Most personalized recommendation methods are designed based on the intuition that users' interests can be inferred from their historical behaviors or from other users with similar preferences [36]. Based on such intuition, enormous personalized recommendation systems (RSs) have been proposed and deployed from all walks of life, which consequently makes people's life more convenient.

To further improve users' click-through rate (CTR), post-click conversion rate (CVR), and time spent in applications, collaborative filtering (CF), which makes use of user-item interactions (e.g., clicks) to learn the relevance between users and items, is proposed and soon becomes the key technique for RSs [1, 4, 5, 15, 18, 28]. In particular, the item-to-item (i2i) recommendation, which recommends items similar to an item of interest in terms of different similarity measures, has been the most widely used CF method under the industrial settings [13, 18, 20], owing to its interpretability and efficiency in real-time personalization [21]. Being different to the i2i method, the item-relationship CF recommendation is another popular research domain. It focuses on discovering relationships among items – such as substitutes or complements – that go beyond traditional item similarities [29].

Generally, a recommendation procedure consists of two phases: matching and ranking. In the matching phase, thousands of candidate items are retrieved by matching models from an item pool with large cardinality. During the ranking phase, a more complex ranking model is used to assign comparable scores to the candidate items, and top-k items are recommended to users. After the big success achieved in computer vision domain [10, 16, 27], the deep neural network [9] has become the most popular technique in both academia and industry circles for solving the matching and ranking problems in RSs [6, 7, 11, 17, 19, 32, 34, 35], due to its better representation and generalization ability compared with CF based solutions [19].

Since the quality of retrieved items in the matching phase directly determines the upper quality limit of the ranking phase, matching model plays a vital role in the whole recommendation procedure. Recent years witness a good many works that optimize the matching phase by use of the deep neural network (called as the deep matching network, DMN for short) and achieves remarkable effectiveness. For example, a deep neural network architecture was designed for both of the matching phase and the ranking phase in YouTube recommendation, named YouTube DNN [6]. In YouTube DNN, items and user are mapped to vectors via embedding layers and a feedforward neural network. At serving time, Youtube DNN computes the most likely k videos and presents them to the user. Multi-interest network with dynamic routing (MIND) [17] then was proposed to deal with users' diverse interests in the matching stage. In MIND, a multi-interest extractor layer based on capsule routing mechanism is designed, which is applicable for clustering users' historical behaviors and extracting users' diverse interests. In [24], Bryan et. al proposed the DeepWalk model, a novel method which employs local information obtained from truncated random walk to learn latent representations of vertices in a network. After that, Wang et. al. proposed to apply the DeepWalk as a DMN to improve the relevance of matching results in the matching phase [32]. In specific, they construct an item graph based on users' historical behaviors, and learns the embeddings of all items in the graph. At the same time, Zhang et. al. presented the ENCORE model [34], which is a neural item-relationship based recommender focusing on matching and presenting "complementary" items to users. Then, a sequential deep matching model, SDM [19], was proposed to capture users' dynamic preferences in the matching phase by combining users' short-term sessions and long-term behaviors. Currently, the idea of collaborative filtering is also connected with the deep neural network. For example, collaborative filtering is formulated as a deep neural network in [31], which jointly performs deep representation learning for the content information and collaborative filtering for the ratings (i.e., feedback) matrix. Moreover, a novel CF model, AutoRec, was proposed based on the autoencoder paradigm [25]. And a neural network based collaborative filtering method, NCF, was presented in [8], which complements the mainstream shallow models for collaborative filtering, opening up a new avenue of research possibilities for recommendation based on deep learning.

Inspired by the big success of deep learning network in the circle of recommendation, our FitNET model follows the deep neural network, and focuses on optimizing the matching phase for online travel platforms (OTPs). Although previous DMN-based works provide many good ideas in improving the matching phase, these works can not well address the challenges faced by online travel applications, namely, sparsity, diversity, and implicitness. In this work, we consider these challenges by introducing the concept of '*itinerary*' and carefully design a series of effective attention mechanisms by being aware of the information of user itinerary. In the following, we present the design of our FitNET model.

3 FITNET MODEL

In this section, we first formalize the personalized matching problem faced by Fliggy Recommendation System (FRS), and then present the FitNET model which effectively solves the problem at Fliggy.

3.1 Problem Formalization

Let $\mathcal{U}$ be the set of users at Fliggy. The objective of personalized matching in FRS is to improve the quality of personalized item set retrieved from an item pool $\mathcal{P}$ of million-scale items related to travel matters for each user $u_i \in \mathcal{U}$, by addressing before-mentioned three challenges facing OTPs: namely, sparsity, diversity, and implicitness. It is known that the quality of the retrieved personalized item set is highly influenced by the inferred travel intention of the user.

To achieve the objective of personalized matching in FRS, this paper designs a novel model named Fliggy itinerary-aware deep matching network (FitNET), which overcomes the challenges facing OTPs by introducing users' itinerary information into the neural network. The concept of *itinerary* is formally defined as

DEFINITION 1. ***Itinerary***. *For a given user $u_i \in \mathcal{U}$, let $\bar{o}_{u_i}$ denotes an order of u_i at Fliggy which has not yet been consumed by u_i. The itinerary of u, denoted by $\mathcal{I}_{u_i}$, is the list of unconsumed orders of u_i at Fliggy (i.e., $\mathcal{I}_{u_i} = \{\bar{o}_{u_i}\}$) which is deemed to correspond to a travel of the user.*

The implicitness problem of OTPs tells that it is quite difficult to infer a user's travel intention with a single order of the user. The introduction of 'itinerary' enables FitNET to be trained based on the granularity of a user's itinerary rather than the granularity of the user's single order. As a result, FitNET can make better prediction to the travel intention of each user, which tackles the implicitness challenge facing OTPs.

Besides the itinerary of each user (i.e., $\mathcal{I}_{u_i}, u_i \in \mathcal{U}$), basic profile of every user is also used as the input, so as to lighten the sparsity challenge due to the inadequate amount of user behaviors collected at OTPs. Let $\mathcal{P}_{u_i}$ denote the basic profile of user u_i (e.g., permanently reside city id, home city id, and etc.), $\mathcal{B}_{u_i}$ be a sequence of items interacted by user u_i before the first order in $\mathcal{I}_{u_i}$ is created by u_i (also known as the user behavior sequence), and $\mathcal{F}_{u_i}$ represent the target item features (e.g., item id, category id, and etc.). Then, every input instance of historical data in FitNET can be represented by a quadruple $(\mathcal{P}_{u_i}, \mathcal{I}_{u_i}, \mathcal{B}_{u_i}, \mathcal{F}_{u_i})$.

The critical task of FitNET is to learn a function for mapping user-related raw inputs into user representations, which can be formulated by

$$v_{u_i} = f_{\mathcal{U}}(e(\mathcal{P}_{u_i}), e(\mathcal{I}_{u_i}), e(\mathcal{B}_{u_i})), \tag{1}$$

where, v_{u_i} is the representation vector of user u_i computed based on the embedding vectors of $\mathcal{P}_{u_i}$, $\mathcal{I}_{u_i}$, and $\mathcal{B}_{u_i}$, represented as $e(\mathcal{P}_{u_i})$, $e(\mathcal{I}_{u_i})$, and $e(\mathcal{B}_{u_i})$, respectively. Besides, let $\mathcal{T}$ be the set of target items, and v_{t_j} be the representation vector of a target item $t_j \in \mathcal{T}$. v_{t_j} can be obtained by the learned mapping function $f_{\mathcal{T}}$ of FitNET with the embedding vector of target item features $\mathcal{F}_{t_j}$, denoted as $e(\mathcal{F}_{t_j})$, as the input (see Equation 2 for details). Specific implementation of the embedding strategy will be elaborated in Section 3.3.

$$v_{t_j} = f_{\mathcal{T}}(e(\mathcal{F}_{t_j})) \tag{2}$$

In FitNET, the representation vectors for both users and items are optimized by the employment of a group of itinerary-aware attention mechanisms, which leverage users' itinerary information to improve the embeddings of intermediate results. These attention mechanisms offer great help not only to the inference of users' travel intentions or preferences, but also to the identification of a user's diverse needs, which addresses the diversity challenge faced by OTPs. When the representation vector of a user $u_i \in \mathcal{U}$ and the representation vectors of items $t_j \in \mathcal{P}$ are learned by FitNET, the top k items in $\mathcal{P}$ are presented as the matched results for the user u_i, according to the score function:

$$f_{score}(v_{u_i}, v_{t_j}) = v_{u_i} \cdot v_{t_j}. \tag{3}$$

where the operator $\cdot$ represents the inner product of two vectors.

In the following we present details of the proposed FitNET model for the personalized matching task at Fliggy. The architecture of the FitNET is illustrated in Figure 4. There are four categories of input data in FitNET. It is worth noting that features of input data are either ID features (e.g., "home_city_id" and "item_id") or dense features (e.g., "gender" and "age_level").

3.2 Feature Representation

Unlike sponsored search, users come into the recommendation interfaces of FRS without clearly expressing their intentions, which matters for the quality of personalized matching. Therefore, features of input data in FRS that depict the intentions of users are the critical elements in FitNET. Besides user basic profile and user behavior sequence whose features are frequently used by popular industrial recommendation systems [19, 35], this work proposes to utilize features of a user's itinerary (i.e., features of the user's unconsumed orders) to better capture the user's travel intention. According to Definition 1, a user's itinerary is a list of unconsumed orders of the user at Fliggy. Therefore, features of a user's itinerary are actually defined as the features of an order. Similarly, features of a user's behavior sequence including a group of interacted items of the user are defined as the features of an item. Table 1 summarizes the whole feature set used by FitNET. Features in the feature set come from four categories, namely, user's basic profile $\mathcal{P}_{u_i}$, user's itinerary $\mathcal{I}_{u_i}$, user's behavior sequence $\mathcal{B}_{u_i}$, and target item features $\mathcal{F}_{t_j}$. Features belonging to the same category consist of a feature group, and they describe the category from different side. In FRS, the value of each feature (e.g, home_city_id="Hangzhou") is transformed into a high-dimensional sparse binary vector (e.g., $[0, \ldots, 1, \ldots, 0]$ denotes the city Hangzhou in China) by encoding techniques [22, 26]. Let x_{ij} be the binary encoding vector of the j-th feature in the i-th feature group, and $k = \sum_n x_{ij}[n]$. Then, vector x_{ij} with $k = 1$ refers to a one-hot encoding type, and vector x_{ij} with $k > 1$ refers to a multi-hot encoding type. Table 1 shows that features of user's itinerary or user's behavior sequence are multi-hot encoding vectors, containing rich information of users' preferences. With the sparse binary representations of each feature, we apply a novel deep neural network – FitNET, to capture the interactions of these features.
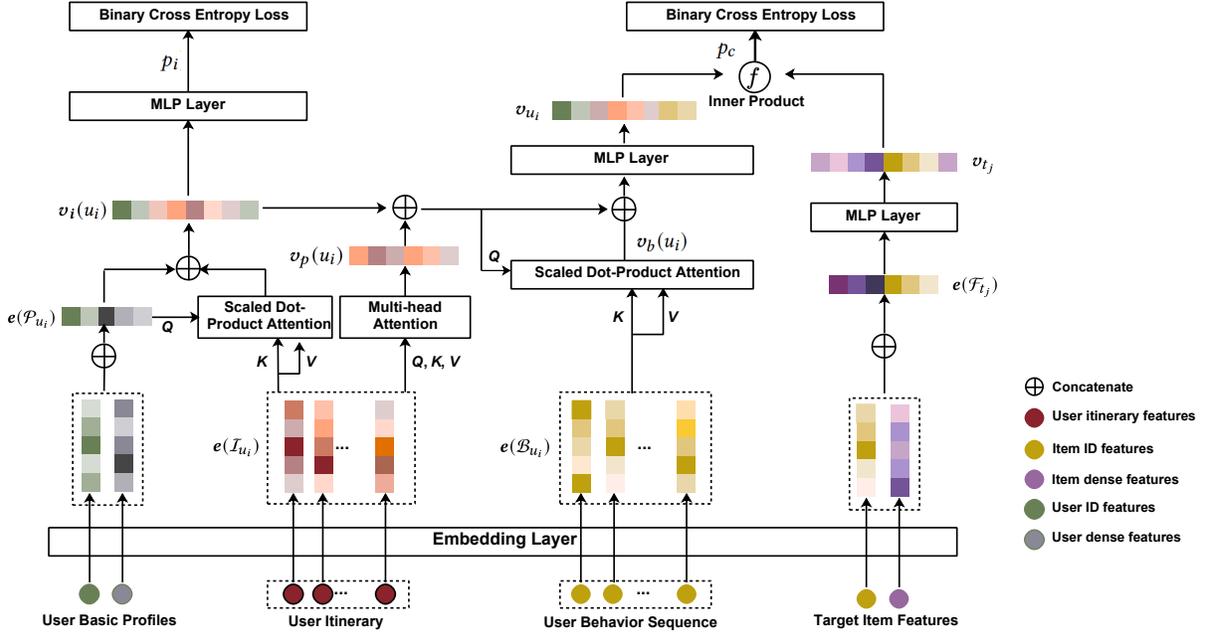
Figure 4: The architecture of the proposed FitNET model.

Table 1: Statistics of features sets used in FitNET. Each feature is encoded into a binary sparse vector.

Category	Feature Group	Dimensionality	Type	# Nonzero Ids per Instance
User's Basic Proflile ($\mathcal{P}_{u_i}$)	gender	2	one-hot	1
	age_level	~ 10	one-hot	1
	reside_city_id	$\sim 10^3$	one-hot	1
	home_city_id	$\sim 10^3$	one-hot	1
	...	...	...	...
User's Itinerary ($\mathcal{I}_{u_i}$)	ordered_item_ids	$\sim 10^6$	multi-hot	~ 10
	ordered_cate_ids	$\sim 10^3$	multi-hot	~ 10
	ordered_dest_city_ids	$\sim 10^3$	multi-hot	~ 10
	...	...	...	...
User's Behavior Sequence ($\mathcal{B}_{u_i}$)	clicked_item_ids	$\sim 10^6$	multi-hot	~ 10
	clicked_cate_ids	$\sim 10^3$	multi-hot	~ 10
	clicked_dest_city_ids	$\sim 10^3$	multi-hot	~ 10
	...	...	...	...
Target Item Features ($\mathcal{F}_{t_j}$)	item_id	$\sim 10^6$	one-hot	1
	cate_id	$\sim 10^3$	one-hot	1
	dest_city_id	$\sim 10^3$	one-hot	1
	...	...	...	...

3.3 Embedding Layer

As shown in Figure 4, the input historical data of FitNET are composed of four categories. Features of input data from each category are transformed into feature representations in the form of high-dimensional sparse binary vectors, especially for the representations of ID features. For instance, the number of item ids is about millions at Fliggy, which motivates us to apply the widely-used embedding technique [23] to embed the feature representations into low-dimensional dense vectors, so as to markedly reduce the

cardinality of parameters and facilitate the training and serving of FitNET.

For the representation vector of the j-th feature in the i-th feature group (i.e., x_{ij}), let $W^{ij} = [w_1^{ij}, \cdots, w_n^{ij}, \cdots, w_{D_{ij}}^{ij}] \in \mathbb{R}^{D_{ij} \times N_{ij}}$ denote the embedding dictionary for the vector, where $w_n^{ij} \in \mathbb{R}^{D_{ij}}$ is an embedding vector with dimensionality of D_{ij}, and N_{ij} is the number of distinct values x_{ij} might be assigned. We segregate feature representations into two types, whose representations are embedded in different ways.

- **One-hot encoding type.** If the vector x_{ij} is one-hot encoding with $x_{ij}[n] = 1$, then the embedding representation of x_{ij}, i.e., $e(x_{ij})$, is the single embedding vector w_n^{ij} in W^{ij}.
- **Multi-hot encoding type.** If the vector x_{ij} is multi-hot encoding with $x_{ij}[n] = 1$ and $n \in \{n_1, \cdots, n_k\}$, then the embedding representation of x_{ij} is the concatenation of a group of embedding vectors in W^{ij}, represented by $e(x_{ij}) = \{w_{n_1}^{ij}, \cdots, w_{n_k}^{ij}\}$.

Embedding vectors of features belonging to the same input category are further concatenated during the embedding process. After embedding, we obtain one embedding vector $e(\mathcal{P}_{u_i}) \in \mathbb{R}^{D_p}$ for the input category "user's basic profile" of each $u_i \in \mathcal{U}$, and one embedding vector $e(\mathcal{T}_{t_j}) \in \mathbb{R}^{D_t}$ for the input category "target item features" of each item $t_j \in \mathcal{P}$. As to the input category "user's itinerary" that consists of a list of user u_i's unconsumed orders, its embedding result is a list of embedding vectors for the unconsumed orders, i.e, $e(\mathcal{I}_{u_i}) = \{e(\bar{o}_{u_i,1}), \cdots, e(\bar{o}_{u_i,l}), \cdots, e(\bar{o}_{u_i,|\mathcal{I}_{u_i}|})\}$ with $e(\bar{o}_{u_i,l}) \in \mathbb{R}^{D_e}$. The embedding vector of an unconsumed order is the concatenation of embedding vectors of all order features. With regard to the input category "user's behavior sequence" that is a sequence of items interacted by user u_i, its embedding result is a sequence of embedding vectors for those items, i.e., $e(\mathcal{B}_{u_i}) = \{e(t_{u_i,1}), \cdots, e(t_{u_i,m}), \cdots, e(t_{u_i,|\mathcal{B}_{u_i}|})\}$ with $e(t_{u_i,m}) \in \mathbb{R}^{D_b}$. The embedding vector of an interacted item is the concatenation of embedding vectors of all its features. Here, $D_* (* \in p, e, b, t)$ denotes the vector dimension and $| * |$ represents the cardinality of the set $*$.

3.4 Itinerary-Aware Attention Mechanisms

Since the information of user itineraries plays an important role in inferring users' travel intentions, a series of itinerary-aware attention mechanisms which learn the interactions between the embedding vectors of users' itineraries and the embedding vectors of other input categories are applied in FitNET.

3.4.1 Interactions of $e(\mathcal{P}_{u_i})$ and $e(\mathcal{I}_{u_i})$. In FitNET, the scaled dot-product attention [30] is applied to the embedding vectors of users' basic profiles and uses' itineraries (i.e., $e(\mathcal{P}_{u_i})$ and $e(\mathcal{I}_{u_i})$ for each $u_i \in \mathcal{U}$), to better understand users' travel intentions on the premise of knowing their basic profiles. Take the "home_city_id" feature in user's basic profile for example. If a user's home city is "Hangzhou" while the user's destination city in an unconsumed order of his/her itinerary is also "Hangzhou", then FitNET can infer the user's travel intention as visiting relatives and thus do not recommend local tourist attractions in Hangzhou to the user. In the attention mechanism, $e(\mathcal{P}_{u_i}) \in \mathbb{R}^{D_p}$ is the query vector and $e(\mathcal{I}_{u_i}) = \{e(\bar{o}_{u_i,l})\}_{l=1}^{|\mathcal{I}_{u_i}|}$ with each $e(\bar{o}_{u_i,l}) \in \mathbb{R}^{D_e}$ is used as both the key vector and the value vector, as shown in Figure 4. The attention function on the set of queries is

$$\textbf{Attention}(Q, K, V) = \textbf{Softmax}\left(\frac{e(\mathcal{P}_{u_i})\left[e(\mathcal{I}_{u_i})\right]^T}{\sqrt{d}}\right) e(\mathcal{I}_{u_i})$$
$$= \sum_{l=1}^{|\mathcal{I}_{u_i}|} \alpha_l \times e(\bar{o}_{u_i,l}), \quad (4)$$

where d is a hyperparameter, and

$$\alpha_l = \frac{exp\left(score(e(\mathcal{P}_{u_i}), e(\bar{o}_{u_i,l}))\right)}{\sum_{l'=1}^{|\mathcal{I}_{u_i}|} exp\left(score(e(\mathcal{P}_{u_i}), e(\bar{o}_{u_i,l'}))\right)},$$

$$score(e(\mathcal{P}_{u_i}), e(\bar{o}_{u_i,l})) = e(\mathcal{P}_{u_i}) W_1 \left[e(\bar{o}_{u_i,l})\right]^T.$$

Here, the matrix $W_1 \in \mathbb{R}^{D_p \times D_e}$ is a learnable parameter in the attention mechanism.

The output vector of this attention mechanism is further concatenated with the embedding vector of user u_i's basic profile to get user u_i's intention vector, denoted by $v_i(u_i)$.

3.4.2 Self-interactions of $e(\mathcal{I}_{u_i})$. Note that a single unconsumed order of a user cannot accurately describe the user's preference for itinerary (such as the preference for destination city or for the item category). Hence, as shown in Figure 4, a multi-head self-attention mechanism [30] is employed in FitNET to capture the self-interactions among embedding vectors of unconsumed orders (i.e., $\{e(\bar{o}_{u_i,l})\}_{l=1}^{|\mathcal{I}_{u_i}|}$) in a user's itinerary, so as to better understand users' purchasing preferences. In the multi-head self-attention mechanism, we employ N heads, where N is set to 8 by following the recommended setting in [30]. The output of each head is concatenated and once again projected using a parameter matrix W^O, resulting in an output vector $v_p(u_i) \in \mathbb{R}^{D_m}$ for each user $u_i \in \mathcal{U}$. The attention function is

$$\textbf{MultiHead}(Q, K, V) = v_p(u_i)$$
$$= \left[v_p^1(u_i) \oplus \cdots \oplus v_p^N(u_i)\right] W^O, \quad (5)$$

where $\oplus$ represents the operation of concatenating two vectors, and the n-th head $v_p^n(u_i) \in \mathbb{R}^{\frac{D_m}{N}}$ in Equation 5 is

$$v_p^n(u_i) = \textbf{Attention}\left(e(\mathcal{I}_{u_i})W_n^Q, \ e(\mathcal{I}_{u_i})W_n^K, \ e(\mathcal{I}_{u_i})W_n^V\right), \quad (6)$$

and $W_n^Q, W_n^K, W_n^V \in \mathbb{R}^{D_e \times \frac{D_m}{N}}$ are the learnable project matrices of the n-th head.

3.4.3 Interactions of $e(\mathcal{P}_{u_i})$, $e(\mathcal{I}_{u_i})$, and $e(\mathcal{B}_{u_i})$. As one of important inputs of FitNET, a user's behavior sequence, however, usually contains causal interactions to items that are unrelated to the user's itinerary. To filter out the negative impact brought by causal interactions on the acquirement of users' diverse needs, as illustrated in Figure 4, FitNET adopts another scaled dot-product attention mechanism to emphasize more important interacted items in a user's behavior sequence, according to the user's itinerary information. As indicated in the figure, the user intention vector $v_i(u_i)$ that encodes both of the basic profile (i.e., $e(\mathcal{P}_{u_i})$) and the travel intention information of user u_i is concatenated with the output vector of the multi-head attention mechanism $v_p(u_i)$ that encodes the preferences of the user, and then the concatenation result (denoted by $v_c = \left[v_i(u_i) \oplus v_p(u_i)\right]$) is used as the query vector in the attention mechanism. The embedding vector of the user u's behavior sequence (i.e., $e(\mathcal{B}_{u_i}) = \{e(t_{u_i,m})\}_{l=1}^{|\mathcal{B}_{u_i}|}$ with each $e(t_{u_i,m}) \in \mathbb{R}^{D_b}$) is used as both the key vector and the value vector. The attention function on the set of queries is

$$\mathbf{Attention}(Q, K, V) = \mathbf{Softmax}\left(\frac{v_c\left[e(\mathcal{B}_{u_i})\right]^T}{\sqrt{d}}\right)e(\mathcal{B}_{u_i})$$
$$= \sum_{m=1}^{|\mathcal{B}_{u_i}|} \beta_m \times e(t_{u_i,m}), \tag{7}$$

where d is a hyperparameter, and the weight β_m for the m-th interacted item in the user behavior sequence is computed by

$$\beta_m = \frac{exp\left(score(v_c, e(t_{u_i,m}))\right)}{\sum_{m'=1}^{|\mathcal{B}_{u_i}|} exp\left(score(v_c, e(o_{u_i,m'}))\right)},$$

$$score(v_c, e(t_{u_i,m})) = v_c \ W_2 \ \left[e(t_{u_i,m})\right]^T,$$

where the matrix $W_2 \in \mathbb{R}^{D_p \times D_b}$ is a learnable parameter in the attention mechanism. The output vector of this attention mechanism is represented as $v_{(u_i)}$ in Figure 4.

To sum up, by introducing users' itineraries into the attentive procedure over user's behavior sequence, users' diverse needs are more precisely embedded, addressing the diversity problem faced by OTPs.

3.5 MLP Layer

The Multi-Layer Perceptron (MLP) layer is a fully-connected feed-forward neural network which employs multiple layers of neurons with nonlinear activation functions, and therefore is more powerful to learn combination of features automatically [2]. As indicated in Figure 4, there are totally three MLP layers in FitNET. The first MLP layer utilizes user intention vectors (i.e., $\{v_i(u_i)\}$ with each $u_i \in \mathcal{U}$) as the input, and is responsible for learning non-linear relationships of features in these input vectors and finally derives a sightseeing probability by a sigmoid function, denoted by p_i, representing the travel intention of user u_i. The input of the second MLP layer is the concatenation of three intermediate result vectors in FitNET, i.e., $[v_i(u_i) \oplus v_p(u_i) \oplus v_b(u_i)]$. Since the vector $v_i(u_i)$ encodes the information of the user's travel intention, the vector $v_p(u_i)$ contains the information of the user's preferences, and the vector $v_b(u_i)$ encompasses the information of the user's interacted items, the purpose of this MLP layer is to learn a high-quality representation for each user, denoted as v_{u_i}, by taking those information into account. Finally, the third MLP layer takes the embedding vectors of target item features as the input, and learns a better representation for each item, denoted as v_{t_j}. After the MLP processing, the learned representations of all users and the learned representations of all target items are executed pairwise calculations according to the inner product operation shown in Equation 3. The inner product between a user u_i's representation v_{u_i} and a target item t_j's representation v_{t_j} is viewed as the probability of the target item t_j being clicked by the user u_i in the recommendation phase, which is denoted by p_c in Figure 4.

3.6 Training & Serving

In the training phase, the FitNET model is trained according to a loss function that consists of two parts, as shown in Equation 8.

$$Loss = Loss_i + Loss_c \tag{8}$$

The first part in the loss function, i.e., $Loss_i$, corresponds to the left 'Binary Cross Entropy Loss' module in Figure 4. It utilizes the pre-known users' travel intentions with respect to each of their orders as labels to compute the loss caused by predicting users' travel intentions. An indicator variable I_i is employed to label each order, with $I_i = 1$ presenting the current user's intention of creating the current order is for sightseeing, and 0 otherwise. Specifically, the loss L_i is calculated by

$$Loss_i = -I_i \log p_i - (1 - I_i) log(1 - p_i), \tag{9}$$

where p_i is the output of the leftmost MLP layer in FitNET, representing the sightseeing probability of a user predicted by FitNET.

The second part of the loss function, i.e., $Loss_c$, corresponds to the right 'Binary Cross Entropy Loss' module in Figure 4. It makes use of users' click behaviors of to recommended items as the labels which are collected before the users consume any order in their itineraries, so as to measure the loss introduced by predicting users' click behaviors to recommended items. Similarly, an indicator variable I_c is used to label each <user-item> pair, with $I_c = 1$ indicating the current user will click the current recommended item, and 0 otherwise. Specifically, the loss $Loss_c$ is computed by

$$Loss_c = -I_c \log p_c - (1 - I_c) log(1 - p_c), \tag{10}$$

where p_c is the output of the inner product operation in FitNET, representing the click probability of a user to a recommended item predicted by FitNET.

After training, FitNET can be used as user representation mapping function $f_{\mathcal{U}}$ and item representation mapping function $f_{\mathcal{T}}$. At serving time, the information of a user's basic profile, itinerary, and behavior sequence are fed into the $f_{\mathcal{U}}$ function to derive the representation vector for each user. Meanwhile, the features of a item are input into the $f_{\mathcal{T}}$ function to get the representation vector for every item in the item pool $\mathcal{P}$. Then, these representation vectors are utilized to match top k items from the item pool for each user in terms of the user's travel intention and preference, according to an approximation nearest neighbor search approach [6]. Note that when a user offers new click behaviors, the user's representation vector will accordingly be altered, which means FitNET enables real-time personalization for the matching phase of FRS.

4 EXPERIMENTS

In this section, we evaluate the effectiveness of the proposed Fit-NET model extensively. At first, FitNET is compared with several existing methods on an offline real-world dataset. Then, an online A/B test of all comparison models is reported at Fliggy's live production environment.

4.1 Comparison Methods

We compare the proposed FitNET with five related methods.

- **FitNET⁻**: FitNET⁻ is a variant of FitNET. Being different from FitNET, only the most recent unconsumed order of a user is used as the input of FitNET⁻, while FitNET takes all unconsumed orders of a user as a whole (i.e., the user's itinerary) and fed them to the network.
- **YouTube DNN**[6]: YouTube DNN model is able to effectively assimilate many signals and model their interaction

with layers of depth, which is one of the most successful deep learning methods used for industrial recommendation systems.

- **DeepWalk**[32]: DeepWalk is one of the most popular methods in commodity recommendation. It builds an item graph from users' behavior history, and learns the embedding of all items by considering the sequential information of items recorded in the graph. Item embeddings are employed to compute pairwise similarity between any two items, which are then utilized to achieve the recommendation.
- **MIND**[17]: MIND is deep neural network with behavior-to-interest dynamic routing which deals with user's diverse interests in the matching stage. MIND is proven to get superior performance than other state-of-the-art methods for recommendation. As the interests of users at OTPs are not as divergent as that in other online e-commerce platforms (such as Taobao in China), the number of extracted interest in MIND is set to 1 (i.e., $k = 1$) in the experiment, which has been verified to be optimal setting for the hyperparamter k at Fliggy platform.
- **OrderDest2i**: OrderDest2i is a weak-personalized recommendation method which is a widely used by industrial recommendation systems. OrderDest2i recommends most popular items with respect to the destination city of a user's order, and matched items are further ranked based on their popularity scores.

4.2 Offline Experiments

4.2.1 Dataset. We employ the Fliggy production dataset in the offline evaluation. The dataset is extracted from **Mobile Fliggy App**, contains one month users' logs at Fliggy collected in April 2020. In the dataset, impression and click samples are labeled as positive samples. On the premise that the number of positive samples is constant, we test the impact of different settings for the amount of negative samples on the performance of FitNET in Section 4.2.3, and list the best setting in Table 2. The dataset is further organized into the training dataset and the testing dataset. It is worth noting that an algorithm is designed to automatically label each itinerary of a user in the training dataset with a travel intention, according to a group of well-defined rules set by our experienced staffs. The effectiveness of the algorithm in identifying a user's travel intention is verified by industrial practices. More details of the Fliggy production dataset are summarized in Table 2.

Table 2: Statistics of Fliggy production dataset. (M - Million)

Dataset features	Training	Testing
# of samples	12.1M	3.4M
# of positive samples	1.3M	0.4M
default # of negative samples	11.8M	3M
# of users	0.27M	0.1M
# of items	1.53M	1.51M
# of orders	0.73M	0.26M
# of itineraries	0.34M	0.11M

4.2.2 Metrics. In the offline setting, two metrics, i.e., ***HitRate***@***k*** and ***Precision***@***k***, are adopted to measure the recommendation performance of different methods, which are also widely used in other related works mentioned in Section 2. *HitRate*@*k* represents the proportion of test cases that have the correctly recommended items in the top k recommendation list of a user, defined as

$$HitRate@k = \frac{\sum_{(u_i,t_j) \in \mathcal{S}_{test}} I(\text{target item occurs in top-}k)}{|\mathcal{S}_{test}|}, \quad (11)$$

where $\mathcal{S}_{test}$ denotes the test set consisting of test cases with each test case being in the form of a user-item pair (u_i, t_j), and I is the indicator function.

Precision@*k* reflects how many items a user click on within the confines of the top-k recommended items for the user. It is calculated by

$$Precision@k = \frac{1}{|\mathcal{U}|} \sum_{u_i \in \mathcal{U}} \frac{|\mathcal{S}_c(u_i)@k \cap \mathcal{S}_{gt}(u_i)|}{k}, \quad (12)$$

where $|\mathcal{U}|$ is the cardinality of the user set $\mathcal{U}$, $\mathcal{S}_c(u_i)@k$ denotes the predicted top-k recommendation items for user u_i, and $\mathcal{S}_{gt}(u_i)$ represents the ground truth items that are clicked by user u_i.

4.2.3 Evaluation of Different Settings of Negative Samples. With regard to the training of FitNET, positive samples can be directly defined as users' click items with fix cardinality of 1.7M (see Table 2 for details), since clicks indicate users' feedback to a impression item (i.e., recommended item) being a likely match to users' intention. However, defining negative samples for the matching phase in a recommendation system is a non-trivial problem [7, 12], for the quality of selected negative samples has great impact on the performance of a recommendation model. For example, a model trained by using the non-click impressions as its negative samples displays the significantly worse performance in retrieving items for users [12], compared with using random samples as negatives. In view of this, before we make the comparison of all methods, we conduct an experiment to evaluate the influence of different settings to negative samples on FitNET. In particular, we verify three settings of negative samples:

Setting 1: For each positive sample, we randomly sample items from the item pool as negative samples.

Setting 2: For any user, the user's negative samples are limited to items whose destination city is consistent with that of the user's itinerary. Besides, we ensure that the categories of 50% negative samples are the same as that of their corresponding positive samples and the categories of the remaining 50% negative samples are different from that of their corresponding positive samples.

Setting 3: The destination city of a user's negative samples is also confined to that of the user's itinerary. And we make sure that only 10% negative samples have the same categories as that of their corresponding positive samples and 90% negative samples have different categories compared with that of their corresponding positive samples.

The experimental results in terms of the above settings are listed in Tables 3-5. We derive three important observations from the three tables:

1) FitNET achieves significantly better performance under the settings 2 and 3. This is because there is no restriction on the destination city of negatives in setting 1, and as a result a good many samples whose destination city is irrelevant to the user's itinerary are selected as the negatives. This indicate that confining the destination city of negative samples to that of the user's itinerary is of vital importance in training a travel-related matching model.
2) The amount of negative samples has apparent impact on the hit rate and the precision of FitNET. In specific, with the increase of negative samples, two metrics are improved at first, since the increase of negative samples enhances the generalization ability of the trained model in differentiating between positives and negatives. However, when the ratio of negative samples across $10/11$, there is not obvious additional gain over the FitNET model. Therefore, $1:10$ is applied as the fix setting to the amount of negative samples in the following experiments.
3) FitNET gains its best performance under the setting 3. Thus we can conclude that by augmenting the ratio of negative samples that have different categories compared with that of their corresponding positive items can improve the ability of the model in distinguishing negative items. In the following experiments, the third setting of negative samples is utilized.

4.2.4 Ablation Study. To evaluate the effectiveness of the proposed three itinerary-aware attention mechanisms in the FitNET model (see Section 3.4 for details), we conducted an ablation study by comparing FitNET with its three variants, named as Var_1, Var_2, and Var_3. Var_1 removes the attention mechanism proposed in Section 3.4.1 from FitNET, which captures the interactions between users' itineraries and their basic profiles. Note that, since the removed attention mechanism is highly related to the auxiliary tower in the FitNET architecture, corresponding to the $Loss_i$ part in the loss function of FitNET (see Equation 8) which is set to minimize the loss of predicting users' travel intentions, the auxiliary tower is deleted from the architecture of FitNET in Var_1. Var_2 removes the multi-head attention mechanism proposed in Section 3.4.2 from FitNET, which learns the self-interactions among embedding vectors of unconsumed orders in a user's itinerary. Finally, Var_3 cuts out the attention mechanism proposed in Section 3.4.3 from FitNET, which emphasizes more important interacted items in a user's behavior sequence, according to the interactions among the user's itinerary, profile, and behavior sequence. Experimental results of the ablation study are listed in Table 6. Table 6 shows that:

1) Each of the three attention mechanisms makes considerably contribution to improve the hit rate and precision, due to the concern of interactions between itineraries and other input categories, which confirms again the important role of itineraries in improving the matching phase at OTPs.
2) Without the attention mechanism proposed in Section 3.4.1 and the auxiliary tower that minimizes the loss in identifying a user's travel intention, there is a relatively bigger decrease in hit rate and precision, which on one hand demonstrates the importance

of retrieving items for users in terms of their travel intentions and on the other the effectiveness of using the itinerary in referring travel intentions.

4.2.5 Evaluation of Comparison Methods. Results on the offline dataset of different methods are shown in Table 7. In general, OrderDest2i which executes personalized matching simply based on the destination city of the order is beaten by all deep learning based methods (i.e., FitNET, FitNET$^-$, MIND, YouTube DNN, and DeepWalk), revealing the power of deep learning for improving the matching phase of recommendation systems. It can also be observed that methods aware of the information of user's itinerary (i.e., FitNET and FitNET$^-$) apparently outperform other methods. This confirms that the consideration of user's itinerary in the design of a deep matching network is of great help in solving the three challenges, namely, sparsity, diversity, and implicitness, facing OTPs. Another important observation from the table is that FitNET gains higher hit rate and precision compared with its variant, FitNET$^-$. Restate that FitNET$^-$ trains the network only based on the most recent unconsumed order of a user, while all unconsumed orders in a user's itinerary are used in FitNET. According to Fliggy's experience, a user's travel intention with respect to the most recent order is probably returning home rather than going sightseeing, and users usually show very low interest in the recommended items relating to their homes. Therefore, compared with FitNET$^-$, FitNET gets on average 1.4% improvement in terms of hit rates, and on average 1.7% growth in terms of precision. This observation once again confirms that all unconsumed orders of a user must be treated as a whole to elevate the prediction accuracy of the user's travel intention. The table also demonstrates that MIND beats YouTube DNN due to the usage of label-aware attention mechanism over the embedding of users' behavior sequences, which decreases the importance of irrelevant interacted items of a user and as a result better identifies a user's diverse interests. On the other hand, YouTube DNN simply uses the average pooling strategy over the embeddings of users' behavior sequences, reducing the quality of its matched items. Finally, it is worth noting that FitNET makes on average 5.2% increase in hit rate and on average 9.1% improvement in precision, compared with the next best method, MIND. Because, FitNET is able to learn more accurate information of users' travel intentions and preferences, by being awareness of users' itinerary information in the design of the matching framework for OTPs.

4.3 Online A/B Test

We conduct online experiments by deploying FitNET to handle real traffic in the personalized interfaces of Mobile Fliggy App (see Figure 1) for seven days in May 2020. We deploy all comparison methods concurrently in the matching phase of Fliggy platform under an A/B test framework, and one thousand of candidate items are retrieved by each method in the matching phase, which then fed to the ranking phase. To guarantee the fairness of comparisons, retrieved items by each method are ranked by the same ranking engine. Moreover, the scheduling engine is revised to make sure the personalized interfaces of each method obtain 20% daily traffic at Fliggy, to make sure that each method has equal traffic.

Table 3: Experiments by varying the amount of negative samples under the setting 1.

+:-	Hit Rate				Precision			
	@3	@10	@20	@50	@3	@10	@20	@50
1:1	0.071	0.157	0.23	0.406	0.029	0.025	0.017	0.011
1:2	0.074	0.163	0.234	0.413	0.031	0.028	0.018	0.013
1:5	0.076	0.165	0.239	0.415	0.034	0.029	0.020	0.015
1:8	0.079	0.167	0.242	0.418	0.037	0.031	0.023	0.018
1:10	**0.08**	**0.171**	0.245	0.42	0.041	**0.032**	**0.025**	**0.0186**
1:15	0.08	0.169	**0.246**	**0.421**	**0.0413**	0.031	0.025	0.0185

Table 4: Experiments by varying the amount of negative samples under the setting 2.

+:-	Hit Rate				Precision			
	@3	@10	@20	@50	@3	@10	@20	@50
1:1	0.154	0.258	0.357	0.509	0.058	0.0415	0.0296	0.0212
1:2	0.158	0.263	0.361	0.517	0.063	0.0423	0.0301	0.0219
1:5	0.162	0.269	0.365	0.525	0.0708	0.0436	0.0309	0.0223
1:8	0.163	0.273	0.37	0.529	0.0715	0.0448	0.0312	0.0227
1:10	**0.168**	0.278	**0.373**	**0.536**	**0.0729**	0.0456	0.0315	**0.0235**
1:15	0.167	**0.280**	0.372	0.535	0.0728	**0.0457**	**0.0317**	0.0235

Table 5: Experiments by varying the amount of negative samples under the setting 3.

+:-	Hit Rate				Precision			
	@3	@10	@20	@50	@3	@10	@20	@50
1:1	0.16	0.265	0.363	0.516	0.063	0.043	0.0303	0.0221
1:2	0.163	0.269	0.371	0.528	0.067	0.0438	0.0307	0.0227
1:5	0.167	0.274	0.375	0.533	0.0714	0.0446	0.0314	0.0231
1:8	0.169	0.281	0.378	0.536	0.0725	0.0454	0.0317	0.0235
1:10	**0.17**	0.286	**0.38**	0.54	0.074	**0.046**	**0.032**	**0.024**
1:15	0.17	**0.287**	0.379	**0.542**	**0.0742**	0.0459	0.032	0.024

Table 6: Ablation Study.

Methods	Hit Rate				Precision			
	@3	@10	@20	@50	@3	@10	@20	@50
FitNET	**0.17**	**0.286**	**0.38**	**0.5403**	**0.074**	**0.046**	**0.032**	**0.024**
Var_1	0.162	0.269	0.368	0.5312	0.0713	0.0441	0.031	0.023
Var_2	0.164	0.267	0.372	0.5346	0.0718	0.044	0.0312	0.0236
Var_3	0.168	0.273	0.376	0.535	0.0723	0.0452	0.0317	0.0238

Table 7: Comparison of different methods on the offline dataset.

Methods	Hit Rate				Precision			
	@3	@10	@20	@50	@3	@10	@20	@50
FitNET	**0.17**	**0.286**	**0.38**	**0.54**	**0.074**	**0.046**	**0.032**	**0.024**
$FitNET^-$	0.166	0.282	0.376	0.537	0.0732	0.0454	0.0313	0.0235
MIND	0.156	0.271	0.365	0.528	0.062	0.042	0.031	0.023
YouTube DNN	0.13	0.25	0.34	0.48	0.051	0.039	0.028	0.021
DeepWalk	0.09	0.2	0.29	0.43	0.047	0.035	0.023	0.016
OrderDest2i	0.045	0.14	0.19	0.35	0.033	0.024	0.016	0.009

click-through rate (*CTR* for short), a widely-used industrial metric [17, 19, 32], is employed to evaluate the performance of recommendation methods for serving online traffic. CTR is calculated by taking the total number of times items are clicked on and dividing it by the measured item impressions, as shown in Equation 13.

$$CTR = \frac{\text{\# of clicks}}{\text{\# of impressions}} \tag{13}$$

The experimental results are shown in Figure 5. It is clear that FitNET outperforms other recommendation methods, which indicates that FitNET generates better representations for users by taking their itinerary information into account. Besides, we make several observations from the figure. First, OrderDest2i exhibits the worst performance as it recommends items to users simply based on the information of destination city. This once again confirm our argument that a user's travel intention or preference can rarely be captured by the information in a single order. Second, MIND which models diverse interests of users achieves competitively CTR, compared with other baseline methods. This indicates that the acquirement of users' diverse needs is vital for enhancing the personalization of the matching phase. As to FitNET, it not only catches the diversity of users' needs by proposing an attentive learning mechanism based on a user's behavior sequence (see Section 3.4.3 for details), but also makes use of the important itinerary information of users to elevate the quality of personalized matching results. We are delighted to observe that the average CTR of FitNET is 1.3% higher than that of the next best method, i.e., MIND, which is an apparent progress made by a industrial recommendation system.
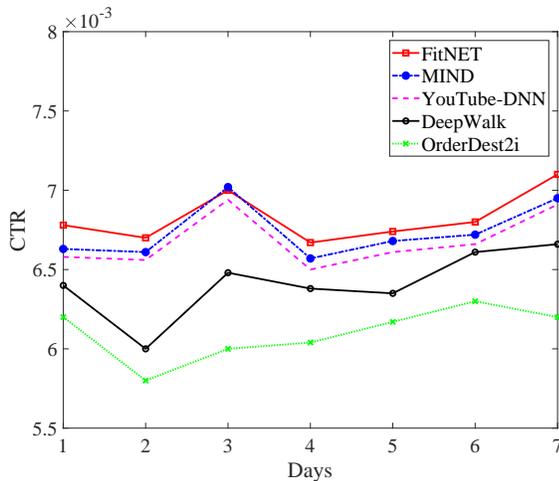

Figure 5: Online CTRs of different methods in seven days in May 2020.

5　CONCLUSION AND FUTURE WORK

Fliggy, one of the most popular online travel platforms (OTPs) in China, has urgent need in helping users find travel-related items they might need or prefer from a item pool with millions of cardinality. Although the popular deep matching networks which retrieve personalized items for users based on deep neural networks

achieve many successes in recent years, tremendous stress still has been put on the recommendation systems of OTPs, in terms of sparsity, diversity and implicitness. To address these challenges, we present a novel Fliggy itinerary-aware deep matching network (FitNET) in this paper, which focuses on improving the accuracy of retrieved items for a user in the matching phase. In particular, to cope with the implicitness problem, we introduce the concept of *itinerary* that contains all unconsumed orders of a user, and uses all orders in a user's itinerary rather than a single order to improve the prediction accuracy to the user's travel intention. Then, to address the diversity and cold start problems, users' profiles are incorporated into the FitNET model. Meanwhile, a group of itinerary-aware attention mechanisms are proposed in FitNET, which are very helpful not only in inferring users' travel intentions or preferences, but also in distinguishing users' diverse needs. Offline experiments are conducted to verify the effectiveness of FitNET in improving recommending accuracy. Online CTRs are also reported to demonstrate the effectiveness and feasibility of FitNET in Fliggy's live production environment. Experimental results show that FitNET achieves superior performance than state-of-the-art methods for recommendation. It is the consideration of users' itinerary information that enables it retrieves more accurate items for each user.

For future work, we will pursue two directions. The first is to incorporate users' comments towards items at Fliggy into our model, so as to achieve a better user representations. The second direction is to upgrade the proposed model by handling more complex scenarios, in which a user may own more than two itineraries.

ACKNOWLEDGMENTS

REFERENCES

[1] Mohamed Hussein Abdi, George Onyango Okeyo, and Ronald Waweru Mwangi. 2018. Matrix Factorization Techniques for Context-Aware Collaborative Filtering Recommender Systems: A Survey. *Comput. Inf. Sci.* 11, 2 (2018), 1–10. https://doi.org/10.5539/cis.v11n2p1

[2] Mutasem Khalil Alsmadi, Khairuddin Bin Omar, Shahrul Azman Noah, and Ibrahim Almarashdah. 2009. Performance Comparison of Multi-layer Perceptron (Back Propagation, Delta Rule and Perceptron) algorithms in Neural Networks. In *IEEE International Advance Computing Conference*.

[3] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2014. Neural Machine Translation by Jointly Learning to Align and Translate. *Computer ence* (2014).

[4] Alex Beutel, Kenton Murray, Christos Faloutsos, and Alexander J. Smola. 2014. CoBaFi: collaborative bayesian filtering. In *23rd International World Wide Web Conference, WWW '14, Seoul, Republic of Korea, April 7-11, 2014*, Chin-Wan Chung, Andrei Z. Broder, Kyuseok Shim, and Torsten Suel (Eds.). ACM, 97–108. https://doi.org/10.1145/2566486.2568040

[5] Tianqi Chen, Weinan Zhang, Qiuxia Lu, Kailong Chen, Zhao Zheng, and Yong Yu. 2012. SVDFeature: a toolkit for feature-based collaborative filtering. *J. Mach. Learn. Res.* 13 (2012), 3619–3622. http://dl.acm.org/citation.cfm?id=2503357

[6] Paul Covington, Jay Adams, and Emre Sargin. 2016. Deep Neural Networks for YouTube Recommendations. In *Proceedings of the 10th ACM Conference*

on Recommender Systems, Boston, MA, USA, September 15-19, 2016. 191–198. https://doi.org/10.1145/2959100.2959190

[7] Mihajlo Grbovic and Haibin Cheng. 2018. Real-time Personalization using Embeddings for Search Ranking at Airbnb. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19-23, 2018*, Yike Guo and Faisal Farooq (Eds.). ACM, 311–320. https://doi.org/10.1145/3219819.3219885

[8] Xiangnan He, Lizi Liao, Hanwang Zhang, Liqiang Nie, Xia Hu, and Tat-Seng Chua. 2017. Neural Collaborative Filtering. In *Proceedings of the 26th International Conference on World Wide Web* (Perth, Australia) *(WWW '17)*. International World Wide Web Conferences Steering Committee, Republic and Canton of Geneva, CHE, 173–182. https://doi.org/10.1145/3038912.3052569

[9] GE Hinton and RR Salakhutdinov. 2006. Reducing the dimensionality of data with neural networks. *Science* 313, 5786 (July 2006), 504–507. https://doi.org/10.1126/science.1127647

[10] Gao Huang, Zhuang Liu, Laurens van der Maaten, and Kilian Q. Weinberger. 2017. Densely Connected Convolutional Networks. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.

[11] Jinhong Huang, Yang Li, Shan Sun, Bufeng Zhang, and Jin Huang. 2020. Personalized Flight Itinerary Ranking at Fliggy. In *Proceedings of the 29th ACM International Conference on Information and Knowledge Management, CIKM 2020, Virtual Event, Ireland, October 19-23, 2019*. 2615–2623. https://doi.org/10.1145/3340531.3412735

[12] Jui-Ting Huang, Ashish Sharma, Shuying Sun, Li Xia, David Zhang, Philip Pronin, Janani Padmanabhan, Giuseppe Ottaviano, and Linjun Yang. 2020. Embedding-based Retrieval in Facebook Search. In *KDD '20: The 26th ACM SIGKDD Conference on Knowledge Discovery and Data Mining, Virtual Event, CA, USA, August 23-27, 2020*, Rajesh Gupta, Yan Liu, Jiliang Tang, and B. Aditya Prakash (Eds.). ACM, 2553–2561. https://dl.acm.org/doi/10.1145/3394486.3403305

[13] Santosh Kabbur, Xia Ning, and George Karypis. 2013. FISM: factored item similarity models for top-N recommender systems. In *The 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD 2013, Chicago, IL, USA, August 11-14, 2013*, Inderjit S. Dhillon, Yehuda Koren, Rayid Ghani, Ted E. Senator, Paul Bradley, Rajesh Parekh, Jingrui He, Robert L. Grossman, and Ramasamy Uthurusamy (Eds.). ACM, 659–667. https://doi.org/10.1145/2487575.2487589

[14] Wang-Cheng Kang and Julian J. McAuley. 2018. Self-Attentive Sequential Recommendation. In *IEEE International Conference on Data Mining, ICDM 2018, Singapore, November 17-20, 2018*. IEEE Computer Society, 197–206. https://doi.org/10.1109/ICDM.2018.00035

[15] Yehuda Koren. 2008. Factorization meets the neighborhood: a multifaceted collaborative filtering model. In *Proceedings of the 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Las Vegas, Nevada, USA, August 24-27, 2008*, Ying Li, Bing Liu, and Sunita Sarawagi (Eds.). ACM, 426–434. https://doi.org/10.1145/1401890.1401944

[16] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E. Hinton. 2012. ImageNet Classification with Deep Convolutional Neural Networks. In *Advances in Neural Information Processing Systems 25: 26th Annual Conference on Neural Information Processing Systems 2012. Proceedings of a meeting held December 3-6, 2012, Lake Tahoe, Nevada, United States*, Peter L. Bartlett, Fernando C. N. Pereira, Christopher J. C. Burges, Léon Bottou, and Kilian Q. Weinberger (Eds.). 1106–1114. http://papers.nips.cc/paper/4824-imagenet-classification-with-deep-convolutional-neural-networks

[17] Chao Li, Zhiyuan Liu, Mengmeng Wu, Yuchi Xu, Huan Zhao, Pipei Huang, Guoliang Kang, Qiwei Chen, Wei Li, and Dik Lun Lee. 2019. Multi-Interest Network with Dynamic Routing for Recommendation at Tmall. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3-7, 2019*, Wenwu Zhu, Dacheng Tao, Xueqi Cheng, Peng Cui, Elke A. Rundensteiner, David Carmel, Qi He, and Jeffrey Xu Yu (Eds.). ACM, 2615–2623. https://doi.org/10.1145/3357384.3357814

[18] Greg Linden, Brent Smith, and Jeremy York. 2003. Amazon.com Recommendations: Item-to-Item Collaborative Filtering. *IEEE Internet Comput.* 7, 1 (2003), 76–80. https://doi.org/10.1109/MIC.2003.1167344

[19] Fuyu Lv, Taiwei Jin, Changlong Yu, Fei Sun, Quan Lin, Keping Yang, and Wilfred Ng. 2019. SDM: Sequential Deep Matching Model for Online Large-scale Recommender System. In *Proceedings of the 28th ACM International Conference on Information and Knowledge Management, CIKM 2019, Beijing, China, November 3-7, 2019*, Wenwu Zhu, Dacheng Tao, Xueqi Cheng, Peng Cui, Elke A. Rundensteiner, David Carmel, Qi He, and Jeffrey Xu Yu (Eds.). ACM, 2635–2643. https://doi.org/10.1145/3357384.3357818

[20] Zequn Lyu, Yu Dong, Chengfu Huo, and Weijun Ren. 2020. Deep Match to Rank Model for Personalized Click-Through Rate Prediction. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. AAAI Press, 156–163. https://aaai.org/ojs/index.php/AAAI/article/view/5346

[21] Zequn Lyu, Yu Dong, Chengfu Huo, and Weijun Ren. 2020. Deep Match to Rank Model for Personalized Click-Through Rate Prediction. In *The Thirty-Fourth AAAI Conference on Artificial Intelligence, AAAI 2020, The Thirty-Second Innovative Applications of Artificial Intelligence Conference, IAAI 2020, The Tenth AAAI Symposium on Educational Advances in Artificial Intelligence, EAAI 2020, New York, NY, USA, February 7-12, 2020*. AAAI Press, 156–163. https://aaai.org/ojs/index.php/AAAI/article/view/5346

[22] H. Brendan McMahan, Gary Holt, D. Sculley, Michael Young, Dietmar Ebner, Julian Grady, Lan Nie, Todd Phillips, Eugene Davydov, Daniel Golovin, Sharat Chikkerur, Dan Liu, Martin Wattenberg, Arnar Mar Hrafnkelsson, Tom Boulos, and Jeremy Kubica. 2013. Ad Click Prediction: A View from the Trenches. In *Proceedings of the 19th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Chicago, Illinois, USA) *(KDD '13)*. Association for Computing Machinery, New York, NY, USA, 1222–1230. https://doi.org/10.1145/2487575.2488200

[23] Tomas Mikolov, Ilya Sutskever, Kai Chen, Gregory S. Corrado, and Jeffrey Dean. 2013. Distributed Representations of Words and Phrases and their Compositionality. In *Advances in Neural Information Processing Systems 26: 27th Annual Conference on Neural Information Processing Systems 2013. Proceedings of a meeting held December 5-8, 2013, Lake Tahoe, Nevada, United States*. 3111–3119. http://papers.nips.cc/paper/5021-distributed-representations-of-words-and-phrases-and-their-co

[24] Bryan Perozzi, Rami Al-Rfou, and Steven Skiena. 2014. DeepWalk: online learning of social representations. In *The 20th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '14, New York, NY, USA - August 24 - 27, 2014*, Sofus A. Macskassy, Claudia Perlich, Jure Leskovec, Wei Wang, and Rayid Ghani (Eds.). ACM, 701–710. https://doi.org/10.1145/2623330.2623732

[25] Suvash Sedhain, Aditya Krishna Menon, Scott Sanner, and Lexing Xie. 2015. AutoRec: Autoencoders Meet Collaborative Filtering. In *Proceedings of the 24th International Conference on World Wide Web Companion, WWW 2015, Florence, Italy, May 18-22, 2015 - Companion Volume*, Aldo Gangemi, Stefano Leonardi, and Alessandro Panconesi (Eds.). ACM, 111–112. https://doi.org/10.1145/2740908.2742726

[26] Ying Shan, T. Ryan Hoens, Jian Jiao, Haijing Wang, and J. C. Mao. 2016. Deep Crossing: Web-Scale Modeling without Manually Crafted Combinatorial Features. In *the 22nd ACM SIGKDD International Conference*.

[27] Karen Simonyan and Andrew Zisserman. 2015. Very Deep Convolutional Networks for Large-Scale Image Recognition. In *3rd International Conference on Learning Representations, ICLR 2015, San Diego, CA, USA, May 7-9, 2015, Conference Track Proceedings*, Yoshua Bengio and Yann LeCun (Eds.). http://arxiv.org/abs/1409.1556

[28] Xiaoyuan Su and Taghi M. Khoshgoftaar. 2009. A Survey of Collaborative Filtering Techniques. *Adv. Artif. Intell.* 2009 (2009), 421425:1–421425:19. https://doi.org/10.1155/2009/421425

[29] Jianshan Sun, Gang Wang, Xusen Cheng, and Yelin Fu. 2015. Mining affective text to improve social media item recommendation. *Inf. Process. Manag.* 51, 4 (2015), 444–457. https://doi.org/10.1016/j.ipm.2014.09.002

[30] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, undefinedukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. In *Proceedings of the 31st International Conference on Neural Information Processing Systems* (Long Beach, California, USA) *(NIPS'17)*. Curran Associates Inc., Red Hook, NY, USA, 6000–6010.

[31] Hao Wang, Naiyan Wang, and Dit-Yan Yeung. 2015. Collaborative Deep Learning for Recommender Systems. In *Proceedings of the 21th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, Sydney, NSW, Australia, August 10-13, 2015*, Longbing Cao, Chengqi Zhang, Thorsten Joachims, Geoffrey I. Webb, Dragos D. Margineantu, and Graham Williams (Eds.). ACM, 1235–1244. https://doi.org/10.1145/2783258.2783273

[32] Jizhe Wang, Pipei Huang, Huan Zhao, Zhibo Zhang, Binqiang Zhao, and Dik Lun Lee. 2018. Billion-scale Commodity Embedding for E-commerce Recommendation in Alibaba. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19-23, 2018*, Yike Guo and Faisal Farooq (Eds.). ACM, 839–848. https://doi.org/10.1145/3219819.3219869

[33] Shuai Zhang, Yi Tay, Lina Yao, Aixin Sun, and Jake An. 2019. Next item recommendation with self-attentive metric learning. In *Thirty-Third AAAI Conference on Artificial Intelligence*, Vol. 9.

[34] Yin Zhang, Haokai Lu, Wei Niu, and James Caverlee. 2018. Quality-aware neural complementary item recommendation. In *Proceedings of the 12th ACM Conference on Recommender Systems, RecSys 2018, Vancouver, BC, Canada, October 2-7, 2018*, Sole Pera, Michael D. Ekstrand, Xavier Amatriain, and John O'Donovan (Eds.). ACM, 77–85. https://doi.org/10.1145/3240323.3240368

[35] Guorui Zhou, Chengru Song, Xiaoqiang Zhu, Ying Fan, Han Zhu, Xiao Ma, Yanghui Yan, Junqi Jin, Han Li, and Kun Gai. 2017. Deep Interest Network for Click-Through Rate Prediction. (2017).

[36] 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 *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD 2018, London, UK, August 19-23, 2018*, Yike Guo and Faisal

Farooq (Eds.). ACM, 1079–1088. https://doi.org/10.1145/3219819.3219826