

Trip Prediction by Leveraging Trip Histories from Neighboring Users

Yuxin Chen¹ and Morteza Haghir Chehreghani²

Abstract—We propose a novel approach for trip prediction by analyzing user’s trip histories. We augment users’ (self-) trip histories by adding “similar” trips from other users, which could be informative and useful for predicting future trips for a given user. This also helps to cope with noisy or sparse trip histories, where the self-history by itself does not provide a reliable prediction of future trips. We show empirical evidence that by enriching the users’ trip histories with additional trips, one can improve the prediction error by 15%~40%, evaluated on multiple subsets of the Nancy2012 dataset. This real-world dataset is collected from public transportation ticket validations in the city of Nancy, France. Our prediction tool is a central component of a trip simulator system designed to analyze the functionality of public transportation in the city of Nancy.

I. INTRODUCTION AND BACKGROUND

This paper is concerned with prediction of future trips according to trip histories. Trip prediction is used, for example, to simulate a public transportation system, to analyze the traffic, to investigate the demand and load, and to identify the bottlenecks and the constraints. In principle, a universal trip planner can be used to predict a trip for a user, by sampling from the predictive distribution of future trips, marginalized over the entire population. Such trips usually reflect the general traveling pattern of all the users in the system, and are useful for estimating the overall statistics of a transportation system. However, passengers do not always travel according to what a trip planner offers to them, due to the fact that they have different trip preferences and therefore behave differently according to various criteria. Thereby, looking at individual historical trips can provide very useful information about trip habits and behavior of them, and hence is helpful for making personalized recommendations.

a) Personalized trip prediction: A partially related task to trip prediction is trip recommendation, where a recommendation system deals with the *choice of the routes* to be taken by the user. A trip predictor, on the other hand, aims at *predicting the trip* that user u will take at time t (i.e. to estimate the origin and the destination). Training a trip prediction model at individual level is essential for building trip planner systems. This is due to the fact that personalized trip planner [25], [30], [31], [12], [13], [4] recommends future trips for a given user, by sampling from the predictive distribution. The sampling is conditioning on the query and the historical behaviors of the user (i.e., $\text{target} \sim \text{Pr}(\text{trip} \mid \text{query}, \text{history})$). The query (e.g., “what’s the most popular

trip on Monday afternoon”, “best route to work in two hours”) is universal for all users, whereas the history usually consists of user’s profiles [28], real-time trip information (e.g., traffic [25] and real-time location [31]), community-contributed meta-data (e.g., geo-tags [23] and photos [26], etc), previously realized trips (e.g., trips realized in a different city [14], [26]) and so on, which encodes user’s bias / preference towards specific trips. Such information is crucial in steering the trip recommender towards making suggestions that are meaningful to the user, as well as predicting future trips. For example, students who travel to schools often take trips with the shortest travel time; while parents who pick up their children at school may very likely choose a different route, so that they can drop by local shops and do grocery shopping on the way. If we observe a direct trip from location A (e.g., home) to B (e.g., school) frequently enough in the travel history of a given user, it is very likely for the same user to take this trip under similar conditions, as we may guess she is a user of some fixed type (e.g., “students”). With sufficient amount of historical data from each user, learning their attributes (i.e., latent variables) [10], [4] from trip histories is helpful for us to train a good trip predictor. In this way, it helps to group users of common interest, which in turn regularizes the model to make meaningful predictions [17].

One particular challenge for trip prediction is the “cold-start” problem: we may want to predict the behavior of very recent or newly-joined users, who do not yet have a long travel history in the system. It is challenging to represent such users (i.e., finding expressive representation in vector space) purely based on histories of their own, since there is not sufficient knowledge about them. As an example, the user in second line of Figure 1 has a rather short trip history, and the single trip it takes on week 9 only contains limited information about the user’s behavior. In such scenarios, one often needs to trade exploration (i.e., collecting more information about the user) and exploitation (i.e., making prediction from the realized trips).¹ However, in this work we focus on *exploitation* of information which is already known and present, and we target at maximizing the relevance of the current predictions. We hope to develop an intelligent trip prediction system that is capable of decoding users’ travel patterns from historical travel data, even when the trip history is not statistically rich enough.

b) Similarity-based approach: A natural way to address such “cold-start” problem is to enrich users’ trip

¹Yuxin Chen is with Department of Computer Science, University of Chicago, USA cheniyuxin@uchicago.edu

²Morteza Haghir Chehreghani is with the Department of Computer Science and Engineering, Chalmers University of Technology, Sweden morteza.chehreghani@chalmers.se

¹One common setting to study this problem is online learning, where one seeks to iteratively predict trips for a given user and explore new knowledge, aiming to maximize the cumulative relevance of the entire set [1], [2], [3].




Fig. 1: Illustration of trips taken by users at some fixed time slots. We can see the histories of trips taken by three different users, one for each line, on the same time slots (say, from 8am to 9am on Mondays) over 9 weeks. The line segments inside the rectangles represent trajectories of the trips; the turning points represent public transport transit stops. A trip predictor aims to predict the trips for each of the users at some given time slot, e.g., from 8am to 9am, on Monday of week 10.

history by looking into historical trips of “similar” users. Similarity-based approach has been popular as an alternative to traditional feature-based learning for numerous machine learning tasks, especially when it is challenging to find satisfactory (vectorial representation of) features for learning purposes [32]. A commonly used approach to similarity-based classification is k -nearest-neighbors (k-NN). Nearest-neighbor learning is the algorithmic analogue of the exemplar model of human learning [18]. Although simple, it is often effective in practice, and empirical studies on a number of benchmark data sets show that it is difficult to surpass the performance of k-NN for similarity data [29].

We build upon and extend the nearest-neighbor paradigm to personalized trip prediction problem. One fundamental problem we face is to develop a proper similarity measurement between different users. Optimizing the similarity measure has been extensively studied in both supervised learning and unsupervised learning setting. Under the supervised learning setting, where target trips are known, one can optimize the similarity measure via metric learning [33]. In unsupervised learning setting, one can regularize the models via, e.g., non-negative matrix factorization [24], [19]. In both settings, however, the goal is to learn a similarity / distance measurement that leads to *optimal clustering*, rather than finding the most useful users which incur the minimal prediction error. In contrast, we take a direct approach, where we directly relate the similarity between two users with the performance w.r.t. the prediction error. We argue that our work is orthogonal to existing methods that learn feature representations, and we can always learn better features via these methods as pre-processing steps for our approach.

One work is relevant to [22], where the goal is to find clusters of users to minimize the prediction error. In [22] the authors make a direct connection between user groups and prediction error: a user is assigned to a group as long as it helps reducing prediction error, and the micro-segmentation of user groups are used *collectively* for targeted and personalized prediction. Despite of that, we do not require partitioning the users into groups; instead, we allow

the similarity measure to be asymmetric, where *neighbors* (i.e., similar users) of a given user do not necessarily have that user as a neighbor. Such relaxation provides us with flexibility of creating arbitrary groups that could help with reducing the expected prediction error.

c) Our contribution: To sum up, our contributions include the following aspects:

1. We develop a general framework for trip prediction, which predicts future trips for given users by incorporating their historical trip choices. Unlike existing works [25], [31], [23], [26] which exploit only the user-specific information, we focus on developing a framework on top of such systems, which aims at improving the performance (i.e., prediction error) of such systems by leveraging historical trips of *similar users*, where our similarity measurement by definition is related to the reduction in prediction error in a separate validation set.
2. We demonstrate the effectiveness of our method on multiple subsets of the Nancy2012 public transportation dataset, and show that by enriching the users’ trip histories with additional trips, one can improve the performance of trip predictor by 15%~40% in terms of prediction error. An important aspect of our study is that the dataset is real-world, i.e., it is collected from real (electronic) trip transactions of the passengers in the city of Nancy, France, during the year 2012. By predicting the trips at an individual level, our prediction tool constitutes an important component of a trip simulator system, which is designed and implemented to analyze the functionality of public transportation and the behavior of users in the city of Nancy. For instance, the trip simulator might query the future trips of all users in the system at a specific time.² This leads to computing sufficient statistics of the overall traffic of the city, as well as zone specific traffic, by aggregating individual predicted trips. Such information can be used, for example, to detect the bottleneck of the system, decide whether bus stops should

²This information can be obtained at different time points which can help to for example understand the dynamics of traffic and the functionality of transportation system [5], [27].

be added / placed, and support strategic decision making (e.g., to build new transportation system).

3. We discuss the possibility that further improvements in performance can be attained by (1) leveraging new, informative and fine-grade features and (2) transferring existing features into more robust representations via e.g. non-negative matrix factorization.

II. NEIGHBOR-BASED TRIP PREDICTION

Our goal is to develop an effective approach for trip prediction based on trip histories. Formally, we define a trip to be a tuple (o, d, v) , where o denotes the (location of the) origin, d denotes the destination, and v denotes the list of transit stops (or via-points). Similarly, we use $(o, d, v)_{u,t}$ to denote a trip taken by user u at t , where t denotes the time of the trip³. Individual users' trip histories might be sparse or noisy, thus, they might not be sufficient to provide a suitable feature representation for predicting future trips. Therefore, we need to augment the individual users' self trip histories by trip histories of other users in order to compute a more robust estimation. However, taking all other user trip histories into account, i.e. averaging over all trips, is not appropriate, since different people might have different trip preferences and thus global averaging discards such a diversity. Therefore, for each user, we need to identify an additional set of trips histories (i.e. neighbors) than the self history which help to improve future trip prediction. To do so, we take two important considerations into account:

I) The users usually make a diverse set of trips during a day. Therefore, it makes sense to divide a day into small (e.g., one-hour) time intervals and consider the trips inside this interval as unit of trip behavior. On the other hand, the trip behavior of user u at time t might be similar to the trip behavior of user u' at a different time t' such that t and t' does not necessary overlap. For example, user u might travel to the city university at time 9am, whereas user u' might take this trip at time 3:00 pm. Therefore, when querying a trip as well as finding appropriate auxiliary trip histories, we parametrize the operations by time point t . Formally, we define the base *entities* that we work with as user-time pairs, i.e., $e := \langle u, t \rangle$. We use T_{ut} to refer to the set of trips associated entity $\langle u, t \rangle$:

$$T_{ut} = \{(o, d, v)_{u',t'} : u' = u \wedge t' = t\}$$

Then, the question is: *for a specific entity $\langle u, t \rangle$ which represents user u at time t , what are the other entities that can be used to obtain a better prediction for the next trip of the user, at time t in future days/weeks?*

II) The usefulness relations are not symmetric, i.e. entity $\langle u', t' \rangle$ may be helpful for entity $\langle u, t \rangle$ to find a better trip in future, however, $\langle u', t' \rangle$ might not need $\langle u, t \rangle$ for this

³For public transportation, we use t to denote the time in the week, e.g., 8am-9am on Monday, since the trips are often repetitive over different weeks. We use o, d to denote locations of the origin and destination stops; this allows us to compute distances between two trips efficiently, and hence we don't need to further discretize the space (e.g., using close-by bus stops or city areas) as is done in [27] (such quasi-discretization may lead to less accurate prediction in addition to extra computation time).




Fig. 2: Computing neighboring entities of $\langle u, t \rangle$: we divide the dataset into training and validation set, each containing a fraction of trips for each entity.

purpose. In particular, such a directional relation can hold whenever the trip history of $\langle u', t' \rangle$ is clean and long enough, but $\langle u, t \rangle$ is very short or noisy. Thus, the methods that work based on grouping or clustering of entities discard this kind of asymmetric relations.

Thereby, we propose a method to compute additional helpful entities to each specific entity. However, we do not have access to the user profile, such as meta-data description of the users' occupation, hobbies, and age group, etc. Instead, we are only given the users' trip histories, based on which we should define a proper time-dependent distance/similarity measure between them. In absence of meta-data feature representation, we follow a similarity-based approach, while relying on the fundamental principle of learning theory: *a good model should perform well on unseen data from the same source*.

For a given entity $\langle u, t \rangle$, we determine its neighbors in a non-parametric way using a separate unseen dataset called *validation set*. More formally, let us start with the simple case, where we are given a dataset of trip entities with equal trip histories L , i.e., $\mathbf{D} = \{\langle u, t \rangle : |T_{u,t}| = L\}$. We then divide the whole dataset into two subsets, training set and validation set, each containing $L/2$ trips per entity. Then, we use the validation set to identify the appropriate neighbors of the entities. For each $\langle u, t \rangle$ in the training set (resp. validation set), we denote the associated set of trips by T_{ut}^{trn} (resp. T_{ut}^{vld}). See Figure 2 for illustration. In case of not availability of enough annotated data, active learning methods can be employed to provide annotated data in a cost-efficient way [15], [21].

To compute the appropriate neighbors of the entity $\langle u, t \rangle$, we investigate which of the training histories are at least equally similar to the validation history compared with the self training history, i.e.,

$$\mathcal{N}_{ut} = \{\langle u', t' \rangle : \text{dist}(T_{u't'}^{trn}, T_{ut}^{vld}) \leq \text{dist}(T_{ut}^{trn}, T_{ut}^{vld})\} \quad (1)$$

where $\text{dist}(\cdot, \cdot)$ denotes the distance between trip histories. In our study, we investigate two options for $\text{dist}(\cdot, \cdot)$:

1) *ordered*, where only the trips at the same positions are compared, i.e.,

$$\text{dist}(p, q) = \frac{2}{L} \sum_{1 \leq i \leq L/2} \text{seuc}(p_i, q_i), \quad (2)$$

Algorithm 1 History-based trip prediction.

Require: The entities and the respective trips.

Ensure: Predicted trip(s) to each entity.

- 1: **for each** entity $\langle u, t \rangle$ **do**
 - 2: Split the trip histories into T_{ut}^{trn} and T_{ut}^{vld} for construction of the training and validation sets.
 - 3: **end for**
 - 4: **for each** entity $\langle u, t \rangle$ **do**
 - 5: $\mathcal{N}_{ut} = \{ \langle u', t' \rangle : \text{dist}(T_{u't'}^{trn}, T_{ut}^{vld}) \leq \text{dist}(T_{ut}^{trn}, T_{ut}^{vld}) \}$.
 - 6: $r_{ut} \in \arg \max_{x \in T(\mathcal{N}_{ut})} \sum_{y \in T(\mathcal{N}_{ut})} f_x \text{sim}(x, y)$.
 - 7: **end for**
 - 8: **return** $\{r_{ut}\}$
-

Hereby, we sort the trips in the trip histories according to their time of realization, and p_i (resp. q_i) indicates the i^{th} trip in trip history p (resp. q). Further, $\text{seuc}(p_i, q_i)$ gives the squared Euclidean distance between trips p_i and q_i . Specifically, for two single-leg trips $p_i := (o_1, d_1, v)$ and $q_i := (o_2, d_2, v)$ where $v = \emptyset$, we have $\text{seuc}(p_i, q_i) = \text{seuc}(\langle o_1, d_1, v \rangle, \langle o_2, d_2, v \rangle) = |o_1 - o_2|^2 + |d_1 - d_2|^2$. Note that this variant requires p and q to include the same number of trips.

- 2) *all2all*: where each trip from one trip history is compared against all trips of the other history, i.e.,

$$\text{dist}(p, q) = \frac{4}{L^2} \sum_{1 \leq i \leq L/2} \sum_{1 \leq j \leq L/2} \text{seuc}(p_i, q_j). \quad (3)$$

One advantage of *all2all* over the *ordered* variant is that p and q do not necessarily need to have the same number of trips. Thus, *all2all* is more general-purpose and can be applied to trip histories with different length.

We note that our framework is general enough to use other distances measures than Euclidean distances as well, e.g. the Minimax distances [6], [7]. In the next step, we employ the members of the neighbor set \mathcal{N}_{ut} to predict a future trip for entity $\langle u, t \rangle$. For this purpose, we consider the total trip histories of all neighbors collected in \mathcal{N}_{ut} (i.e. including both training and validation trips) and compute the representative trip(s) as the trip(s) r_{ut} with the minimal average distance (or maximal average similarity) with other trips. More precisely,

$$r_{ut} \in \arg \max_{x \in T(\mathcal{N}_{ut})} \sum_{y \in T(\mathcal{N}_{ut})} f_x \cdot \text{sim}(x, y), \quad (4)$$

where $T(\mathcal{N}_{ut})$ indicates the set of all trips of all entities in \mathcal{N}_{ut} , and f_x denotes the frequency of trip x in this set. $\text{sim}(x, y)$ measures the pairwise similarity between the two trips x and y , which is obtained by $\text{const} - \text{seuc}(x, y)$. We select const as the minimal value for which the pairwise similarities become nonnegative. Finally, r_{ut} is predicted as the next trip of the user under investigation. Note that r_{ut} might not be deterministic (if there are ties among multiple trips). Algorithm 1 describes the whole procedure in detail.

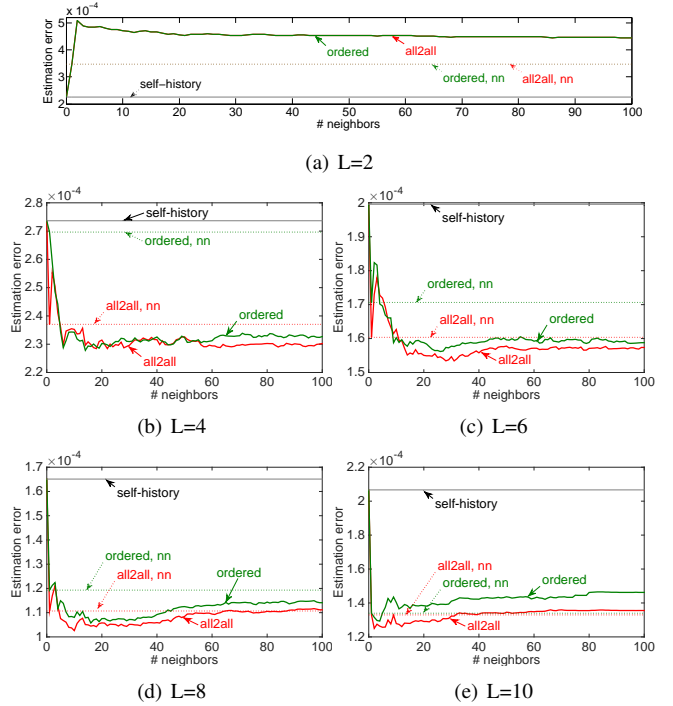


Fig. 3: Estimation error on entities with even L .




Fig. 4: Estimation error of different experimental settings.

III. EXPERIMENTS

In this section, we conduct extensive empirical evaluation of our approach on a real-world public transportation management platform, and demonstrate its superior performance compared to the alternative baselines (i.e., prediction using “singleton” entity features, and prediction using the single nearest neighbor together with self history). In specific, we show that our framework can improve the performance of existing trip prediction algorithms via our similarity-based data refinement process. Moreover, we investigate the impact of transforming the origin and destination features to a new set of features via non-negative matrix factorization.

A. Experimental Setup

a) *Dataset*: We perform our experiments on the real-world trip specification data Nancy2012 collected from the city of Nancy in France [11], [12]. This data is prepared from e-card validation collection. We query trip histories with different lengths, i.e. $L = 2, 3, 4, 5, 6, 7, 8, 9, 10$, to




Fig. 5: Estimation error on entities with different lengths of trip history when L is an odd number (i.e., $L = 3, 5, 7, 9$).

produce different datasets.⁴ For each L we collect 2,000 entities from the database, unless there are less entities for a specific L (For $L = 10$, we could collect only 740 entities). We consider single-leg trips in our evaluations.⁵ Thus, each trip is specified by four elements: the longitude and the latitude of the origin and the longitude and the latitude of the destination.

We split each dataset into training and validation sets. Moreover, we have an additional trip (*test* trip) for each entity which will be used as the ground truth (i.e. T_{ut}^{tst}) in order to investigate the correctness of our estimation/prediction.

b) *Evaluation criteria*: We compare the ground-truth and the predicted trips and compute the mean squared error

$$e\hat{r}r = \frac{1}{|\{\{u, t\}\}|} \sum_{\langle u, t \rangle} \text{seuc}(r_{ut}, T_{ut}^{tst}), \quad (5)$$

where $|\{\{u, t\}\}|$ shows the number of test cases (entities).

B. Numerical Analysis

a) *Results with different trip history lengths*: Figures 3, 4 and 5 illustrate the estimation error of computing the neighbors respectively for $L = 2, 4, 6, 8, 10$ (i.e. when L is an even number) and $L = 3, 5, 7, 9$ (i.e. when L is an odd number). The neighbors are sorted according to their usefulness on validation set (i.e., their distance to the entity of interest, as defined in Equation (2) and Equation (3)). We investigate different number of neighbors per user, where *no. of neighbors* = 0, indicating the use of only self history for computing representative trip and prediction. Thus, this setting constitutes our baseline. Another baseline is to use the single nearest neighbor with self history. In Figure 4(a) and Figure 4(b), we plot the prediction error using self-history only, nearest-neighbor, and the *optimal set* of neighbors, respectively, for the two options of distance function. We observe, i) except for $L = 2$, our approach always leads to reducing the estimation error. For $L = 2$, there is only one trip for training and one for validation. Thus, due to noise and sparsity, we are not able to select informative and reliable neighbors. However, once we increase the number of trips for training and validation sets, e.g. $L = 3, 4, 5, 6, 7, 8, 9, 10$, then, our approach yields computing better neighbors and a

better representative trip among them, which thereby leads to reduce the estimation error by 15%~40%. ii) As we increase the number of trips in history, i.e. the L , then we can better compute the neighbors and obtain a more reliable representative trip. Thus, a larger L yields better performance in trip prediction, as well as the results are smoother. iii) The results are very much consistent between *all2all* and *ordered* variants, which also indicates lack of any significant temporal trip behavior. However, the advantage of the *all2all* variant is that it can be employed even when there are entities with varying number of trips. Thus, the *all2all* variant can replace whenever the *ordered* variant is desired. Moreover, *all2all* in general outperforms *ordered*, if we are allowed to include a fixed number of neighbors for each entity. Such observation suggest that we should consider to use *all2all* whenever possible as the more reliable similarity measurement. We note that in the case of superiority of the *ordered* variant, methods such as recurrent neural networks could be employed to better extract the temporal aspects of driving/transport trajectories [16], [20].

b) *Pre-processing trips via non-negative matrix factorization*: We then investigate that how the use of matrix factorization methods affects the prediction accuracy. In particular, we perform non-negative matrix factorization on the feature matrices, in order to transform the original features into another type of features which might be more suitable. This technique is very common in recommendations and collaborative filtering. We repeat the experiments for different number of hidden components and choose the best results. In our experiments, the optimal number of components is 4. The results are shown in Figure 6 for $L = 6$ and $L = 8$. For the other values of L , we observe very consistent results. We observe that transforming the original features into the new features leads to a significant increase in the prediction error. This observation implies that the original features are sufficient and informative enough to be used for the purpose of learning and prediction. Intuitively, this makes sense, because the original features are orthogonal (non-redundant) and sufficiently describe the origin and destination points.

c) *Enriching short-history entities with longer history neighbors*: In the following, we investigate whether augmenting short histories with long histories can help to predict more accurate trips. In particular, we study the case of $L = 2$, i.e. the only case where our approach fails to improve prediction accuracy. We choose different number of entities

⁴Note that the *ordered* variant requires that the two trip histories (training and validation sets) must be aligned, i.e. they should have the same lengths. Thus, we perform $L = 3, 5, 7, 9$ only for *all2all* variant.

⁵Note that over 95% of the trips in our dataset are single leg.

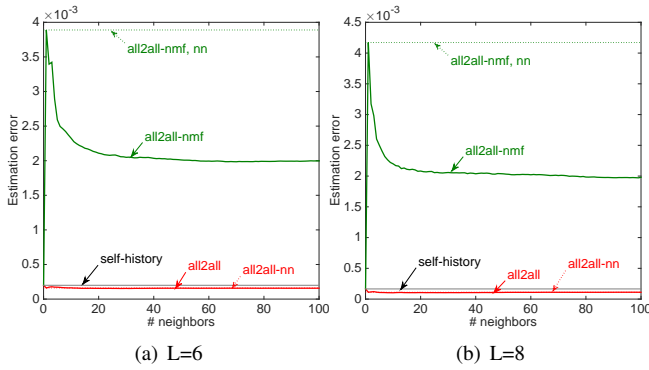


Fig. 6: Estimation error on entities with/without non-negative matrix factorization, for $L = 6, 8$.




Fig. 7: Estimation error when different number of entities with $L = 2$ are augmented with 2000 entities with $L = 8$.

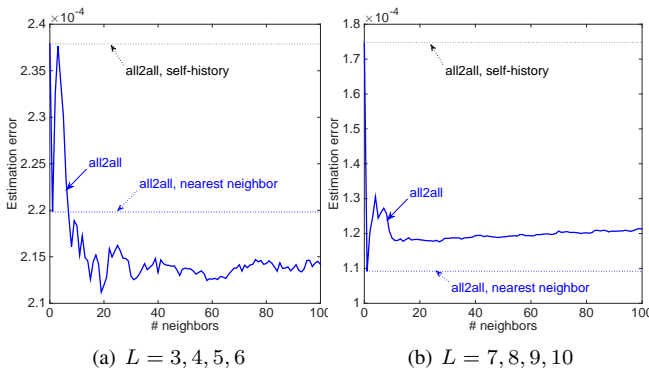


Fig. 8: Estimation error when entities with different number of trips are combined.

with $L = 2$ (e.g. 100, 500, 1000 and 2000) and combine them with 2000 entities whose L is 8. That is, we merge the training set (resp. validation set) of trips with $L = 8$ with training set (resp. validation set) of trips with $L = 2$. Figure 7 illustrates the results. We compute the estimation error only for the entities with $L = 2$. We observe that, i) the impact of very short histories (i.e. $L = 2$) is very crucial, such that even augmenting them by very long histories dose not help much. We still see that using only the self-histories is a better choice for this particular case. ii) As we increase the ratio of the number of long histories to the number of short histories, then we obtain better and more reliable neighbors

such the estimation error decreases. We particularly observe this behavior when the number of entities changes from 2000 to 1000, 500 and finally to 100.

d) *Combining entities with different number of trips:*

Finally, we consider combination of entities with different number of trips, i.e. with varying L . In the first case (Figure 8(a)) we combine entities with $L = 3, 4, 5, 6$ trips, where the dataset contains 500 entities from each category. In the second case (Figure 8(b)), we consider $L = 7, 8, 9, 10$, and collect 500 entities from each category. For this setting, we employ *all2all* measure for computing appropriate neighbors, since the entities have different number of trips. Figure 8 illustrates the results. We observe, i) in both cases, our method helps to compute appropriate neighbors and thereby to reduce the estimation error. ii) For the second case, i.e. when $L = 7, 8, 9, 10$, the estimation error is smaller (and smoother) than the first case. The reason is that in the second case the trip histories are longer, thus the representative trip can be computed in a more robust way.

IV. DISCUSSION AND FUTURE WORK

In this section, we discuss some potential directions on improving the prediction accuracy. One defining factor for performance is the trips' initial feature representations. As previously discussed, the performance of the predictor relies on our definition of the distance function $dist(\cdot, \cdot)$, which currently is defined as a function of the (pairwise-) squared Euclidean distances between trips. However, the geographical information about a trip is more than just the origin and destination stop. As demonstrated in Figure 9, straight line distance between (O, D) pairs hardly reflects the scales of the difference between different trips. In Figure 9(b), Trip B and Trip C represent the same service line in different hours of the day. They are almost identical except for the last stop. Trip A and Trip B (or C) are very different, though they still share a common stop which could be a popular transit stop for 2-leg trips (i.e., some users travel on Trip A may transfer to B (or C) at the intersecting point). To capture such potentially useful information, we propose a new distance measure between trips, $tripd(\cdot, \cdot)$, defined as follows:

$$tripd(p_i, q_i) = \left(1 - \frac{p_i \cap q_i}{p_i \cup q_i}\right) * seuc(p_i, q_i) \quad (6)$$

where the first term on the R.H.S. represents the Jaccard distance between trip p_i and q_i if we view them as sets of intermediate stops. This heuristic captures the intuition that if two trips share many common stops, even though the ending stops are far apart, we may still want to treat them as “somewhat similar” since they can belong to different segments of the same service line, or the two trips can be potential transfer trips for each other.

Since our method for proposing neighbors is orthogonal to the feature extraction/engineering component, we may preprocess the current features by transferring them into more robust, noise-resilient features, via commonly used techniques such as non-negative matrix factorization or truncated SVD. This black-box feature engineering component




Fig. 9: Routes and Stops in Nancy2012.

might be particularly helpful when we use more complicated type of features or a combination of different criteria. Another choice would be minimax distance which computes the transitive relations and elongated patterns in a nonparametric way [9], [8].

V. CONCLUSION

We propose a new method for trip prediction by taking into account the user trip history. We augment users trip history with trips taken by similar users, where the similarity between users are directly guided by the prediction error. We show that by incorporating the augmented trip history, one can improve the informativeness of users' self histories, and hence improve the overall performance of the trip prediction system. We perform experiments on a real-world dataset collected from real trip transactions of the passengers in the city of Nancy in France.

Acknowledgment: The work is mainly done at Xerox Research Centre Europe. We thank for the insightful discussions and the help in working with Nancy2012.

REFERENCES

- [1] Niklas Åkerblom, Yuxin Chen, and Morteza Haghir Chehreghani. An online learning framework for energy-efficient navigation of electric vehicles. In *Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI*, pages 2051–2057, 2020.
- [2] Niklas Åkerblom, Yuxin Chen, and Morteza Haghir Chehreghani. Online learning of energy consumption for navigation of electric vehicles. *CoRR*, abs/2111.02314, 2021.
- [3] Niklas Åkerblom, Fazeleh Sadat Hoseini, and Morteza Haghir Chehreghani. Online learning of network bottlenecks via minimax paths. *Machine Learning*, 2022.
- [4] Theo A Arentze. Adaptive personalized travel information systems: a bayesian method to learn users' personal preferences in multimodal transport networks. *IEEE Transactions on intelligent transportation systems*, 14(4):1957–1966, 2013.
- [5] Oded Cats. Dynamic modelling of transit operations and passenger decisions. 2011.
- [6] Morteza Haghir Chehreghani. K-nearest neighbor search and outlier detection via minimax distances. In *Proceedings of the 2016 SIAM International Conference on Data Mining*, pages 405–413, 2016.
- [7] Morteza Haghir Chehreghani. Classification with minimax distance measures. In *Proceedings of the Thirty-First AAAI Conference on Artificial Intelligence (AAAI)*, pages 1784–1790, 2017.
- [8] Morteza Haghir Chehreghani. Feature-oriented analysis of user profile completion problem. In *39th European Conference on Information Retrieval (ECIR)*, volume 10193, pages 304–316, 2017.
- [9] Morteza Haghir Chehreghani. Unsupervised representation learning with minimax distance measures. *Machine Learning*, 109(11):2063–2097, 2020.
- [10] An-Jung Cheng, Yan-Ying Chen, Yen-Ta Huang, Winston H. Hsu, and Hong-Yuan Mark Liao. Personalized travel recommendation by mining people attributes from community-contributed photos. In *ACM Multimedia*, pages 83–92, 2011.
- [11] Boris Chidlovskii. A method for estimating dynamic origin-destination matrices from the atv data. Technical report, XRCE, no. 2012158, 2012.
- [12] Boris Chidlovskii. Improved trip planning by learning from travelers' choices. In *ICML workshop on Mining Urban Data*, pages 17–26, 2015.
- [13] Caspar G. Chorus, Eric J. E. Molin, and Bert Van Wee. Use and effects of advanced traveller information services (atis): A review of the literature. *Transport Reviews*, 26(2):127–149, 2006.
- [14] Maarten Clements, Pavel Serdyukov, Arjen P de Vries, and Marcel JT Reinders. Personalised travel recommendation based on location co-occurrence. *arXiv preprint arXiv:1106.5213*, 2011.
- [15] Federica Comuni, Christopher Mészáros, Niklas Åkerblom, and Morteza Haghir Chehreghani. Passive and active learning of driver behavior from electric vehicles. In *25th IEEE International Intelligent Transportation Systems Conference, ITSC*. IEEE, 2022.
- [16] Andreas Demetriou, Henrik Allsvåg, Sadegh Rahrovani, and Morteza Haghir Chehreghani. Generation of driving scenario trajectories with generative adversarial networks. In *23rd IEEE International Conference on Intelligent Transportation Systems, ITSC*, pages 1–6. IEEE, 2020.
- [17] Michael D Ekstrand, John T Riedl, and Joseph A Konstan. Collaborative filtering recommender systems. *Foundations and Trends in HCI*, 4(2):81–173, 2011.
- [18] Robert L Goldstone and Alan Kersten. Concepts and categorization. *Handbook of psychology*, 2003.
- [19] Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *Unsupervised learning*. Springer, 2009.
- [20] Fazeleh Sadat Hoseini, Sadegh Rahrovani, and Morteza Haghir Chehreghani. Vehicle motion trajectories clustering via embedding transitive relations. In *24th IEEE International Intelligent Transportation Systems Conference, ITSC*, pages 1314–1321. IEEE, 2021.
- [21] Sanna Jarl, Linus Aronsson, Sadegh Rahrovani, and Morteza Haghir Chehreghani. Active learning of driving scenario trajectories. *Engineering Applications of Artificial Intelligence*, 113, 2022.
- [22] Tianyi Jiang and Alexander Tuzhilin. Improving personalization solutions through optimal segmentation of customer bases. *TKDE*, 21(3):305–320, 2009.
- [23] Takeshi Kurashima, Tomoharu Iwata, Go Irie, and Ko Fujimura. Travel route recommendation using geotags in photo sharing sites. In *CIKM*, 2010.
- [24] Daniel D Lee and H Sebastian Seung. Algorithms for non-negative matrix factorization. In *NIPS*, 2001.
- [25] Julia Letchner, John Krumm, and Eric Horvitz. Trip router with individualized preferences (trip): Incorporating personalization into route planning. 2006.
- [26] Abdul Majid, Ling Chen, Gencai Chen, Hamid Turab Mirza, Ibrar Hussain, and John Woodward. A context-aware personalized travel recommendation system based on geotagged social media data mining. *Int. J. Geogr. Inf. Sci.*, 2013.
- [27] Luís Moreira-Matias, João Gama, Michel Ferreira, João Mendes-Moreira, and Luis Damas. Time-evolving od matrix estimation using high-speed gps data streams. *Expert Systems with Applications*, 44:275–288, 2016.
- [28] Moon-Hee Park, Jin-Hyuk Hong, and Sung-Bae Cho. Location-based recommendation system using bayesian user's preference model in mobile devices. In *Ubiquitous Intelligence and Computing*, pages 1130–1139. Springer, 2007.
- [29] Elzbieta Pekalska, Pavel Paclik, and Robert PW Duin. A generalized kernel approach to dissimilarity-based classification. *JMLR*, 2:175–211, 2002.
- [30] Francesco Ricci, Lior Rokach, and Bracha Shapira. *Introduction to recommender systems handbook*. Springer, 2011.
- [31] Henan Wang, Guoliang Li, Huiqi Hu, Shuo Chen, Bingwen Shen, Hao Wu, Wen-Syan Li, and Kian-Lee Tan. R3: A real-time route recommendation system. *Proc. VLDB Endow.*, 7(13):1549–1552, August 2014.
- [32] D Randall Wilson and Tony R Martinez. Reduction techniques for instance-based learning algorithms. *Machine learning*, 38(3):257–286, 2000.
- [33] Liu Yang. Distance metric learning: A comprehensive survey. 2006.