

Destination Prediction Based on Partial Trajectory Data

Patrick Ebel¹, Ibrahim Emre Göl¹, Christoph Lingenfelder² and Andreas Vogelsang¹

Abstract—Two-thirds of the people who buy a new car prefer to use a substitute instead of the built-in navigation system. However, for many applications, knowledge about a user’s intended destination and route is crucial. For example, suggestions for available parking spots close to the destination can be made or ride-sharing opportunities along the route are facilitated. Our approach predicts probable destinations and routes of a vehicle, based on the most recent partial trajectory and additional contextual data. The approach follows a three-step procedure: First, a k -d tree-based space discretization is performed, mapping GPS locations to discrete regions. Secondly, a recurrent neural network is trained to predict the destination based on partial sequences of trajectories. The neural network produces destination scores, signifying the probability of each region being the destination. Finally, the routes to the most probable destinations are calculated. To evaluate the method, we compare multiple neural architectures and present the experimental results of the destination prediction. The experiments are based on two public datasets of non-personalized, timestamped GPS locations of taxi trips. The best performing models were able to predict the destination of a vehicle with a mean error of 1.3 km and 1.43 km respectively.

I. INTRODUCTION

Knowledge about a user’s intended route and destination provides many opportunities to improve the driving experience. A smart travel assistance system that knows in advance where the user is likely to go may offer the following functionalities:

- Make smart suggestions for gas or charging stations.
- Advise the user to change the planned route in order to avoid a congested road or area.
- Suggest Points of Interest (POI).
- Show possible parking spots close to the destination.

If the driver is using the built-in navigation system, the planned route and destination are known. In such a case the previously named tasks can be solved in a straightforward manner. However, referring to a study [1] taken by J.D. Power with more than 13,000 consumers, two-thirds of the people who buy a new car prefer to use a substitute instead of the built-in navigation system. In those cases, the planned route and destination are unavailable for the in-car systems, which makes the route and destination prediction necessary to enable the above-listed features.

In this paper, we propose a Long Short-Term Memory (LSTM)-based model using a k -d tree-based space partitioning method to predict the future destination and route of a

¹ Technische Universität Berlin, Germany [patrick.ebel, andreas.vogelsang]@tu-berlin.de, ibrahim.e.goel@campus.tu-berlin.de

² Mercedes-Benz Innovation Lab, MBition GmbH, Berlin, Germany christoph.lingenfelder@daimler.com




Fig. 1: Use Case: Route and Destination Prediction.

car. To perform a prediction, the model requires a partial trajectory in the form of a sequence of Global Positioning System (GPS) coordinates with associated timestamps. The prediction follows a three-step procedure: In the first step, a k -d tree-based space discretization is performed, transforming the analyzed area into a set of discrete regions. Thus, each trip is not represented as a sequence of GPS locations but as a sequence of regions. This sequence is then, along with additionally retrieved metadata, fed to the LSTM. The neural network outputs destination scores, signifying the probability for each region being the destination of the partial trajectory. Subsequently, the highest-scoring, i.e. most probable, routes for the near future can be estimated (see Fig. 1). Our approach does not require any personalized data and is, therefore, able to make a forecast based on data collected from a crowd of anonymous users with no knowledge about personal patterns or regularities. The model is evaluated on two datasets containing trajectory information of taxis in two different cities, namely Porto and San Francisco. Additionally, the model is evaluated on the test set provided in the ECML/PKDD 2015 Kaggle Challenge [2] obtaining a score that would have ranked first out of 381 submissions.

II. RELATED WORK

The recent approaches in the field of destination prediction can be divided into *personalized destination prediction* and *generic destination prediction*. The former approaches [3], [4], [5], [6] try to model the mobility patterns of a specific user. Due to the high degree of regularity in personal mobility patterns, they mostly accomplish a very high prediction accuracy. However, the necessary datasets are very limited in their volume and availability. Additionally, recording and saving trajectories directly connected to a person is highly critical due to privacy concerns. It also limits usefulness, because it will fail exactly when a user needs advice most, namely when doing something new.

This work focuses on generic destination prediction where the data is collected from a crowd of anonymous drivers without any knowledge of personal patterns or regularities. The majority of these approaches divide the space into discrete subsets in the first step. This is necessary because the available datasets are too sparse to cover a sufficient amount of query routes [7]. There are mainly two different strategies applied, *road network mapping* and *spatial partitioning*.

In road network mapping, space gets discretized based on an underlying road network. The GPS locations are mapped to road segments with a unique identifier. The main drawback of these approaches [4], [8], [9], [10] is the high amount of road links, which increases the data sparsity problem.

Spatial partitioning is the process of dividing space into multiple, non-overlapping regions. The spatial partitioning itself can be divided into two approaches, *space-based partitioning* and *trajectory-based partitioning*. The approaches using space-based partitioning [3], [5], [11], [12] map the locations to an overlaying uniform grid, dividing the space into a set of congruent cells. This has its advantages in its simplicity, fast implementation, and intuitive understanding. The essential downside is that the spatial distribution of trajectories is not considered. Hence, the distribution of data points across the cells or regions might be very imbalanced which can lead to a loss in prediction accuracy [13]. To solve this problem, Xue et. al [13] propose a quantile-based as well as a k -d tree-based partitioning strategy, that divide the space based on the density of data points. Due to the more uniform distribution of data points across the grid cells, both result in higher prediction accuracy. Wang et al. introduce a method [7] where space is first divided into uniform grid cells in order to synthesize the nearest cells in the second step. Thus, discrete regions of variable shape are generated.

Besides the different spatial partitioning methods, multiple machine learning approaches are applied in recent work. Most use probabilistic models and especially Markov Models (MMs), where each state represents a location [8], [9], [10], [13]. Thus, each trajectory is modeled as a series of state transitions. Knowing the transitions and their probabilities, the probability of reaching a certain location is calculated. To avoid the data sparsity problem most MMs are based on low-order Markov processes and, therefore, only incorporate the latest time-steps which limits their capability to model long-term dependencies.

In contrast to probabilistic models, recent efforts have been made using Artificial Neural Networks (ANNs) for destination prediction. Endo et al. [11] propose an LSTM, computing transition probabilities between the grid cells for the next time step is applied. To estimate destination probabilities for destinations further in the future, the authors apply the Monte Carlo principle. The main advantages of LSTM-based models lie in their capability to model long-term dependencies and to overcome the data sparsity problem. Liu et al. [14] propose an approach called Spatial Temporal Recurrent Neural Network (ST-RNN), incorporating distance-specific transition matrices to model geographical dependencies. Brébisson et al. [15] propose a Multi-Layer Perceptron

(MLP)-based solution that won the Kaggle challenge on taxi destination prediction [2]. A LSTM-based method which predicts the destination solely based on the individual pick-up and drop-off points of the taxi drivers is presented by Rossi et al. [16].

Another solution to the problem are matching-based algorithms. The main idea is to match a query trajectory with recorded trajectories. The predicted destination or route is then equal to the destination or route of the recorded trajectory having the highest similarity to the query trajectory. In the work of Lam et al. [17] and the early work of Froehlich and Krumm [18], trip matching approaches are applied. A hybrid model that combines MMs with Prediction by Partial Matchings (PPMs) is proposed by Dantas Nobre Neto et al. [4], [19]. However, the main drawback of all matching-based methods is that only routes and destinations which exist in the historical data can be predicted.

III. PROPOSED APPROACH

This work combines a trajectory-based space partitioning with an LSTM-based multi-input model destination prediction model.

A. Definitions and Problem Statement

Definition 1: Let T be the set of trajectories t . A trajectory t is defined as a sequence $(p_n)_{n=1}^N$, where $p_n \in \mathbb{R}^2$ is a single observation and N denotes the length of the trajectory. Each observation $p_n = (\phi, \lambda)_n$ is composed of its latitudinal ϕ and longitudinal λ GPS coordinates. A trajectory's destination is defined as $y_{\text{GPS}} = p_N$.

Definition 2: A partial trajectory $t_p = (p_n)_{n=1}^{N_p}$, is defined as a sub-trajectory of t , where N_p is a random variable following the discrete uniform distribution over $\mathbb{N} \cap [2, N-1]$.

Definition 3: The haversine distance $D_h(p_A, p_B)$ measures the distance between two locations $p_A, p_B \in \mathbb{R}^2$ on a sphere and is computed as follows:

$$D_h(p_A, p_B) = 2 \cdot r \cdot \arctan \left(\sqrt{\frac{a}{1-a}} \right), \text{ with} \quad (1)$$

$$a = \sin^2 \left(\frac{\phi_B - \phi_A}{2} \right) + \cos(\phi_A) \cos(\phi_B) \sin^2 \left(\frac{\lambda_B - \lambda_A}{2} \right),$$

where r is the earth's radius.

Problem: We define the destination prediction task as the problem of predicting y_{GPS} given a partial trajectory t_p .

B. Space Discretization

To overcome the data sparsity problem, a k -d tree-based partitioning approach is applied. A k -d tree is a binary tree in which each node represents a coordinate space in dimension k . Each non-leaf node divides the data space of its parent node into two subspaces of equal size. Thus, a k -d tree can be used to recursively divide the data space into partitions until a defined number of data points per leaf n_{ppr} is reached. In our case, with $k = 2$, each coordinate represents a location and each partition represents a region. The number of resulting




Fig. 2: Spatial partitioning of the Porto area.

regions n_r and their size too, are dependent on n_{ppr} . For both datasets n_{ppr} was heuristically defined such that the regions are small enough in order to achieve a satisfactory accuracy but big enough to avoid data sparsity.

After the discretization each latitude/longitude pair is converted to an integer value $i_r = 1, \dots, n_r$, signifying the region which the point is located in. The final partitioning of the Porto dataset is shown in Fig. 2. We can observe that the regions are smaller in areas with a high density of data points, e.g. the city center or the airport area and larger in more remote areas where less data was collected.

C. Modeling

Due to their ability to keep a memory of previous inputs, LSTMs are considered to be efficient for time-series prediction. Their main advantage to model long-term dependencies is especially important for the destination prediction task at hand. In contrast to matching-based algorithms, LSTMs are able to overcome the data sparsity problem, due to their ability to generalize.

In addition to the trajectory information, our approach also processes contextual information, namely the time of the day, the day of the week, the temperature and the precipitation¹, which are assumed to be constant for each trip. To process the constant contextual information as well as time-series data, namely the trajectory data, the model needs to process the two inputs separately. The architecture of the multi-input model is shown in Fig. 3.

The mapped trajectory data is fed to an *Embedding layer* which turns integer values into dense vectors of a fixed size. Thus, each index and therefore each region identifier i_r is mapped to a vector of size $s_{embedTrip}$. These vectors are initialized randomly but since the so-called embedding table is part of the model parameters they are tuned during training. The embedding table holds all vectors describing the regions and thus is of size $s_{embedTrip} \times n_r$. The idea of using embeddings to represent integer values was inspired

¹We retrieved the weather-related data for San Francisco from www.frontierweather.com and for Porto from www.meteoblue.com

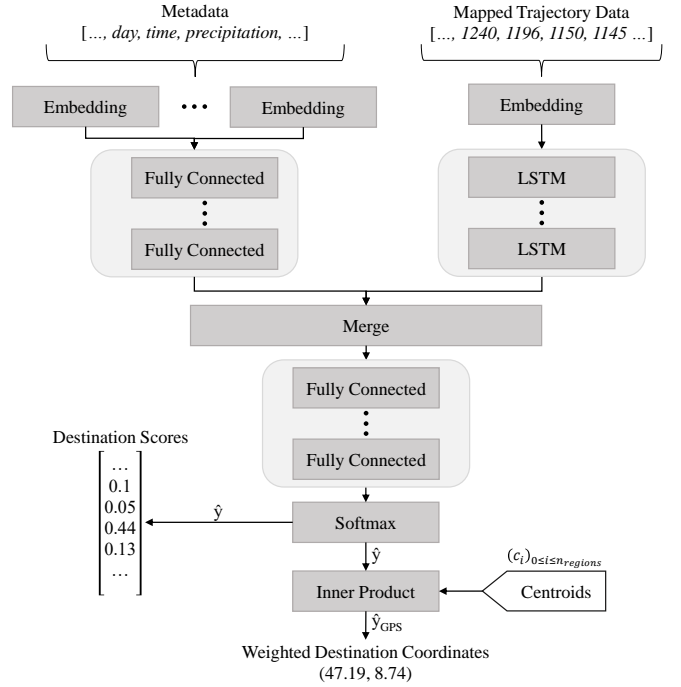


Fig. 3: Architecture of the LSTM-based multi-input model.

by Natural Language Processing (NLP) [22]. The intention to embed the regions is to learn their spatial information and their relationship with each other. The output of the embedding layer is a matrix describing the query trajectory as a sequence of embedding vectors. This matrix is then processed by a stack of n_{LSTM} many-to-one LSTM-layers.

Similar to the trajectory, each attribute of contextual information is mapped to an embedding vector of size $s_{embedMeta}$. Afterward, the metadata embeddings are fed into a stack of $n_{denseMeta}$ fully connected layers. The output of this MLP stack is concatenated with the output of the LSTM stack, and in turn, fed to a stack of n_{dense} fully connected layers. The next layer is a *softmax* layer which normalizes its output inputs such that all outputs add up to 1. Each entry \hat{y}_i of the output vector \hat{y} represents the destination score for region i . Thus, \hat{y} can be interpreted as the probability distribution over all regions.

Additionally, the inner product of \hat{y} and $(c_i)_{i=1}^{n_r}$, where $c_i = (\phi_i^c, \lambda_i^c)$ represents the centroid coordinates of region i , is computed as follows:

$$\hat{y}_{GPS} = \begin{pmatrix} \phi \\ \lambda \end{pmatrix}_{pred} = \begin{pmatrix} \sum_{i=0}^{n_r} \hat{y}_i \phi_i^c \\ \sum_{i=0}^{n_r} \hat{y}_i \lambda_i^c \end{pmatrix},$$

and represents a weighted destination prediction.

D. Test Design

Having two outputs \hat{y}_{GPS} and \hat{y} , the model can be optimized to serve different purposes. Optimizing the model solely regarding \hat{y}_{GPS} leads to a superior top-1 destination prediction as it is required in the mentioned Kaggle competition. However, for many intelligent travel assistance systems,

a probability distribution over multiple destinations is required. In this case, an optimization regarding \hat{y} is needed. To measure the performance and optimize the model regarding the two purposes, two error measures are specified. As in most similar works [15], [17], [16], [20] the *Mean Haversine Distance* serves as the basis for both error measures. Inspired by Besse et al. [20], the first error measure $E_{\text{pred}}^1(\theta)$ is defined as the mean of the haversine distance between the location of the true destination y_{GPS} of the trajectory t and the location of the weighted prediction $\hat{y}_{\text{GPS}}(t_p, \theta)$, where θ is the set of model parameters, adjusted during training. Thus, $E_{\text{pred}}^1(\theta)$ reads as follows:

$$E_{\text{pred}}^1(\theta) = \frac{1}{|T|} \sum_{t_p \in T} D_h(\hat{y}_{\text{GPS}}(t_p, \theta), y_{\text{GPS}}). \quad (2)$$

However, the optimization of \hat{y}_{GPS} using Eq.2 does not imply an optimization of \hat{y} . For illustrating the potential flaw, we assume that y_{GPS} is equal to the location of region A's centroid c_A and that c_A is the midpoint between c_B and c_C being the centroids of region B and C. During training, the penalty for assigning a score of 0.5 to regions B and C would be zero since \hat{y}_{GPS} would be exactly at c_A . This also holds for assigning a score of 1 to region A (which would be the wanted solution). Hence, the network has no reason to shift weight to region A. Thus, for justifying the weights as probability scores and being able to optimize the model with regards to \hat{y} , a second error measure $E_{\text{pred}}^2(\theta)$ is introduced. In $E_{\text{pred}}^2(\theta)$ the distance from each regions' centroid c_i to the true destination y_{GPS} is calculated and weighted based on \hat{y}_i :

$$E_{\text{pred}}^2(\theta) = \frac{1}{|T|} \sum_{t_p \in T} \sum_{i=0}^{n_r} \hat{y}_i(t_p, \theta) D_h(c_i, y_{\text{GPS}}). \quad (3)$$

In the following, the models are evaluated against E_{pred}^1 as well as against E_{pred}^2 and optimized using a weighted combination $E_{\text{pred}}(\theta)$:

$$E_{\text{pred}}(\theta) = \alpha E_{\text{pred}}^1(\theta) + (1 - \alpha) E_{\text{pred}}^2(\theta), \quad (4)$$

with $0 \leq \alpha \leq 1$ being a hyperparameter which controls the importance of the error measures during training.

E. Route Prediction

The route prediction is based on the assumption that the driver is likely to take the best possible route. Having calculated \hat{y} , the destination scores for all regions are known. The route prediction algorithm calculates the route from the last known position to the top- n destinations. Subsequently, each route gets assigned the score of the destination it leads to. Since in most cases, the top- n destinations are in the same area, the calculated routes overlap up to a certain distance. The score of visiting the overlapping parts of the routes is therefore assumed to be equal to the sum of the scores of the individual routes. Therefore, it is possible to retrieve higher scores for near future routes. This is important since, for example, a recommendation system for gas stations will only present a prediction to the driver when the possibility that

the driver is taking the assumed route is over some threshold. In that case, it is not necessary to exactly know the final destination but to know the route in the near future.

IV. EVALUATION

In this section, the datasets, experimental results of the introduced models are presented.

A. Datasets

Both datasets used in this work are publicly available and contain data collected from taxis. The first dataset, the *Porto dataset*, was published on Kaggle as the basis for a taxi trajectory prediction challenge [2] and is also used in several other works [11], [15], [16], [17], [20]. It contains 1,710,670 trajectories of 442 taxis operating in the city of Porto, Portugal. The recording of the data took place over a period of one year starting in July 2013. The second dataset [21] contains 927,976 trajectories of 536 taxis collected over 30 days in San Francisco, USA. This dataset was also processed in different other works [16], [20] and is further referred to as the *San Francisco dataset*. In contrast to the Porto dataset where the trajectories are given as a univariate time-series with an interval of 15 s, the update interval varies for the observations in the San Francisco dataset with a mean of 63.3 s and a standard deviation of 52.58 s.

B. Data Preprocessing

Before the data can be fed to a model it needs to be processed. The data preprocessing consist of four different steps (see Table I): (1) Initially, all trips that are either extremely short ($t < 2$ min), extremely long ($t > 120$ min) or consist of only a single datapoint are deleted. This is done based on the assumption that most of these irregularities occur due to recording issues. (2) To further enhance the data quality, trips containing erroneous data points, e.g. due to GPS errors or incorrectly handled taximeters are conditioned as follows: If the assumed speed between two consecutive points exceeds 240 km/h, the outliers are smoothed by applying a moving median filter. (3) Afterward, all trips that still contain locations outside of the defined area (exemplarily displayed for Porto in Fig. 2) are deleted. (4) Roundtrips or sightseeing trips have no value for destination prediction models but are existent especially in taxi data. To clean those trips we introduce a *roundtrip factor* τ , describing the relation between the length of a trip and the linear distance between start and destination:

$$\tau = \frac{\sum_{n=1}^N D_H(p_n, p_{n+1})}{D_H(p_1, p_N)}. \quad (5)$$

The city topology, as well as the frequency of the update intervals, affects what can be considered as a roundtrip. Thus, the threshold for τ needs to be chosen separately for each dataset. For the Porto dataset, we choose $\tau_P = 3.5$, which corresponds to the 95th percentile of the distribution of τ_P over all trips in the dataset at this preprocessing step. Thus, all trips that are longer than 3.5 times the beeline between their start and destination are deleted. For the San Francisco

TABLE I: Data preprocessing steps.

| Step | Number of trips | |
|------|---------------------|-------------------|
| | Porto | San Francisco |
| - | 1,710,670 (100.0 %) | 927,976 (100.0 %) |
| (1) | 1,638,681 (95.79 %) | 820,108 (88.37 %) |
| (2) | 1,638,681 (95.79 %) | 820,108 (88.37 %) |
| (3) | 1,630,112 (95.29 %) | 815,403 (87.87 %) |
| (4) | 1,545,240 (90.33 %) | 700,197 (75.44 %) |

dataset, we choose $\tau_{SF} = 2.65$, being more restrictive, due to the longer update intervals between two consecutive data points. After preprocessing, the Porto dataset is reduced to 1,545,240 trajectories which accounts for 90.33 % of the data. The San Francisco dataset is reduced to 700,197 trajectories (75.44 %).

C. Hyperparameter Optimization

To find the best performing set of parameters², a hyperparameter optimization is performed for all models on two NVIDIA Quadro P5000 Graphics Processing Units (GPUs). To evaluate the models and to provide the necessary comparability, all models are trained, evaluated and tested on the same preprocessed datasets. Due to the large amount of data the models are trained on 90 % of the data randomly sampled from the datasets. The remaining 10 % are equally split in validation and test set which still results in a set size of 35.000 trajectories for the smaller San Francisco dataset.

During optimization we found that if $\alpha \neq 1$ at the beginning of the training process, E_{pred}^1 and E_{pred}^2 are improving very slowly. However, a good method to prevent this behavior and to accelerate the training process is to set $\alpha = 1$ for the first e epochs. This leads to a fast reduction of E_{pred}^1 and an accompanying slower reduction of E_{pred}^2 . If α is decreased after e epochs, E_{pred}^2 converges close to E_{pred}^1 .

D. Experimental Results

In addition to the LSTM-based approach, three alternative approaches are evaluated. The preprocessing procedure as well as the space discretization process are equal for all the approaches.

The first approach, called *Baseline Algorithm*, is solely based on the trigonometrical relationship between the partial trajectory and the centroid coordinates. First, a set of destination candidates, consisting of the top- k most visited regions is calculated. The predicted destination is then equal to the centroid coordinates of the destination candidate, closest to the extension of the straight line going through the first and last point of the partial trajectory.

The second approach is based on an MLP. Due to the architectural requirements of MLPs, the input has to be a fixed-size vector. As in the work of Brebisson et al. [15], this problem is overcome by feeding the model with the first and last j locations, respectively regions, of the query trajectory.

²<https://doi.org/10.6084/m9.figshare.11698500>

TABLE II: Comparison of the different models.

| Model | Porto | | San Francisco | |
|--------------------|---------------------|---------------------|---------------------|---------------------|
| | E_{pred}^1 | E_{pred}^2 | E_{pred}^1 | E_{pred}^2 |
| Baseline Algorithm | 2504 m | - | 2103 m | - |
| Single-Input MLP | 1595 m | 1635 m | 1573 m | 1672 m |
| Single-Input LSTM | 1460 m | 1567 m | 1300 m | 1462 m |
| Multi-Input LSTM | 1430 m | 1480 m | 1315 m | 1388 m |

Additionally, a single-input, LSTM-based approach is evaluated, that only takes the query trajectory as input and serves as a reference to quantify whether the prediction accuracy can be improved by considering contextual information.

Table II shows the results achieved on the test sets. Since the Baseline Algorithm only outputs a single destination prediction, it is only evaluated with respect to E_{pred}^1 .

In general the models behave similarly on both datasets with E_{pred}^2 being slightly higher than E_{pred}^1 . This originates from the fact, that scores for regions that are on opposite sides from the true destination eliminate each other for some degree when \hat{y}_{GPS} is calculated which results in a lower E_{pred}^1 .

Regarding the different approaches, both LSTM-based models outperform the MLP-based model. This supports the assumption that LSTM-based models are superior when it comes to handling long-term dependencies. However, the small difference in performance implies that the first and last j locations of a query trajectory hold most of the information regarding the final destination. As displayed, concerning E_{pred}^2 , the multi-input LSTM-based models are superior to the single-input model, which only considers the trajectory information. This strengthens the presumption that the consideration of contextual information contributes to better prediction performance. However, the difference in performance between the two approaches is less than 6 % for E_{pred}^2 and for E_{pred}^1 the single-input model even outperforms the multi-input model on the San Francisco dataset. This may be due to the limitation to taxi data. For private cars, the correlation between the metadata and the destination may be higher, due to regularities introduced by commuting patterns.

Fig. 4 shows E_{pred}^1 and E_{pred}^2 of the final three ANN-based models and the Baseline Algorithm according to the given proportion p of the full trajectory. If only 5 % of the trajectory is given, both errors of the ANN-based models are between 2.9 km and 3 km on the Porto dataset and between 3.5 km and 4.3 km on the San Francisco dataset. If each query trajectory consists of 50 % of its full trajectory, on average, the weighted predictions of the multi-input models are 1.237 km (Porto) and 1.188 km (San Francisco) away from the true destination. For all approaches, E_{pred}^1 and E_{pred}^2 decrease with an increasing length of the given partial trajectory. The gap between E_{pred}^1 and E_{pred}^2 increases accordingly for each of the models. However, there is a larger difference for the single-input LSTM compared to the multi-input LSTM.

Fig. 5 shows the distribution of E_{pred}^1 of the final multi-input LSTM-based models according to the trajectory com-

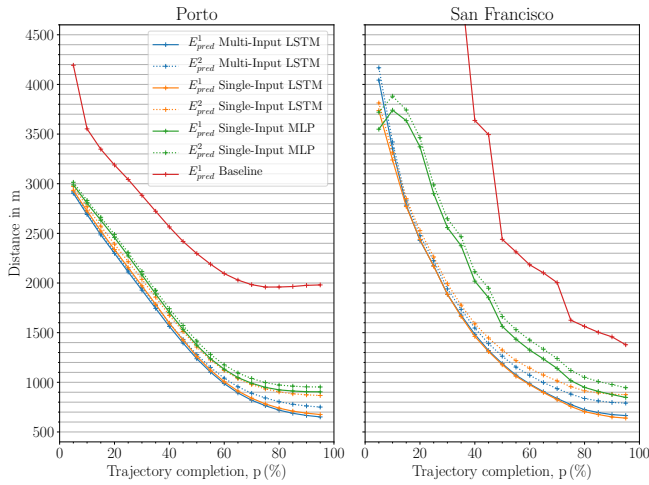


Fig. 4: E^1_{pred} and E^2_{pred} according to trajectory completion.

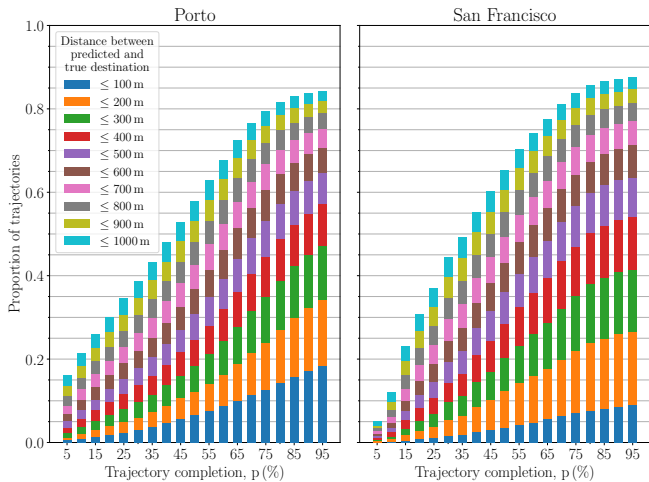


Fig. 5: Distribution of E^1_{pred} according to trajectory completion.

pletion. It can be observed that the prediction accuracy improves regularly with the length of the given trajectory. Comparing the datasets, we can observe that, for low completion rates, the predictions are more accurate on the Porto dataset and that the proportion of predictions that are closer than 100 m to the true destination is twice as high as for the San Francisco dataset.

We additionally evaluated the models based on a snippet of length t , which is randomly sampled from the full trajectory is given. If 2 minutes of the trip are given, E^1_{pred} and E^2_{pred} of the LSTM-based models are about 1.5 km for both datasets.

Fig. 6 shows a route and destination prediction for a query trajectory, starting in the city center (green marker) and ending at the airport (blue marker). In Fig. 6a, the car has only traveled a small distance (orange marker) and therefore the top-5 predicted regions (green squares) are widely spread. However, the traveling direction of the car is already roughly determined. At the time of the second prediction (Fig. 6b), made close to the highway exit which leads to the harbor,

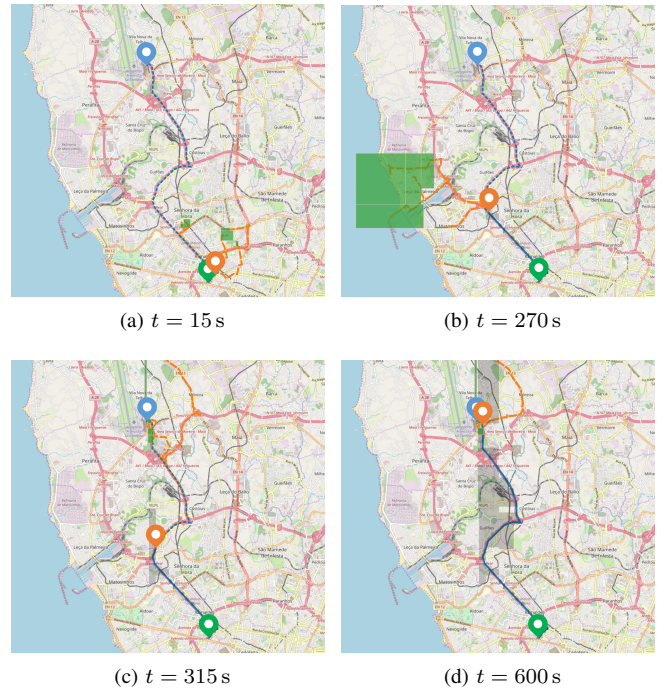


Fig. 6: Route and destination prediction.

all predicted regions lie in the harbor area. Shortly after passing the exit (Fig. 6b), the top-5 predictions jump to the airport area and remain there until the car arrives. Thus, after roughly 50% of the route, the prediction is already very accurate and the top-5 predicted routes (dotted orange lines) match the true ongoing trajectory (dotted blue line) to a high degree. Thus, the approach is not only able to determine that the driver is very likely to go to the airport, but also to predict the route he is going to take.

E. Kaggle

The Kaggle competition is already over, but it is still possible to submit results and receive a ranking on an unknown test set. Our best multi-input model, optimized only against E^1_{pred} achieved a mean haversine distance of 1.995 km and would have ranked first out of 381 submissions.

V. ADVANTAGES OF THE PROPOSED METHOD

Our method is designed to predict the route and destination of a car without using personalized location data. In our approach, data is collected from a crowd of anonymous users with no knowledge about personal patterns or regularities. Therefore, no personalized data is needed to train the models or to make predictions. Personalized destination prediction approaches [3], [4], [5], [6], [16] try to learn the mobility patterns of a specific user to predict future movements. Compared to those approaches, our method is more broadly applicable in practice, especially when it comes to private cars since the handling of personalized location data from customers is highly critical and raises privacy concerns. Additionally, our model can make a prediction even if the user has never been to the city before.

The presented approach not only solves the top-1 destination prediction problem but can also predict multiple destinations and their probabilities. The probability assignment is an important advantage when it comes to practical use cases. It allows recommendation systems to consider whether a suggestion based on the destination prediction should be made to the driver or not. This is important since too many inappropriate suggestions based on unsure destination predictions would lead to the user no longer using the respective system.

Further, our approach achieves good prediction accuracies not only for partial trajectories of any length but also for partial trajectories where the starting point is not known. This may be necessary because a constant recording of the trajectory may not be permitted in practice due to privacy regulations.

VI. CONCLUSIONS AND FUTURE WORK

For many intelligent applications, which aim to improve the driving experience, knowledge about a driver's intended route and destination is crucial. This work introduces three destination prediction models based on ANNs, able to solve this problem. The chosen datasets are analyzed, cleaned and processed. The location data is transformed using a k -d tree-based spatial partitioning approach. The results, achieved on the Porto and San Francisco datasets, show that the best performing models are able to predict the destination based on randomly cut query trajectories with an average accuracy of 1.43 km and 1.3 km respectively. Even without considering the additionally given contextual data of the Kaggle competition, the multi-input LSTM-based model would have scored first out of 381 approaches. Additionally, the models can predict multiple destinations and their probabilities at any time of the trajectory. In combination with the introduced route prediction, valuable input for multiple in-car applications can be produced.

One potential limitation of our work is that we were only able to evaluate it on data collected from taxis. This may introduce bias compared to the results we would have achieved for private cars. In the next steps, we want to evaluate the method on data collected from private cars and analyze the impact of the metadata used for prediction.

As for most deep-learning approaches, our models lack explainability. For future work, it may be interesting to implement an attention mechanism. The neural attention mechanism has the possibility to enhance the interpretability [23] and would therefore allow us to draw conclusions which parts of the trajectories are especially important for predicting the destination.

REFERENCES

- [1] P. Valdes-Dapena, "Most drivers who own cars with built-in GPS systems use phones for directions," 2016. [Online]. Available: <https://money.cnn.com/2016/10/10/autos/car-navigation-frustration/index.html>
- [2] Kaggle Inc., "ECML/PKDD 15: Taxi Trajectory Prediction," 2015. [Online]. Available: <https://www.kaggle.com/c/pkdd-15-predict-taxi-service-trajectory-i/data>
- [3] J. Krumm and E. Horvitz, "Predestination: Inferring Destinations from Partial Trajectories," in *UbiComp*, 2006, pp. 243–260.
- [4] F. Dantas Nobre Neto, C. d. S. Baptista, and C. E. C. Campelo, "Combining Markov model and Prediction by Partial Matching compression technique for route and destination prediction," *Knowledge-Based Systems*, vol. 154, pp. 81–92, 2018.
- [5] C. Manasseh and R. Sengupta, "Predicting driver destination using machine learning techniques," in *16th International IEEE Conference on Intelligent Transportation Systems (ITSC 2013)*, 2013, pp. 142–147.
- [6] R. A. Stegmann, I. Žliobaitė, T. Tolvanen, J. Hollmén, and J. Read, "A survey of evaluation methods for personal route and destination prediction from mobility traces," *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery*, p. e1237, 2018.
- [7] L. Wang, M. Wang, T. Ku, Y. Cheng, and X. Guo, "A hybrid model towards moving route prediction under data sparsity," in *2017 20th International Conference on Information Fusion*, 2017, pp. 1–8.
- [8] Y. Lassoued, J. Monteil, Y. Gu, G. Russo, R. Shorten, and M. Mevisen, "A hidden Markov model for route and destination prediction," in *2017 IEEE 20th International Conference on Intelligent Transportation Systems (ITSC)*, 2017, pp. 1–8.
- [9] X. Li, M. Li, Y.-J. Gong, X.-L. Zhang, and J. Yin, "T-DesP: Destination Prediction Based on Big Trajectory Data," *IEEE Transactions on Intelligent Transportation Systems*, pp. 2344–2354, 2016.
- [10] R. Simmons, B. Browning, Y. Zhang, and V. Sadekar, "Learning to Predict Driver Route and Destination Intent," in *2006 IEEE Intelligent Transportation Systems Conference*, 2006, pp. 127–132.
- [11] Y. Endo, K. Nishida, H. Toda, and H. Sawada, "Predicting Destinations from Partial Trajectories Using Recurrent Neural Network," in *Advances in Knowledge Discovery and Data Mining*, 2017, pp. 160–172.
- [12] P. Pecher, M. Hunter, and R. Fujimoto, "Data-Driven Vehicle Trajectory Prediction," in *Proceedings of the 2016 ACM SIGSIM Conference on Principles of Advanced Discrete Simulation*, ser. SIGSIM-PADS '16. New York, NY, USA: ACM, 2016, pp. 13–22.
- [13] A. Y. Xue, J. Qi, X. Xie, R. Zhang, J. Huang, and Y. Li, "Solving the data sparsity problem in destination prediction," *The VLDB Journal*, vol. 24, no. 2, pp. 219–243, 2015.
- [14] Q. Liu, S. Wu, L. Wang, and T. Tan, "Predicting the Next Location: A Recurrent Model with Spatial and Temporal Contexts," in *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, ser. AAAI'16. AAAI Press, 2016, pp. 194–200.
- [15] A. d. Brébisson, É. Simon, A. Auvolet, P. Vincent, and Y. Bengio, "Artificial Neural Networks Applied to Taxi Destination Prediction," *CoRR*, vol. abs/1508.00021, 2015.
- [16] A. Rossi, G. Barlacchi, M. Bianchini, and B. Lepri, "Modelling Taxi Drivers' Behaviour for the Next Destination Prediction," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–10, 2019.
- [17] H. T. Lam, E. Diaz-Aviles, A. Pascale, Y. Gkoufas, and B. Chen, "(Blue) Taxi Destination and Trip Time Prediction from Partial Trajectories," in *Proceedings of the 2015th International Conference on ECML PKDD Discovery Challenge*, 2015, pp. 63–74.
- [18] J. Froehlich and J. Krumm, "Route Prediction from Trip Observations," in *Society of Automotive Engineers World Congress*, 2008.
- [19] F. Dantas Nobre Neto, C. d. S. Baptista, and C. E. C. Campelo, "A user-personalized model for real time destination and route prediction," in *2016 IEEE 19th International Conference on Intelligent Transportation Systems (ITSC)*, 2016, pp. 401–407.
- [20] P. C. Besse, B. Guillouet, J.-M. Loubes, and F. Royer, "Destination Prediction by Trajectory Distribution-Based Model," *IEEE Transactions on Intelligent Transportation Systems*, pp. 1–12, 2017.
- [21] M. Piorowski, N. Sarafjanovic-Djukic, and M. Grossglauser, "CRAWDAD dataset epfl/mobility (v. 2009-02-24)," 2009.
- [22] Y. Bengio, R. Ducharme, P. Vincent, and C. Janvin, "A Neural Probabilistic Language Model," *J. Mach. Learn. Res.*, vol. 3, pp. 1137–1155, 2003.
- [23] K. Xu, J. Ba, R. Kiros, K. Cho, A. Courville, R. Salakhudinov, R. Zemel, and Y. Bengio, "Show, Attend and Tell: Neural Image Caption Generation with Visual Attention," in *Proceedings of the 32nd International Conference on Machine Learning*, 2015, pp. 2048–2057.