

(Blue) Taxi Destination and Trip Time Prediction from Partial Trajectories

Hoang Thanh Lam, Ernesto Diaz-Aviles,
Alessandra Pascale, Yiannis Gkoufas, and Bei Chen

IBM Research – Ireland

{t.l.hoang, e.diaz-aviles, apascale, yiannisg, beichen2}@ie.ibm.com

Abstract Real-time estimation of destination and travel time for taxis is of great importance for existing electronic dispatch systems. We present an approach based on trip matching and ensemble learning, in which we leverage the patterns observed in a dataset of roughly 1.7 million taxi journeys to predict the corresponding final destination and travel time for ongoing taxi trips, as a solution for the ECML/PKDD Discovery Challenge 2015 competition. The results of our empirical evaluation show that our approach is effective and very robust, which led our team – BlueTaxi – to the 3rd and 7th position of the final rankings for the trip time and destination prediction tasks, respectively. Given the fact that the final rankings were computed using a very small test set (with only 320 trips) we believe that our approach is one of the most robust solutions for the challenge based on the consistency of our good results across the test sets.

Keywords: Taxi trajectories; Trip matching; Ensemble learning; Kaggle; ECML–PKDD Discovery Challenge

1 Introduction

Taxi dispatch planning and coordination is challenging given the traffic dynamics in modern urban cities. The improvement of electronic dispatching of taxis can have a positive impact in a city, not only by reducing traffic congestion and environmental impact, but also improving the local economy, e.g., by mobilizing more customers in cities where the demand of taxis is very high.

Current electronic dispatch systems capture localization and taximeter state from mobile data terminals installed in GPS-enabled taxis. The collective traffic dynamics extracted from the taxis’ GPS trajectories and the stream of data collected in real-time represent a valuable source of information to understand and predict the behavior of future trips. In particular, we are interested in improving taxi electronic dispatch by predicting the final destination and the total trip time of ongoing taxi rides based on their initial partial trajectories. A solution to this challenge is very valuable given that in most cases busy taxi drivers do not

report the final destination of the ride, which makes electronic dispatch planning more difficult and sub-optimal, e.g., ideally the destination of a ride should be very close to the start of the next one.

These two tasks, namely taxi destination and trip time prediction, are the ones that correspond to the Discovery Challenge part of ECML/PKDD 2015 [2]. This paper details our data-driven solution to the challenge that placed our team, *BlueTaxi*, in the 3rd and 7th position in the final ranking for the time and destination prediction task, respectively.

In particular, we are given a data set with roughly 1.7 million taxi journeys from 442 taxis operating in the city of Porto, Portugal, for a period of one year. Besides polylines with GPS updates (at a resolution of 15 seconds), for every taxi trip we are provided with additional information that includes the taxi id, origin call id, taxi stand id, starting timestamp of the ride, etc. The competition is hosted on Kaggle’s platform¹, which allows participants to submit predictions for a given test set whose ground-truth is held by the competition organizers. Submissions are scored immediately (based on their predictive performance relative to a 50% of the hidden test dataset) and summarized on a live *public* leaderboard (LB). At the end of the competition, the other 50% of the held-out test set, i.e., from which the participants did not receive any feedback, is used to compute the final rankings, also known as the *private* LB. The challenge details and dataset characteristics can be found in [2].

The challenge organizers explicitly asked for prediction of trip destination and duration at 5 specific time snapshots, namely 2014-08-14 18:00:00, 2014-09-30 08:30:00, 2014-10-06 17:45:00, 2014-11-01 04:00:00, and 2014-12-21 14:30:00.

Our solution is based on a simple intuition: two trips with similar route likely end at the same or nearby destinations. For example, Figure 1 shows two complete trips extracted from the data. Both trips ended at the Porto airport. Although the first parts of the trips are very different, the last parts of the trips are the same. Therefore, using the destinations of similar trips in the past we can predict the destination of a test trip.

This simple intuition serves as the key idea behind our approach in which we extract features (destination and trip time) from similar trips and then build a model ensemble based on the constructed features to predict taxis’ destination and travel time. In the next sections we detail the feature extraction process, the predictive models of our approach, and present the experimental results that show the effectiveness of our solution.

2 Feature Extraction

From our initial data exploration, we found that the dataset contains missing GPS updates and erroneous information. In fact, the *missing value* column part of the dataset was erroneous: most values of the column were *False*, even though the records included missing information. We detected missing GPS points in the

¹ <https://www.kaggle.com/>



Figure 1. Two trips with different starting points but with the same destination (Porto airport). The final part of the trajectories are very close to each other, which helps to estimate the destination of similar trips.

trajectory by observing large jumps between two consecutive GPS updates, i.e., by considering distances exceeding the distance that a taxi would have travelled at a speed limit (e.g., 160 km/h).

We also found that the trip start timestamp information was unreliable, e.g., some trips in the test set have starting timestamps 5 hours before the cut-off timestamps of the corresponding snapshot but they contain only a few GPS updates. Moreover, some taxi trips have very unusual trajectories, e.g., we noted that some taxi drivers forgot to turn the meter off after completing a trip, for instance in the way back to the city center after dropping a passenger at the airport, such *turning back* trajectories were very difficult to predict.

Considering these issues, we preprocessed the data accordingly and extracted the features that are detailed in this section, which will serve as input to build our predictive models for destination and trip time prediction.

2.1 Feature Extraction for Destination Prediction

For predicting the destination of a given incomplete test trip A , our method is based on trip matching which uses destinations of trips in the training data with similar trajectories to predict the final destination of the given test trip. Figure 1 shows an example of two trips with different starting points both going toward the Porto airport. As we can see, the final part of the trips are very similar, both trips took the same highways before ending at the same destination. We captures such pattern by a set of features described as follows.

- Final destination coordinates and Haversine distance to 10 nearest neighbors: for a given test trip A we search for 10 trips in the training data that are closest to it. Similarly, for each pair of trips A and B we compute the mean Haversine distance (Equation 1) between the corresponding points in their trajectories. We ignored trips B that have fewer GPS updates than the ones in A .
- Kernel regression (KR) as a smooth version of k-NN regression method. Our previous work on bus arrival time prediction shows that KR gives better results than k-NN [6,8]. Destinations calculated by KR were used as features to estimate the final destination of a test trip. KR requires to set the bandwidth parameter, in our experiment, we set it to 0.005, 0.05 and 0.5 to obtain three different KR-based destination predictions corresponding to these values.
- The aforementioned features from KR are sensitive to the metric used for evaluating the similarity between trips. We experimented with dynamic time warping and mean Haversine distance, but only the latter metric was chosen given its efficiency in computation. Moreover, some trips may have very different initial trajectory but still share the same destination, e.g., see Figure 1. Therefore, besides computing KR predictions for the full trip, we also compute them using only the last d meters of the ride during the trip matching step, where $d \in \{100, 200, 300, 400, 500, 700, 1000, 1200, 1500\}$. The models showed that the last 500 to 700 meters of the trips are the most important features for predicting both latitude and longitude of the final destination.
- For a pair of trips A and B , we also consider the distance metric that looks for the best match along their trajectories without alignment. The empirical results showed that KR predictions by best matching gives more accurate results than first aligning the initial part of A 's trajectory with B 's.
- The features described so far do not consider contextual information such as the taxi id, call id, taxi stand, time of day, or day of the week, which are very important, consider for example Figure 2, which shows a heatmap of the destinations for all trips with origin call id 59708 in the training data, as we can see the destination is quite regular, although overall, there are 155 trips, only 8 major destinations were identified. In our work, we exploit this information by using KR to match a test trip with only trips with the same call id, taxi id, day of the week, hour of the day, or taxi stand id. We called the KR models based on these features *contextual KR*. When a contextual field of the test trip has a missing value, the corresponding contextual KR was replaced by the KR prediction with no contextual information. The empirical results showed that contextual KR-based prediction on the call id gave the best results among all KR contextual related features.
- We observed that there are many GPS updates that are erroneous, e.g., coordinates completely outside the trajectory shape or some GPS updates are very close together due to traffic jam or when a taxi is parked or stays still at the same location with the meter on. These type of noisy GPS updates influence the performance of the KR prediction, especially the KR using the mean Haversine distance metrics which requires point to point comparison across trips. Therefore we used the RDP algorithm [3,7] to simplify the trips with parameter ϵ equal




Figure 2. Heatmap of destinations of 155 trips with call id 59708. There are only 8 major destinations. Using call id information one can narrow down the possible destinations of a given taxi.

to 1×10^{-6} , 5×10^{-6} , and 5×10^{-5} . From the simplified trips we extracted all KR-related features as we did for the *raw* trip trajectories.

Besides features extracted via trip matching we also added features extracted directly from the partially observed trips:

- Euclidean distance traveled and Haversine distance between the first and the last GPS update.
- Direction: to define whether the taxi is moving outside of the city or vice versa. We compared the distance between the city center and the first and the last point of the trips. If the former is larger the taxi is considered as moving toward the city center and moving away (e.g., to the countryside) otherwise.
- Time gap between the cut-off timestamps and the starting timestamps. We observed that this feature is not very reliable because the starting timestamps are quite noisy.
- Number of GPS updates (this feature is also noisy due to missing values).
- Day of the week. We observed that the prediction error was higher for some days of the week, but the Random Forest model used in our experiments did not rank this feature high for destination prediction (cf. Section 4).
- The first and the last GPS location.

2.2 Feature Extraction for Trip Time Prediction

The set of features for the time prediction task is very similar to the set of features for destination prediction, with the difference that travel time of closest trips were considered as the target variable instead of destination. The features extracted for this task are as follows.

- Travel time and Haversine distance to 10 nearest neighbors.
- Kernel regression features. All KR related features for destination prediction were also extracted for time prediction, with the difference that *travel time* of closest trips was considered as target variable instead of destination.

In addition to, all features extracted from the partially observed trips as described in the previous section, we also considered the following additional particular features for time prediction which were extracted directly from the incomplete trips observed (i.e., trips that are still ongoing):

- Average speed calculated on the last d meters of the partial trajectory observed so far and on the entire incomplete trip, where $d \in \{10, 20, 50, 100, 200\}$. These features convey up-to-date traffic condition at the moment of making a prediction.
- Average speed calculated on all the incomplete trips observed so far with the starting time no more than an hour apart from the cut-off time-stamp. These features convey information about the traffic condition around the snapshot timestamps.
- Average acceleration calculated on the last d meters of the incomplete trips observed so far and on the entire incomplete trips, where $d \in \{10, 20, 50, 100, 200\}$.
- Shape complexity: the ratio between the (Euclidean) traveled distance and the Haversine distance between the first and the last GPS location. Trips with higher complexity (e.g., zig-zag trips) tend to be trips for which the taxi drivers were driving around the city to search for passengers. Zig-zag trips tend to have longer travel time so it is reasonable to identify those trips beforehand.
- Missing values in the GPS trace were identified by calculating the speed between any pair of consecutive GPS updates. If the estimated speed is over the speed limit \hat{v} km/h even for only one pair of consecutive GPS updates in the partially observed trip, the trip was labelled as a trip with missing GPS updates. We used speed limits $\hat{v} \in \{100, 120, 140, 160\}$ km/h. Trips with missing values tend to have longer travel time.

In total, we have 66 features for the trip time prediction tasks.

3 Predictive Modelling

In order to train the predictive models, first, we created a local training dataset by considering 5 time snapshots from weekdays that resemble the disclosed test set and was fixed for all trained models.

Overall we extracted 13301 trips from these snapshots in the training data, which is a small fraction of the 1.7 million trips in the original training set. We also considered roughly 12000 additional trips from 5 snapshots with start timestamps one hour after the specified test set snapshots. However, models trained on these combined sets yielded improvement only locally but not on the public leaderboard. Therefore, we decided to ignore the additional training trips although in practice to obtain robust results one should consider these additional data for training the models as well.

We observed that feature extraction for the training set is not very efficient because most features were extracted based on nearest neighbor search. In order to speed up this process, we propose to use an index structure based on *geohash*².

² <https://en.wikipedia.org/wiki/Geohash>

We first represent each GPS by its geohash, then we search for the nearest trips via range queries to retrieve trips with a maximum distance threshold of 1km from the first point of the test trip. This simple indexing technique speeds up the nearest neighbor search significantly since range queries are very efficient with geohashes. It is important to note that this technique is a heuristic solution which does not guarantee exact nearest neighbors results. However, we did not observe significant differences in the prediction results when exact or approximated nearest neighbors were used.

In the rest of this section, we describe in detail the different models for each of the prediction tasks.

3.1 Models for Destination Prediction

Given the set of features, any regression model can be used to predict latitude and longitude independently. To this end, we chose Random Forest (RF) [1] given the robustness of its predictions observed in both the local validation set and in the public LB test set. Besides RF models, we also experimented with Support Vector Regression (SVR), but the results on the public LB did not differ much, so we used only RF for the final destination prediction. Moreover, with a RF model we can easily assess the contribution of each feature on the final prediction. With this insight, we know whether a new set of features is relevant or not every time it is added to the model. Thanks to this we could perform feature selection using the `rfcv` function provided with the `randomForest` package in R. With feature selection the obtained results did not differ on the public LB but the results were more robust, in fact it needs fewer trees to obtain similar prediction results.

To handle outliers, we first trained an initial RF model with 2000 trees and removed from training set all trips that prediction error was greater than 90% of error quantile. Subsequently, a new RF model re-trained with the new training set became our final model.

3.2 Ensemble for Trip Time Prediction

For the trip time prediction, before training the models we removed outliers with a trip travel time that exceeded a median absolute deviation of 3.5. Furthermore, we trained the models using as target variable the (log-transformed) *delta time* between the cut-off time point of our training snapshot and the timestamp associated to the last point of the trajectory. That is, our goal is to train models to predict the trip time remaining for a given ongoing trip. This preprocessing strategy led to significantly better results.

Then, we used the features described in Section 2.2 to train a model ensemble for a robust prediction. The individual members of the ensemble include the following regression models: Gradient Boosted Regression Trees (GBRT) [4], Random Forest regressor (RF) [1], and Extremely Randomized Trees regressor (ERT) [5].

In order to produce a single predictor we follow a *stacked generalization (stacking)* [10] approach summarized as follows:

- Remove from our training data, D_T , a subset of the samples (i.e., trips) and split them equally to form a validation set D_v and (local) test set D_t , i.e., $|D_v| = |D_t|$.
- Train n models on the remaining training data $D'_T = D_T \setminus (D_v \cup D_t)$.
- For each model, compute the predictions for the validation set \mathbf{P}^v .
- For the validation set the corresponding true trip time is known, which leads to a supervised learning problem: using the predictions \mathbf{P}_n^v as the input and the correct responses in D_v as the output, train a meta-regressor $M_e(\mathbf{P}_n^v)$ to ensemble the predictions.
- Insert the validation set D_v back into the training data D'_T .
- Train the models again using the $D'_T \cup D_v$ training dataset.
- Predict for the test set D_t to obtain \mathbf{P}^t .
- Ensemble the final prediction for the test set using $M_e(\mathbf{P}_n^t)$.

Table 2 details the models parametrization and predictive performance in terms of the Root Mean Squared Logarithmic Error (RMSLE) defined in Equation 2. We experimented with regularized linear regression and with simple average for the meta-regressor to produce the final prediction. The results of our approach are presented in Section 4.2.

4 Experimental Results

During the competition, the performance on the ground-truth test set held by the organizers was only available for a 50% of the test trips and via a limited number of submissions to the Kaggle’s website. Therefore, in order to analyze the predictions results locally, we split our subset of training data with 13301 trips into two parts: local training and validation datasets. This section reports the experimental results for the public and private LB test sets as well as for the local validation set.

4.1 Destination Prediction

The evaluation metric for destination prediction is Mean Haversine Distance (MHD), which measures distances between two points on a sphere based on their latitude and longitude. The MHD can be computed as follows.

$$\text{MHD} = 2 \cdot r \cdot \arctan \left(\sqrt{\frac{a}{1-a}} \right), \text{ where} \tag{1}$$

$$a = \sin^2 \left(\frac{\phi_2 - \phi_1}{2} \right) + \cos(\phi_1) \cdot \cos(\phi_2) \cdot \sin^2 \left(\frac{\lambda_2 - \lambda_1}{2} \right)$$

and ϕ , λ , $r = 6371\text{km}$ are the latitude, longitude, and the Earth’s radius, respectively.

Table 1(a) shows the mean Haversine distance from the predicted destination to the correct destination on three different test sets. Note that for this task we

Table 1. Destination prediction error in terms of the Mean Haversine Distance (MHD).

(a)		(b)	
Dataset	MHD (km)	Day of the week and time	MHD (km)
Public LB	2.27	Monday (17:45)	2.24
Private LB	2.14	Tuesday (08:45)	2.40
Local validation set	2.39	Thursday (18:00)	2.48
		Saturday (04:00)	2.21
		Sunday (14:30)	2.95

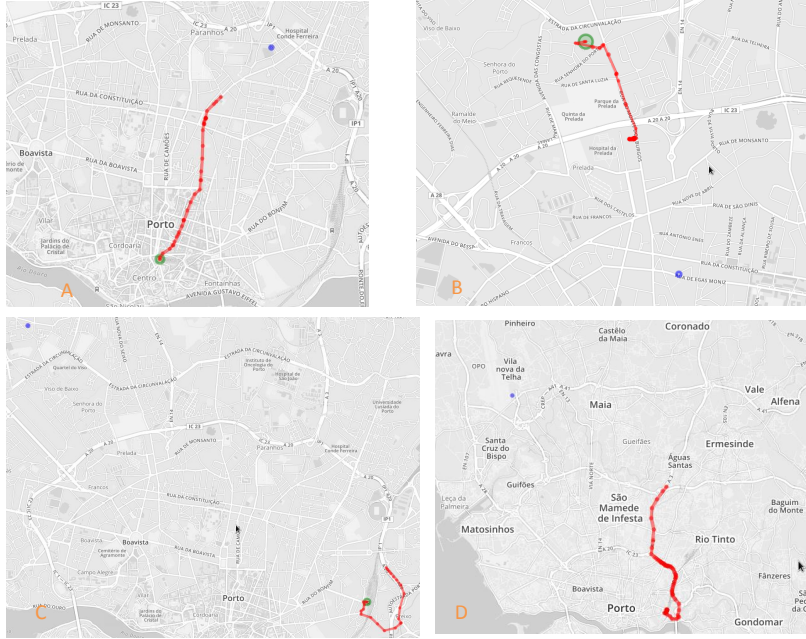


Figure 3. Examples of trips on the LB test set and our prediction (blue circle).

used a 66%–34% training–validation split. We can see that the results are quite consistent and the prediction errors on three different test sets are small.

Table 1(b) reports the prediction error on the local test set at different snapshots. As we can see, among the snapshots, trips in early morning on Saturday are easier to predict while trips on Sunday afternoon are very difficult to predict with significantly higher prediction error. On Sundays, taxis tend to commit longer trips with irregular destinations while early morning Saturdays taxis usually go to popular destinations like airports, train stations, or hospitals.

Figure 3 shows four trips on the LB test set and our prediction. Some predictions are just about one block away from the last GPS location (cases A and B), while some predictions are very far away from the last location (cases C and D). The later cases usually concern trips that took the highway to go to the other side of the city by skipping the city center or go to the airport via A3 motorway.

4.2 Trip Time Prediction

For the trip time task, predictions are evaluated using the Root Mean Squared Logarithmic Error (RMSLE), which is defined as follows:

$$\text{RMSLE} = \sqrt{\frac{1}{n} \sum_{i=1}^n (\ln(p_i + 1) - \ln(a_i + 1))^2} \quad (2)$$

where n is the total number of observations in the test data set, p_i is the predicted value, a_i is the actual value for travel time for trip i , and \ln is the natural logarithm.

Table 2 summarizes the predictive performance of the individual model members of our ensemble locally using a 80%–20% training–validation split.

Table 2. Regression models for trip time prediction and the corresponding performance (RMSLE) in our local test dataset. The models include Gradient Boosted Regression Trees (GBRT), Random Forest regressor (RF), and Extremely Randomized Trees regressor (ERT).

Model	Parameters	RMSLE
GBRT	Default parameter settings used for the GBRT models: learning_rate=0.1, max_depth=3, max_features=n_features, min_samples_leaf=3, min_samples_split=3, n_estimators=128	
GBRT-01	loss='squared_loss', subsample=1.0	0.41508
GBRT-02	loss='squared_loss', subsample=0.8	0.41498
GBRT-03	loss='least_absolute_deviation', subsample=1.0	0.40886
GBRT-04	loss='least_absolute_deviation', subsample=0.8	0.40952
GBRT-05	loss='huber', subsample=1.0, alpha_quantile=0.9	0.41218
GBRT-06	loss='huber', subsample=0.8, alpha_quantile=0.9	0.41108
GBRT-07	loss='huber', subsample=1.0, alpha_quantile=0.5	0.41000
GBRT-08	loss='huber', subsample=0.8, alpha_quantile=0.5	0.40705
GBRT-09	loss='quantile', subsample=1.0, alpha_quantile=0.5	0.40798
GBRT-10	loss='quantile', subsample=0.8, alpha_quantile=0.5	0.40959
RF	Default parameter settings used for the RF models: max_depth=None, n_estimators=2500, use_out_of_bag_samples=False	
RF-01	max_features=n_features, min_samples_leaf=4, min_samples_split=2	0.41674
RF-02	max_features=n_features, min_samples_leaf=1, min_samples_split=1	0.41872
RF-03	max_features=sqrt(n_features), min_samples_leaf=4, min_samples_split=2	0.41737
RF-04	max_features=log ₂ (n_features), min_samples_leaf=4, min_samples_split=2	0.41777
RF-05	max_features=n_features, min_samples_leaf=1, min_samples_split=1, use_out_of_bag_samples=True	0.41865
RF-06	max_features=sqrt(n_features), min_samples_leaf=4, min_samples_split=2, use_out_of_bag_samples=True	0.41731
RF-07	max_features=log ₂ (n_features), min_samples_leaf=4, min_samples_split=2, use_out_of_bag_samples=True	0.41790
ERT	Default parameter settings used for the RF models: max_depth=None, n_estimators=1000, use_out_of_bag_samples=False	
ERT-01	max_features=n_features, min_samples_leaf=1, min_samples_split=1	0.41676
ERT-02	max_features=n_features, min_samples_leaf=1, min_samples_split=2	0.41726
ERT-03	max_features=n_features, min_samples_leaf=1, min_samples_split=2, use_out_of_bag_samples=True	0.41708
ERT-04	max_features=n_features, min_samples_leaf=1, min_samples_split=1, use_out_of_bag_samples=True, n_estimators=3000	0.41735

Table 3. Trip time prediction results on the public and private leaderboards in terms of RMSLE, where the lower the value, the better.

Model ensemble description	RMSLE		
	Public LB	Private LB	Average
ME1: Model ensemble using the average to compute the final prediction.	0.49627	0.53097	0.51362
ME2: Model ensemble using regularized linear regression with L2 penalty as meta-regressor to compute the final prediction.	0.51925	0.51327	0.51626
ME3: Model ensemble using regularized linear regression with L1 penalty (Lasso) as meta-regressor to obtain the final prediction.	0.49855	0.52491	0.51173
ME4: Model ensemble trained using trips with instances whose number of GPS updates is between 2 and 612, which reassemble the distribution observed in the partial trajectories of the test set. The final prediction is computed using mean of the individual predictions.	0.49418	0.54083	0.51750
ME5: Average prediction of model ensembles ME1, ME2, M3, and M4 (3rd place on the public LB).	0.49408	0.53563	0.51485
ME6: Average prediction of model ensembles ME1, ME2 and M3 (3rd place on the private LB).	0.49539	0.53097	0.51318

In Table 3 we present our most relevant results in the official test set of the competition. The table shows the performance of our strategies in the public and private leaderboards, which corresponds to a 50%–50% split of the hidden test set available only to the organizers. Remember that the final ranking is solely based on the private leaderboard from which the participants did not receive any feedback during the competition.

We achieved the best performance in the public leaderboard (RMSLE=0.49408) using ME5, which corresponds to an average of model ensembles as specified in Table 3, this result placed our team at the 3rd place. For the private leaderboard, we also achieved the top-3 final position using ME6 (RMSLE=0.53097). ME6 is similar to the ME5 approach, but it does not include in its average the predictions of ME4, which corresponds to the models trained using a subset of the trips with number of GPS updates in the range between 2 and 612 points.

Note that our ME2 strategy achieved a RMSLE of 0.51925 and 0.51327 in the public and private leaderboards, respectively. The RMSLE of this entry in the leaderboard would have achieved the top-1 position in the final rankings. The competition rules allowed to select two final sets of predictions for the final evaluation. Unfortunately, given the performance observed in the public leaderboard for which we had feedback, we did not select the ME2 approach as one of final two candidate submissions.

The model ensemble that uses regularized linear regression with L1 penalty (Lasso) [9] achieves the best average performance for the whole test dataset with an RMSLE=0.51173.

We used NumPy³ and scikit-learn⁴ to implement our approach for trip time prediction.

³ <http://www.numpy.org/>

⁴ <http://scikit-learn.org/>

5 Conclusions and Future work

We propose a data-driven approach for taxi destination and trip time prediction based on trip matching. The experimental results show that our models exhibit good performance for both prediction tasks and that our approach is very robust when compared to competitors' solutions to the ECML/PKDD Discovery Challenge. For future work, we plan to generalize our current approach to automatically adapt to contexts and select a particular set of features that would improve predictive performance within the given context.

References

1. Breiman, L.: Random forests. *Machine Learning* 45(1), 5–32 (2001)
2. Discovery Challenge: On Learning from Taxi GPS Traces: ECML-PKDD. <http://www.geolink.pt/ecmlpkdd2015-challenge/> (2015)
3. Douglas, D.H., Peucker, T.K.: Algorithms for the reduction of the number of points required to represent a digitized line or its caricature. *Cartographica: The International Journal for Geographic Information and Geovisualization* 10(2), 112–122 (1973), doi:10.3138/FM57-6770-U75U-7727
4. Friedman, J.H.: Greedy function approximation: A gradient boosting machine. *The Annals of Statistics* 29(5) (2001)
5. Geurts, P., Ernst, D., Wehenkel, L.: Extremely randomized trees. *Machine Learning* 63(1), 3–42 (2006)
6. Lam, H.T., Bouillet, E.: Flexible sliding windows for kernel regression based bus arrival time prediction. In: In proceedings of ECML PKDD (2015)
7. Ramer, U.: An iterative procedure for the polygonal approximation of plane curves. *Computer Graphics and Image Processing* 1(3), 244 – 256 (1972)
8. Sinn, M., Yoon, J.W., Calabrese, F., Bouillet, E.: Predicting arrival times of buses using real-time gps measurements. In: *Intelligent Transportation Systems (ITSC), 2012 15th International IEEE Conference on*. pp. 1227–1232. IEEE (2012)
9. Tibshirani, R.: Regression shrinkage and selection via the lasso: a retrospective. *Journal of the Royal Statistical Society: Series B (Statistical Methodology)* 73(3), 273–282 (2011), <http://dx.doi.org/10.1111/j.1467-9868.2011.00771.x>
10. Wolpert, D.H.: Stacked Generalization. *Neural Networks* 5(2), 241 – 259 (1992)