

Online Learning of User-Specific Destination Prediction Models

Erfan Davami
Department of EECS
University of Central Florida
Orlando, FL 32816
Email: erfand@knights.ucf.edu

Gita Sukthankar*
Department of EECS
University of Central Florida
Orlando, FL 32816
Email: gitars@eecs.ucf.edu

ABSTRACT

In this paper, we introduce and evaluate two different mechanisms for efficient online updating of user-specific destination prediction models. Although users can experience long periods of regular behavior during which it is possible to leverage the visitation time to learn a static user-specific model of transportation patterns, many users exhibit a substantial amount of variability in their travel patterns, either because their habits slowly change over time or they oscillate between several different routines. Our methods combat this problem by doing an online modification of the contribution of past data to account for this drift in user behavior. By learning model updates, our proposed mechanisms, *Discount Factor* updating and *Dynamic Conditional Probability Table* assignment, can improve on the prediction accuracy of the best non updating methods on two challenging location-based social networking datasets while remaining robust to the effects of missing check-in data.

I INTRODUCTION

Mechanisms for learning predictive models of human transportation patterns are often foiled by the conflict between two forces: 1) strong correlations between destination and visitation time 2) long periods of disruptions when regular habits are not observed. People are often at work at 10:00 am, in bed at 1:00 am, and have numerous regular periodic commitments. This characteristic can dominate the error metric on the training set, and most feature selection paradigms will identify time and day as important features for predicting destinations.

However, there always exist long periods of disruption when regular habits are not observed. Users go on trips, experience deviations in their work and home routines, or change their lifestyles. In some cases, their behavior patterns will return to the learned baseline, but often the disruption represents a permanent change. During this period of time, visitation time and temporal dependencies will not be informa-

tive, and overreliance on those cues is punished. In this case unless the model can adapt to these changes in behavior, the accuracy will plummet since the majority of samples will be predicted incorrectly. In the case of a non-adaptive model, the learning mechanism will attempt to learn the model that predicts the majority of the samples, effectively sacrificing the samples that occur during those period of time.

To combat this problem, we introduce methods for online learning of user-specific destination prediction models, *Discount Factor* updating and *Dynamic Conditional Probability Table* assignment. The key to our methods is the use of efficient online updating procedures that modify the contribution of past data to the current prediction of the user's behavior. The baseline non-adaptive learning mechanism used in this paper is a Bayes net, which for these location-based social networking datasets achieves comparable performance to the a set of specialized methods for modeling human mobility [2]. Our two adaptation mechanisms perform online modifications of the conditional probability tables used for the inference to model the user's current transportation patterns; however the ideas behind the adaptive mechanisms could be generalized to other types of classifiers as well. This paper demonstrates that the use of online adaptation can offer significant improvements in prediction accuracy, particularly for users with certain mobility profiles.

This paper is organized as follows. Section II presents a selection of related work on learning models of human transportation patterns. Section III describes the location-based social media datasets and our proposed online learning methods. In Section IV, we present an evaluation of our proposed methods against several specialized methods for learning human mobility patterns before concluding the paper.

II RELATED WORK

Learning techniques that leverage temporal dependencies between subsequent locations can perform well at modeling human transportation patterns from

GPS data. Although the assignment of GPS readings to road segments can be a noisy process, GPS generally provides a good continuous stream of data that can be used to learn a variety of models such as dynamic Bayesian networks [8], hidden Markov models [7], or conditional random fields [9]. The problem can also be formulated as an inverse reinforcement learning problem [12] in which the users are attempting to select trajectories that maximize an unknown reward function. Another predictive assumption that can be made is that the users are operating according to a steering model that minimizes velocity changes; this model can be combined with hidden state estimation techniques to predict future user positions [11].

However, in this paper, the datasets that we are using contain user check-ins collected from defunct location-based social networking sites (part of the Stanford Large Network Dataset Collection [1]). Unlike in the Reality Mining dataset [3] or the Microsoft Multiperson Location Survey (MSMLS) [6], the user must voluntarily check-in to the social media site to announce his/her presence to other users. If the user doesn't check in, no data is collected. Thus, there are often significant discontinuities in the data when the user neglects to check in, and it is likely that the users opt to underreport their presence at certain locations. For this type of dataset, we found that the dynamic Bayes network which utilizes temporal dependencies actually performs slightly worse than the simple Bayes net used as the baseline for our model. Our proposed techniques are dynamic in the sense that they change their predictions over time, but they do not explicitly maintain conditional dependencies between subsequent check-ins.

Rather than trying to learn temporal dependencies, our aim is to use the visitation time as the key feature, which is less sensitive to discontinuous data but very sensitive to local changes in the users' habits. These patterns can be discovered by doing an eigen-decomposition analysis of the data [4], and interestingly can be predictive of users' activities several years into the future as shown in [10]. Cho et al. [2] demonstrate that a large section of this dataset can be fitted using a two-state mixture of Gaussians with a time-dependent state prior (Periodic Mobility Model), which we use as one of our comparison benchmarks; the two latent states in their model correspond to the user's home and work locations. The main contribution of this paper is to demonstrate how online learning can improve destination prediction by making the learned models more robust to temporary disruptions in user behavior patterns.

III METHOD

This section describes:

1. the location-based social network datasets used to learn and evaluate our destination prediction models;
2. our baseline non-adaptive Bayes net model;
3. our first proposed method, *Dynamic Conditional Probability Table* assignment (DCPTA), for creating multiple region-specific models for each user;
4. *Discount Factor* adaptation (DF), our second proposed method for diminishing the effects of stale data in the conditional probability tables with a discount factor.

1 DATASETS

The datasets used in this research were extracted from two location-based social networking websites called Gowalla and Brightkite. Cho et al. [2] have made both datasets publicly available at the Stanford Large Network Dataset Collection [1]. Gowalla (2007-2012), gave the users the option to check in at locations through either their mobile app or their website, and Brightkite was a similar social networking website that was active from 2007 to 2011. The data from these two websites consists of one user record per check-in that stores the user ID, exact time and date of the check-in, along with the ID and coordinates of the check-in location. Table 1 shows some features of these datasets, and Figure 1 shows a map of user activity within the United States.

Table 1: Location-based Social Media Datasets

Dataset	Gowalla	Brightkite
Records	6,442,857	4,492,538
Users	107,092	50,687
Average check-ins per user	60.16	88.63
Median check-ins per user	25	11



Figure 1: The scope of user check-ins across the United States for the Brightkite location-based social networking dataset. This location-based service was primarily active in the United States, Europe, and Japan between 2007 and 2011.

2 BASELINE MODEL

For our non-adaptive model, we implemented a simple Bayes net with our modified version of the Bayes Net toolbox in Matlab. A Bayes net is a probabilistic graphical model that represents random variables and their conditional dependencies in the form of a directed acyclic graph. Figure 2 shows the Bayes net structure that we identified after experimenting with other more complicated model structures and dynamic Bayes networks in which the variables were conditioned on their values from the previous time step.




Figure 2: Structure of the Bayes net used as the baseline model for inferring the user's latitude and longitude from the check-in day and time.

In this paper we use a fast simple method for training the network and extracting the most probable values of the output variables (the latitude and longitude nodes). The data structure of the network consists of $h \times d \times l$ matrices for the CPTs (Conditional Probability Tables) in which h and d are respectively the hour and day of the week at which the observation occurs and l is the list of possible check-in locations.

For parameter learning, the corresponding cells of the CPT of the output nodes are incremented; predictions are made by looking up the argmax latitude and longitude values for the user's location based on the check-in time. This method is feasible given the simple independence assumptions in this model and the large size of the dataset.

The main problem with the non-adaptive model is the large distortions which occur in the probability table when the user makes a long-range trip. Imagine a particular user being at some specific location, and following a repetitive pattern of activities for some months. If the user goes on vacation for a month, then the non-adaptive model will deliver a series of incorrect predictions based on the previously learned CPT, only slowly adapting to the new situation. Even once the user is back from the vacation, the effect of the probability distortion (caused by check-ins during the trip) is still clearly visible. We propose two new online learning algorithms capable of overcoming this problem, described in the next sections.

3 DYNAMIC CONDITIONAL PROBABILITY TABLE ASSIGNMENT (DCPTA)

The movement pattern of most users in the dataset consists of a regular pattern of periodic short-range movements punctuated by occasional long-range movements. Figure 3 shows the movement pattern of one randomly selected user in the Brightkite dataset.

The average distance between subsequent check-ins ends up being a good measure of the user's mobility. When the user's movement exceeds twice the average distance between check-ins, it generally signals the start of a new mobility pattern. DCPTA (Dynamic Conditional Probability Table Assignment) uses this measure to determine when to learn a new user profile. By dividing the data into sections each time this jump in movement occurs, we can segment the movement of any user into sections with a relatively low variance which are stored in separate conditional probability tables and can be recovered if the user returns to those regions. Algorithm 1 describes how




Figure 3: Movement history of a user in the Brightkite dataset. (left) Initially, the latitude and longitude of the user’s check-ins are converted to a single distance measurement relative to the first recorded check-in. (middle) The movement history can be divided into sections of low variance for learning the user’s transportation pattern in a particular region by segmenting the data stream based on movement jumps that exceed twice the average distance between check-ins. (right) DCPTA learns a separate conditional probability table for each segment; these tables correspond to a different aspect of the user’s routine.

the DCPTA algorithm works.

Data: Check-ins of a particular user

Result: Dynamic Conditional Probability Table Assignment

Let \mathcal{D} be the set of observed check-in distances

Let \mathcal{S} be the set of observed stored segments

for every new check-in do

 Determine the distance of the current check-in from the initial check-in;

$d_i = \text{dist}(\text{coordinates}_i - \text{coordinates}_0)$;

if $d_i \leq 2 * \text{mean}(\mathcal{D})$ **then**

 load $CPT(\arg \min_{s \in \mathcal{S}} |d_i - s|)$

else

$\mathcal{S} \leftarrow \mathcal{S} \cup CPT(d_i)$;

end

end

Algorithm 1: DCPTA (Dynamic Conditional Probability Table Assignment). This algorithm maintains a running average of the user’s movements relative to an initial location and creates a new location-specific conditional probability table whenever the user’s relative movements exceed a certain threshold.

4 DISCOUNT FACTOR ADAPTATION (DF)

DCPTA is most effective when the user returns to regions governed by previously learned conditional probability tables, and least effective when the user keeps changing his/her habits. For instance, users who are unemployed have a greater flexibility in their daily schedule which translates into a data series with a less defined temporal structure. To learn prediction models for users that exhibit erratic check-in behaviors, we introduce a discount factor, γ , into the

process of updating the CPT such that the existing entry is discounted before incrementing the entry for the new observation. γ can range between 0 and 1; our results indicate that the use of the discount factor improves the online learning but that the learning is relatively insensitive to the magnitude of the parameter. Algorithm 2 gives the procedure for discounting conditional probability tables.

Data: Check-ins of a particular user

Result: Discounted Conditional Probability Table

for $\forall h, d, l$ **do**

$CPT_{\text{Latitude}}(h, d, l) = * \gamma$;

$CPT_{\text{Longitude}}(h, d, l) = * \gamma$;

end

Algorithm 2: DF (Discount Factor Adaptation). Before the CPT is updated with the incoming observation, the discounting procedure is applied. Discounting the conditional probability table reduces the effect of older check-ins on future predictions. This technique works well if the user’s behavior changes slowly over time, rather than rapidly switching between destination-specific transportation patterns.

The discount factor reduces the effect of previous observations on the network, making the most recent check-ins more influential on the location prediction procedure. The advantage of this method compared to the previous proposed method is its lower computational and programming complexity. Applying the discount factor limits the location prediction to a few previous observations while discarding the stale data from older check-ins.

IV RESULTS

We employ two datasets, Gowalla and Brightkite, containing data from real users' check-in information [2]. Our evaluations are performed over the subset of users with greater than 100 check-ins, corresponding to 7600 and 8800 from Brightkite and Gowalla, respectively. We directly compare our methods against the techniques proposed by Cho et al. [2] and Gonzalez et al. [5].

As an additional baseline, we performed location prediction using the Bayes net (BN) described in Section III. This network consists of the four nodes shown in Figure 2, where the predicted latitudes and longitudes are conditionally dependent on the weekday and hour of observation. For each user, the Bayes net is first trained using 15% of check-in data so that the Bayes net can gain some information about the periodic and geographical movement patterns of the user. The test results for this strategy and also the DCPTA strategy are shown in Figure 4.

1 APPLYING DISCOUNT FACTOR ON THE CPTS (THE DF METHOD)

As discussed above, applying a time-dependent discount factor can be a useful way of eliminating travel distortion over the conditional probability tables of our Bayes net. A discount factor of 0 implies a stateless system where counts are reset after each observation; conversely, a factor of 1 applies no temporal decay to the system. The first question we address is whether this discount factor dataset-specific, and how it impacts prediction accuracy. Figure 5 shows the effect of varying the discount factor from 0 to 1. We observe that the performance is relatively insensitive to the precise value of the discount factor and that a discount factor of 0.5 maximizes prediction accuracy for either dataset and is used in subsequent experiments.

Figure 4 shows how the prediction rate of the original Bayes net improves with the enhancement of dynamic conditional probability table assignment (DCPTA) and discount factor (DF). Specifically, we examine how accuracy varies with tolerance, which is defined as the level of error that is acceptable (considered as correct), expressed as a fraction of the total distance traveled by the user. For instance, a tolerance of 0.05 specifies that a prediction must lie within 5% of a check-in to be counted as correct. In this figure, we see that the proposed enhancements improve over the baseline BN over the entire curve, and add approx-

imately 5% to the prediction rate, with DF slightly outperforming DCPTA, over the entire curve.

Figure 6 compares the location prediction results from our methods to the following five recent methods described in the literature:

1. Periodic mobility model [2], denoted as PMM;
2. Periodic and social mobility model [2], denoted as PSSM;
3. Gaussian Mixture Model [5], denoted as G;
4. Last-known location model [2], denoted as RW;
5. Most frequent location model [2], denoted as MF.




Figure 6: Prediction performance of the Bayes net predictor and the proposed enhancements on the Bayes net along with the performance of prior art on the Brightkite dataset.

The *Periodic Mobility Model (PMM)* assumes the majority of the human movement in a network is based on a periodic movement between a small set of locations. The *Periodic and Social Mobility Model (PSMM)* also adds additional parameters to model movement driven by one's social relationships with other members of the network.

Our Bayes net methods are denoted as BN, BN&DCPTA, and BN&DF, and the comparison employs a tolerance level of 2.7%.

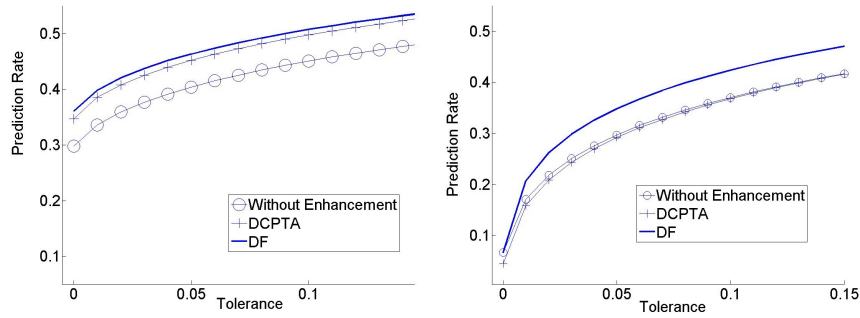


Figure 4: Prediction performance of the Bayes net predictor and the proposed enhancements on the Bayes net (DCPTA and DF) on the Brightkite dataset (left) and the Gowalla dataset (right). Tolerance is the fraction of the total distance traveled by a user that is considered the acceptable distance of the prediction and the actual location of the user in every check-in.




Figure 5: Prediction performance using the proposed Bayes net vs. discount factor on the Brightkite dataset (left) and the Gowalla dataset (right). We observe that the prediction accuracy is relatively insensitive to the choice of discount factor and that a factor near 0.5 maximizes performance on both datasets.




Figure 7: Comparing the movement pattern of different users in the Brightkite dataset. The top two patterns are better predicted using the DCPTA method however the DF method performs better at predicting the bottom two patterns. A possible hypothesis is that DCPTA performs better for users who have multiple short trips, compared to the DF updating.

We observe that the BN without enhancement (36%) performs almost as well as the best of the state-of-the-art approaches, PMM (36.5%) and PSMM (36.3%). However, with our enhancements, we see that accuracy increases by almost 6%, with BN&DF at 42%, slightly outperforming BN&DCPTA at 41%. The remaining baselines (B, RW, MF) are not competitive.

The results shown in Figure 6 are averages over many predictions. Figure 7 provides a more detailed look at some specific instances, and we see that DCPTA does occasionally outperform DF on users with certain features. The top row of the figure shows examples of users for whom DCPTA performs best while the bottom row shows some for whom DF is a better predictor. We hypothesize that users who spend time shuttling between a small set of locations and relatively little time on infrequent long-range trips are better predicted using DCPTA; conversely, DF is better able to handle users who go on long trips and make frequent check-ins away from home.

2 HANDLING MISSING DATA

Missing data within the dataset can be a severe problem for location prediction algorithms. The algorithms used for prediction of GPS data often will not work as well when dealing with check-in data due to the high inconsistency of datapoints. Unfortunately, due to the relatively high overhead imposed on users by a check-in action, the chance of collecting data with missing check-ins is inevitable.

In this section we examine the robustness of the proposed algorithms towards missing data. Seven experiments were conducted using both datasets in which a percentage of check-in data was randomly withheld from the dataset. Figure 8 summarizes the prediction results on each dataset. All of the proposed methods are quite robust to missing data, with the best (DF) showing a drop of only 10% for 70% missing data on the Brightkite dataset (left) and negligible loss on the Gowalla dataset (right). This confirms our belief that there is significant redundancy in the second dataset that can be exploited. Somewhat surprisingly, we observe a slight *improvement* in DCPTA's performance with missing data on the Gowalla dataset. We attribute this to the fact that withholding data has the effect of reducing check-ins corresponding to long-range travels, which results in a reduction of such outliers.




Figure 8: Prediction rates of proposed algorithms when applied to datasets with missing data. Some fraction of the check-in data is randomly withheld and then predicted using the belief network and the proposed enhancements. Our approaches exhibit robustness to missing data.

3 COMPLEXITY

We briefly summarize the computational complexity and storage requirements for the proposed methods. The core data structure behind our methods is a conditional probability table (described in Section 3B). Storing such a discretized table, even at double-precision, is cheap: a table with $24 \times 7 \times 700$ double-precision cells requires less than 1MB of memory.

Conditional probability tables are also computationally efficient, affording constant-time updates. Finding the maximum in the table employs an exhaustive scan that is linear $O(N)$ with respect to the number of cells, N ; in practice, since the number of cells is around 100K, this remains very efficient.

The DCPTA method (Section 3.C) requires multiple CPTs for every segment of the users' movement pattern, thus requiring a memory growth of $O(s)$ where s denotes the number of segments. In terms of computational complexity, the algorithm must search s CPTs in order to load the right CPT for future use, resulting a computational complexity of $O(sN)$.

Finally, the DF method simply multiplies the CPT by a real number (discount factor). This procedure has no impact on the memory usage of the belief network however, but increases the computational complexity equivalent to a scalar matrix multiply, which is theoretically $O(N)$ but very efficient on current hardware.

Table 2 presents measured running times for each method on both datasets. The processing was done using Matlab 2012a, an Intel Quad-core Xeon Processor and 18GB of memory.

Table 2: Computational time required for each dataset (minutes)

Name of Dataset	Gowalla	Brightkite
Processed Users	8800	7600
Processed check-ins	2,694,344	3,399,651
Belief Network	9	12
BN&DF	10	12
BN&DCPTA	19	20

V CONCLUSION

In this paper we present two new algorithms for on-line learning of user-specific destination prediction models, Dynamic Conditional Probability Table Assignment (DCPTA) and Discount Factor updating (DF). Although we describe the use of our online update procedures for a Bayes net model, the same intuitions behind the discounting of stale data and threshold switching between multiple models can be applied toward online learning procedures for other types of classifiers. Our proposed destination prediction model leverages the predictive power of visitation times while rapidly adapting to schedule changes by the users. Adapting to changing user habits allows our model to achieve better predictive performance than the best static models which are continually penalized by non-stationary user behavior.

ACKNOWLEDGMENTS

This research was funded by AFOSR YIP FA9550-09-1-0525 and DARPA award N10AP20027.

References

- [1] Stanford Large Network Dataset Collection. <http://snap.stanford.edu/data/index.html>.
- [2] E. Cho, S. Meyers, and J. Leskovec. Friendship and mobility: user-movement in location based social networks. In *Proceedings of the International Conference on Knowledge Discovery and Data Mining*, 2011.
- [3] N. Eagle and A. Pentland. Reality mining: Sensing complex social systems. *Journal of Personal and Ubiquitous Computing*, 2005.
- [4] N. Eagle and A. Pentland. Eigenbehaviors: Identifying structure in routine. In *Proceedings of the ACM International Conference on Ubiquitous Computing*, 2006.
- [5] M. Gonzalez, C. Hidalgo, and A. Barabasi. Understanding individual human mobility patterns. *Nature*, 453(7196):779–782, 2008.
- [6] J. Krumm and E. Horvitz. Predestination: Inferring destinations from partial trajectories. In *Proceedings of the ACM International Conference on Ubiquitous Computing*, 2006.
- [7] J. Letchner, J. Krumm, and E. Horvitz. Trip router with individualized preferences (TRIP): Incorporating personalization into route planning. In *Proceedings of Conference on Innovative Applications of Artificial Intelligence*, 2006.
- [8] L. Liao, D. Fox, and H. Kautz. Learning and inferring transportation routines. In *Proceedings of National Conference on Artificial Intelligence*, 2004.
- [9] L. Liao, D. Fox, and H. Kautz. Extracting places and activities from GPS traces using hierarchical conditional random fields. *International Journal of Robotics Research*, 26:119–134, 2006.
- [10] A. Sadilek and J. Krumm. Far out: Predicting long-term human mobility. In *Proceedings of National Conference on Artificial Intelligence*, 2012.
- [11] B. Tastan and G. Sukthankar. Leveraging human behavior models to improve path prediction and tracking in indoor environments. *Pervasive and Mobile Computing*, 7:319–330, 2011.
- [12] B. Ziebart, A. Maas, J. A. Bagnell, and A. Dey. Maximum entropy inverse reinforcement learning. In *Proceedings of National Conference on Artificial Intelligence*, volume 3, pages 1433–1438, 2008.