Bayesian Network Movement Model

Andreas Züfle
Department of Geography and Geoinformation Science, George Mason University, USA

Abstract

To unify and integrate various types, resolutions, and levels of uncertainty of user location data, we employ a Bayesian learning approach, which allows to iteratively add new knowledge to refine a model describing the motion of an object in space and time. By explicitly modelling uncertainty, we maintain a notion of data reliability, such that the confidence of any query and mining results can be assessed. Based on this motion model for individual users, we present our algorithms for estimating the continuous location of users based on sparse observations. Our approach uses a global traffic model using a Markov-chain as an apriori-model, which is learned from all available historic trajectory data. Starting from this apriori-model, we use observations of individual users to adapt this model, using a forward backward approach to add new knowledge to the model. This yields as user-specific aposteriori-model, which captures information of their observation, and uses the apriori-knowledge to model the error and uncertainty in-between these observations. As verified by our empirical study on real trajectory data, this model allows to predict the location of objects in-between discrete observation much more accurate than competing approaches, thus significantly reducing the uncertainty in spatio-temporal data.

1 Introduction

Both the current trends in technology such as smart phones, general mobile devices, stationary sensors and satellites as well as a new user mentality of utilizing this technology to voluntarily share information produce a huge ood of geo-spatial and geo-spatio-temporal data. This data flood provides a tremendous potential of discovering new and possibly useful knowledge. In addition to the fact that measurements are imprecise, due to the physical limitation of the devices, some form of interpolation is needed in-between discrete time instances. These issues introduce the notion of uncertainty in the context of spatio-temporal data management - an aspect raising an imminent need for scalable and flexible data management.

Managing, querying and mining spatio-temporal data has received a large amount of research interest in the past years. In most of these works however, the assumption is made that the trajectory, i.e., the function of a spatio-temporal object that maps each point in time to a position in space, is known entirely without any uncertainty, such as the trajectory depicted in Figure 1(a). A survey on techniques for managing, querying and mining trajectory data without uncertainty is given in [23]. However, it is not viable to maintain positions of moving objects continuously for each point in time for a series of reasons:

- Positions are captured sparsely to save energy and preserve battery of the capturing devices.
- Storing arbitrary continuous functions at high accuracy would require a very high space complexity, thus having detrimental impact on query performance.
- In most applications, measurements and observations are taken at discrete times only, e.g. due to limited bandwidth or energy constraints (e.g. in applications using GPS technology), thus the information of the full trajectory is not available.
- Some systems only allow to track the position of an object at predefined spatial positions (e.g. systems using RFID technology).

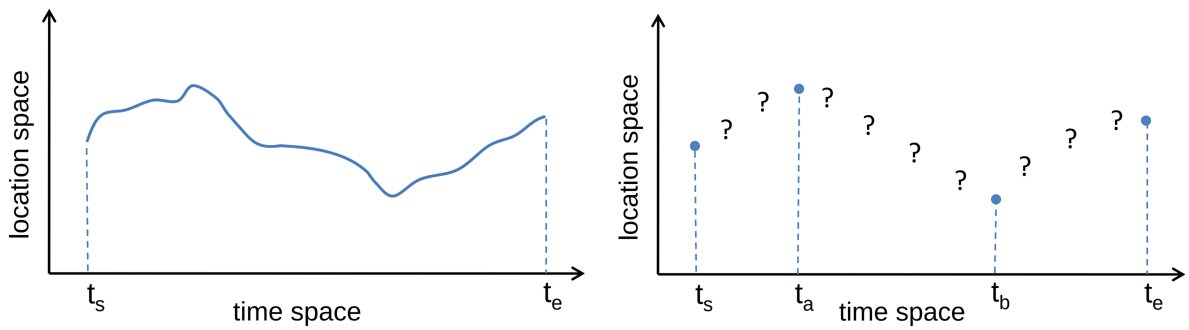

(a) Ideal Case

(b) Practical Case

Figure 1: Spatio-Temporal Data

For practical applications it is thus mandatory to consider uncertainty in spatio-temporal data, such as depicted in Figure 1(b). As an exemplary application, consider the problem of monitoring iceberg activity in the North Atlantic. Ships transiting between Europe and east coast ports of North America traverse a great circular route that brings them into the vicinity of icebergs carried south by the cold Labrador Current near the Grand Banks. It was here that the R.M.S. Titanic sank in 1912, after it struck an iceberg. This disaster resulted in the loss of 1517 lives and led directly to the founding of the The International Ice Patrol (IIP) in 1914. The mission of the IIP is to monitor iceberg danger near the Grand Banks of Newfoundland and provide the locations of all known ice to the maritime community. The IIP does this by sighting icebergs, using visual observations from ships and aircrafts, as well as data from buoys and radars. A database stores the recorded positions and extents of observed icebergs and data models are used to predict their movement. Ignoring the uncertainty in such data, e.g., by treating the expected position of an iceberg as truth, may lead to disasters when icebergs deviate from the expected position, running into ships having a false feeling of safety.

The aim of this work is to explicitly model the uncertainty in spatio-temporal data. Therefore, we model the motion of an object in-between discrete observations using a Bayesian learning approach. This approach uses global traffic patterns, learned from all available historical data, to build a prior model. This prior model can be adapted using observations of an individual model. The following Section 2 surveys the related work, and their lack of ability of assess the probability distribution of the uncertainty region of objects. Section 3 surveys the data model that we emply, which is based on previously published material [3]. To adapt this model to account for observations of individual users, Section 4 presents a Bayesian learning approach found in more details in [7]. An experimental evaluation, showing that this model can vastly improve the location prediction in spatio-temporal databases, is given in Section 5.

2 State-of-the-Art

In recent years, the research community often employed linear interpolation to model the movement of spatio-temporal objects [9, 11, 15], as depicted in Figure 2(a). Based on this model, [13] and [14] addressed the problem of query processing. Clearly, whenever the sampling rate becomes low, the true trajectory of an object might deviate greatly from the linear interpolation. Approaches using more complex models than linear interpolation suffer from the same problem [12]. To mitigate this problem, *Conservative Space-Time Approximation Models* capture the uncertainty explicitly, by bounding all possible (time, location)-pairs of an object by a simple geometric structure in time and space. The result is a spatio-temporal approximation that is based on previous knowledge such as the maximum speed and the maximum acceleration of objects. An example for the one-dimensional case is shown in Figure 2(b), using the maximum speed of objects to obtain a conservative two-dimensional (space, time)-approximation. Generally, such approaches utilize various geometrical shapes, such

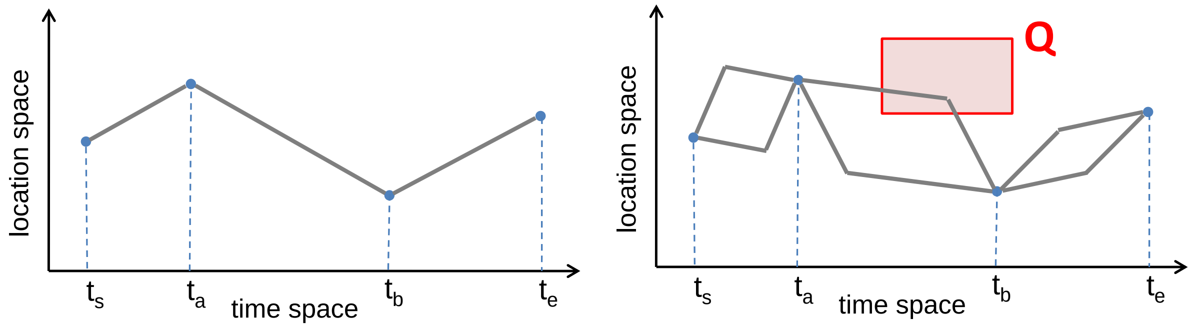

(a) Linear Interpolation (b) Spatio-Temporal Space Approximation

Figure 2: Interpolation between observations

as sheared cylinders [17, 18, 19], diamonds [6] and so called beads [5, 16] for the case of two spatial dimensions (the third dimension is time). The main problem of all these approaches is that no probability information is given for any object approximation. Thus, it is not possible to assess the probability of an object to satisfy some query predicate: For example, for the query window Q depicted in Figure 2(b), the only conclusion that can be made is that the approximated object may possibly intersect Q. However no information about the likelihood of this event can be assessed.

Furthermore, solutions have been presented to estimate the location distribution of objects [6, 1, 20] at each point in time. However, by describing the distribution of an object independently at each point of time, time dependencies of trajectories are ignored. By disregarding these correlations between points in time, completely impossible trajectories are modelled.

Example 1: For example, if the last observed location of a vehicle at time t was just before an intersection going North or South. Such models may predict that at times $t+1$, $t+2$, and $t+3$, the object are predicted to be North with a probability of 50% and South with a probability of 50%. Such models predict impossible future trajectories where the object warps forth and back between the northern and the southern road. The reason is that the temporal correlation, that the object must be North at time $t+2$ given the object is North at time $t+1$ is discarded.

3 Modeling Spatio-Temporal Motion

The key idea of the proposed approach is, similar to [10]: to model possible object trajectories by stochastic processes, more precisely a Markov chain. Employing the Markov chain model for representing spatio-temporal data has three major advantages over previous work:

1 It allows answering queries such that results are associated with corresponding probabilities.

2 Dependencies between object locations at consecutive points in time are taken into account.

3 It is possible to reduce all queries on this model to simple matrix operations, because transitions between spatial entities over time can be performed by matrix multiplications, for which there exist efficient solutions.

Following state-of-the-art models, discrete space and time domains S and $\mathcal{T}$ are assumed, where space is defined by coordinates and time denotes points in time. In practice, discretization techniques (e.g. a unidistant grid) can be used to satisfy this assumption.

Formally, let $S = \{s_1, ...s_{|S|}\} \subseteq \mathbb{R}^d$ be a finite set of possible locations in space which we call *states* and let $\mathcal{T} = \mathbb{N}_0^+$ be the time-space. Consequently, a (certain) object o that moves in space is represented by a *trajectory* given by a function $o : \mathcal{T} \rightarrow S$ that defines the location $o(t) \in S$ of o at a certain point of time $t \in \mathcal{T}$.

Definition 1 (Spatio-Temporal Object): Given a d-dimensional spatial domain $\mathcal{S}$ and a time domain $\mathcal{T}$, a spatio-temporal object comprises a set Θ^o of pairs $\Theta_i^o = (\Theta_i^o.s, \Theta_i^o.t) \in \mathcal{S} \times \mathcal{T}$.

Semantically, each pair $\Theta_i^o \in \Theta^o$ corresponds to the observation that o is located at location $\Theta_i^o.s$ at time $\Theta_i^o.t$.

Following the tradition of uncertain spatial non-temporal databases, we suppose that the locations of an uncertain spatio-temporal object $o \in \mathcal{DB}$ at time t are realizations of a random variable $o(t)$. This model implies that a spatio-temporal objects is described by a series of random variable $(o(t))_{t \in T}$. This consideration directly equals the definition of a *stochastic process* [4]:

Definition 2 (Stochastic Process): A stochastic process X is a family of random variables $(X(t) \in S)_{t \in \mathcal{T} \subseteq \mathbb{R}}$. If $T \subseteq \mathbb{N}$, we say that X is a discrete stochastic process.

Note that in general, these random variables $X(t_1), X(t_2), t_1, t_2 \in T$ are stochastically dependent.

Definition 3 (Uncertain Object Trajectory): Given the spatial domain $\mathcal{S}$ and the time domain $\mathcal{T}$, an *uncertain object trajectory* $o(t) \in S$ of an object $o \in \mathcal{DB}$ is a stochastic process $(o(t) \in S)_{t \in \mathcal{T}}$.

An example of an uncertain object trajectory of an object $o \in \mathcal{DB}$ is illustrated in Figure 3. The raster models all possible locations (i.e., states) in S, shown for the time sequence $\langle t_0, \ldots, t_3 \rangle$. Here we assume that object o has been observed at time t_0; all locations of o that follow are uncertain. Consequently, the uncertain trajectory of o comprises all possible trajectories starting at $o(t_0)$. According to Definition 3, the uncertain motion of an object is defined as a stochastic process. In this work, the motion of moving objects through space and time is model by a first-order Markov-Chain model. Formally:

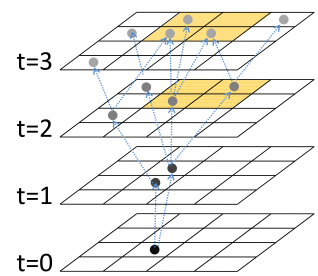

Figure 3: Transitions

Definition 4: A stochastic process $o(t), t \in \mathcal{T}$ is called a Markov-Chain if and only if

$$\forall t \in \mathbb{N}_0 \forall s_j, s_i, s_{t-1}, ...s_0 \in S : P(o(t+1) = s_j | o(t) = s_i, o(t-1) = s_{t-1}, ..., o(0) = s_0) = P(o_{t+1} = s_j | o_t = s_i)$$

The above equality is called Markov assumption or Markov property. The conditional probability

$$P_{i,j}(t) := P(o(t+1) = s_j | o(t) = s_i)$$

is the (single-step) *transition probability* of state s_i to state s_j at time t.

In this work we assume that the transition probabilities $P_{i,j}(t)$ are given, e.g. derived from expert knowledge or derived from historical data. For example, the current of water in the Atlantic ocean can be used to infer the transitions of icebergs. For traffic data, the transition probabilities at road intersections can be estimated using historical data [2, 7].

Let $P(o,t)$ be the distribution vector of an object o at time t, such that $(P(o,t) = s_1), \ldots, P(o,t) = s_{|\mathcal{S}|})$, $p_i \in P(o,t)$ corresponds to the probability that o is located at state s_i at time t. The distribution vector of o at time $t+1$ can be inferred from $P(o,t)$ as follows:$P(o,t+1) = P(o,t) \cdot M$

The *m-step transition probability* $P_{i,j}^t$ is the probability that an uncertain object o that is located at state s_i at time t, will be located at state s_j at time $t+m$ and can be computed exploiting the Equations of Chapman-Kolmogorov ([8]) as follows: $(P_{i,j}^m) = M^m$

Given the probability distribution $P(o,t)$ of an uncertain object o at time t, the probability distribution $P(o,t+m)$ of o at time $t+m$ can be computed by $P(o,t+m) = P(o,t) \cdot M^m$.

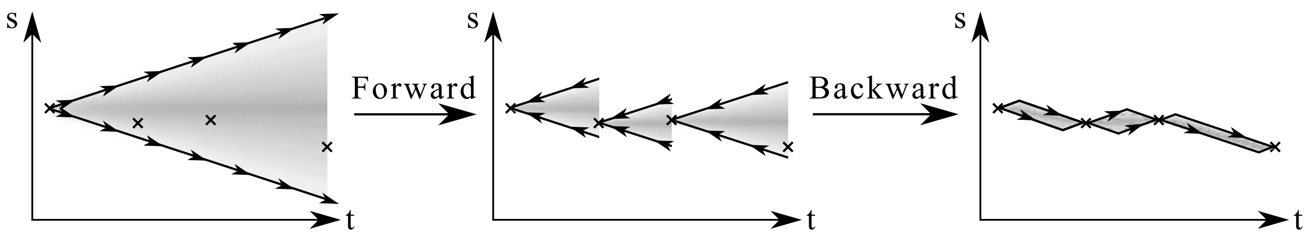

Figure 4: An overview over our forward-backward-algorithm.

To estimate the location of an object o, the Markov-model requires the position of an object o at some time in the past. In a nutshell, the Markov-model allows to carry information about the position of an object o at some time t over to future points of time $t + k$. The quality of this information, and thus the predictive power of the model, diminishes as the time horizon k increases. This fact is attributed to cumulative prediction errors incurred over time. Furthermore, the Markov-model, by itself, is not well-suited to model the motion of real entities, usually follow an optimal routes (shortest, fastest, etc.), rather than randomly walking the network until they accidently find their destination. In the next section, this problem will be solved by employing a Bayesian learning approach to adapt transition probabilities of an object given all observations in the past, in the current, and in the future. This allows to account for all observations of an entity, using the Markov-model only to capture the error in-between discrete observations.

4 Adapting the Model to Observations

In the following, let $\Theta^o = \{< t_1^o, \theta_1^o >, ..., < t_m^o, \theta_m^o >\}$ denote the set of observations of a spatio-temporal object o. Each observation $< t_i^o, \theta_i^o >$ of o is a pair, containing the time t_i^o at which o has been observed in state θ_i^o. This set is assumed to be sorted, i.e., $j > i \Rightarrow t_j^o > t^o$. The traditional sampling approach uses the transition probabilities $P(o(t+1) = s_j|o(t) = s_i)$ given by the Markov chain to create sample trajectories. To incorporate the knowledge given by a set of observations Θ^o of an object o, we need to obtain the probability

$$P(o(t+1) = s_j|o(t) = s_i, \Theta^o),$$

that is the probability that object o transitions to state s_j from state s_i at time t, given all observations.

In a nutshell, this problem can be solved using a forward-backward approach.

Forward-Run: Starting at the time of the first observation t_1^o with the initial observation θ_1^o, we perform transitions of object o using the original Markov chain of o until the final observation at time $t_{|\Theta^o|}$ is reached. During this *Forward*-run, Bayesian inference is used to construct a time-reversed Markov-model R_t^o of o at time t given observations in the past, i.e., a model that describes the probability $R_{ij}^o(t) := P(o(t-1) = s_j|o(t) = s_i, past^o(t))$ of coming from a state s_j at time $t-1$, given being at state s_i at time t and given the observations $past^o(t) := \{\theta_i^o | t_i^o < t)\}$ in the past. Using the Theorem of Bayes, this transition probability can be rewritten as follows

$$R^o(t)_{ij} := P(o(t-1) = s_j|o(t) = s_i, past^o(t)) = \frac{P(o(t) = s_i|o(t-1) = s_j, past^o(t)) \cdot P(o(t-1) = s_j|past^o(t))}{P(o(t) = s_i|past^o(t))}$$

The probability $P(o(t) = s_i|o(t-1) = s_j, past^o(t))$ can be rewritten as $P(o(t) = s_i|o(t-1) = s_j)$, exploiting the Markov property (Eq. 4). This probability is given by the original Markov-chain $T^o(t)$. Furthermore, both priors $P(o(t-1) = s_j|past^o(t))$ and $P(o(t) = s_i|past^o(t))$ can be computed in a single run, as shown in more detail in [7, 24].

Backward-Run: Then, in a second step, we traverse time backwards, from time $t_{|\Theta^o|}$ to t_1, by employing the time-reversed Markov-model $R^o(t)$ constructed in the forward step. Again, Bayesian inference is used to construct a new Markov model $F^o(t-1)$ that is further adapted to incorporate knowledge about observations in the future. This new Markov model contains the transition probabilities

$$F^o_{ij}(t-1) := P(o(t) = s_j | o(t-1) = s_i, \Theta^o). \tag{1}$$

for each point of time t, given all observations, i.e., in the past, the present and the future. Using the Theorem of Bayes the same way as before, we switch the direction of time back to forward, by rewritting:

$$P(o(t+1) = s_j | o(t) = s_i, \Theta^o) = \frac{P(o(t) = s_i | o(t+1) = s_j, \Theta^o) \cdot P(o(t+1) = s_j | \Theta^o)}{P(o(t) = s_i | \Theta^o)}$$

The backward probability $P(o(t) = s_i | o(t{+}1) = s_j, \Theta^o)$ is available directly in matrix F^o that we computed in the Forward-run. The priors $P(o(t+1) = s_j | \Theta^o)$ and $P(o(t) = s_i | \Theta^o)$ can be computed in a single run as again shown in [7, 24].

To illustrate the idea of our Forward-Backward learning, consider the following example.

Example 2: Figure 4 visualizes a one-dimensional uniform random walk, i.e. a very simple Markov chain where the object may move to adjacent states with a uniform distribution. Figure 4(a) shows the initial model, using knowledge about the first observation only. In this case, a large set of ($time$,$location$) pairs can be reached with a probability greater than zero. The shading of reachable ($time$,$location$) pairs indicates the likelihood of these pairs, assuming that a detour is less likely than a direct path.[1] The adapted model after the forward phase is depicted in Figure 4(b), significantly reducing the space of reachable ($time$,$location$) pairs and adapting respective probabilities, thus drastically improving the model. Since the model of Figure 4 (b) is obtained trivially, the main contribution of this section is the backward-phase which adapts the Markov model to observations in the future. This task is not trivial, since the Markov-property does not hold for the future, i.e., the past is *not* conditionally independent of the future given the present. During the backward phase, we traverse the Markov chain backwards, from time $t_{|\Theta^o|}$ to t_1, by employing the information acquired in the forward-phase. This phase yields the final transition matrices for each point in time, such that all observations Θ^o are taken into account for the adapted model. The resulting final transition matrices $F^o(t-1)$ contain the transition probabilities $F^o(t-1)_{ij} = P(o(t) = s_j | o(t-1) = s_i, \Theta^o)$. Figure 4(c) shows the resulting final model after the backward phase.

5 Experimental Evaluation

Our evaluation uses the $T - drive$ dataset containing GPS trajectories of more than 10,000 taxis in the city of Beijing [22, 21]. We remove all GPS trajectories having a sampling frequency of less than one observation per 10 seconds, thus reducing the dataset by about 20% in size. The roat-network we use is generated using the map-construction techniques of [2] to obtain both a set of possible states (corresponding to crossroads) and a transition matrix reflecting the possible movements of the cabs. This process yields a state space consisting of about 3000 states and the corresponding transition matrices and direct edges between states. We assume that a-priori, all objects utilize the same Markov model M. In this data set, both the underlying road network as well as the estimated transition probabilities are taken from real trajectories.

We evaluate the effectiveness of the forward-backward model adaption in comparison to other approaches. Therefore, we select random trajectories from our dataset, using only one of the GPS observations every 100 seconds. Then, we compare the predicted location of all approaches by the ground-truth given by the more frequently sampled raw trajectory.

[1] This assumption is only made for illustration of this example. In general, a detour via a highway may be more likely than a direct path through a lake. These likelihoods are captured by the given Markov-model.

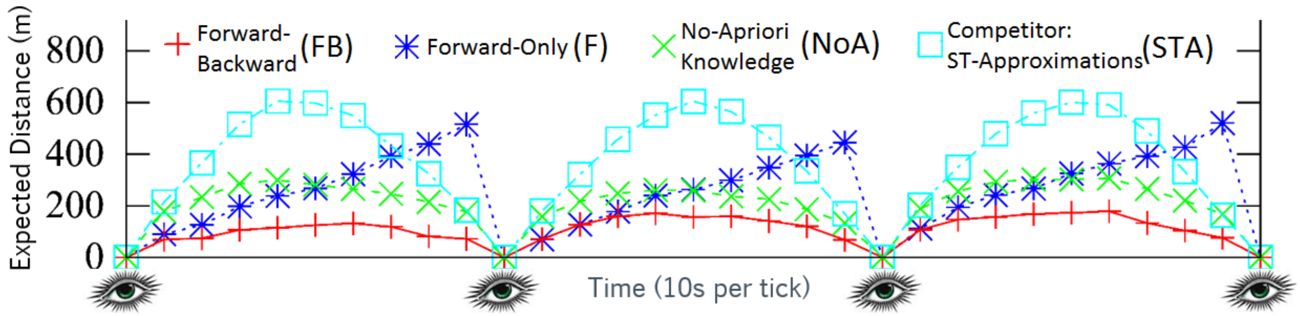

Figure 5: Effectiveness of the Model Adaption

Figure 5 shows the mean error of these approaches, computed during each point of time, evaluated over a time interval of 300 seconds using 10-second ticks for time discretization. We use the ground-truth observation of each objects at the initial time, and after every 10 ticks (100 seconds). The mean error has been computed in leave-one-out manner, i.e. trajectories for computing the error have not been used to train the model in order to avoid overfitting. The figure visualizes the error of the the model adapted by the full **forward-backward-adapted a-priori-model (FB)** and the approach using the **forward-model adaption only (F)**. We further implemented two additional approaches:

Spatio-Temporal approximations (STA), a competitor corresponding to [18, 16], which identifies a conservative approximations of possible locations of an object at each point of time. Due to a lack of better knowledge, this approach assumes all reachable states at a given time to have a uniform probability. The difference to the cylinders and beads approximation models presented in [18, 16] is that these models use conservative approximations that may include some (time, state) pairs actually having a zero probability for an object to be located at. Thus, our competitor approach here, using "perfect approximation", is at least as good as the cylinders and beads approximation models in terms of effectiveness, regardless of the approximation type used.

No apriori-knowledge (NoA), an approach that is equivalent to our forward-backward approach, however turning probabilities in the apriori transition matrix are equally distributed instead of learning the exact transition probabilities from the underlying map data. This corresponds to the case where no trajectory data is available to learn traffic patterns from.

We see that the forward-only approach (F) yields high error, especially directly before an observation. This problem is solved by the forward-backward approach (FB), by using future observations to adapt the current model. It is promising to see that, in the case where the Markov chain is assumed to be uniformly distributed (NoA), the results are still good, but worse than with the actual learned probabilities (FB). This is good news, as it shows that even a non-optimally learned Markov chain can lead to useful results, however with a slightly higher error. This good performance comes from the fact that with a uniform transition distribution the diamond-shaped space of possible time-state pairs still has high probabilities in the center of the diamond, since trajectories near the center of the bead will have a higher likelihood than trajectories close to the beads boundary. This stands in contrast to the uniform approach (STA) that models all states at the diamonds border to have the same probabilities as the states in the diamonds center; explaining why STA performs worse than NoA. Most importantly however, it is clear that our approach is capable of reducing the expected distance between the corresponding model and the ground-truth location to only 150meters for taxi cabs in Beijing, having observations only once every 100seconds. This is a drastic reduction compared to the state-of-the-art (STA), which approximates only the reachable space of an object. Without the ability to model the probability distribution on this space, this approach yields an expected error of as large as 600meters. To conclude, combining concrete observations of an individual with a global traffic patterns learned from the crowd produces the most accurate spatio-temporal uncertainty models.

References

[1] J. Aßfalg, H.-P. Kriegel, P. Kröger, and M. Renz. Probabilistic similarity search for uncertain time series. pages 435–443, 2009.

[2] Z. Chen, H. T. Shen, and X. Zhou. Discovering popular routes from trajectories. pages 900–911, 2011.

[3] T. Emrich, H.-P. Kriegel, N. Mamoulis, M. Renz, and A. Züfle. Indexing uncertain spatio-temporal data. pages 395–404, 2012.

[4] S. Karlin and H. M. Taylor. *A First Course in Stochastic Processes*, volume 2. Academic Pr Inc, 1975.

[5] B. Kuijpers and W. Othman. Trajectory databases: Data models, uncertainty and complete query languages. In *JCSS*, pages 538–560, 2010.

[6] H. Mokhtar and J. Su. Universal trajectory queries for moving object databases. In *Mobile Data Management (MDM)*, pages 133–144, 2004.

[7] J. Niedermayer, A. Züfle, T. Emrich, M. Renz, N. Mamoulis, L. Chen, and H.-P. Kriegel. Probabilistic nearest neighbor queries on uncertain moving object trajectories. In *VLDB*, 2014.

[8] A. Papoulis. *Probability, Random Variables, and Stochastic Processes, 2nd ed.* McGraw-Hill, 1984.

[9] D. Pfoser, C. S. Jensen, and Y. Theodoridis. Novel approaches to the indexing of moving object trajectories. pages 396–406, 2000.

[10] C. Ré, J. Letchner, M. Balazinksa, and D. Suciu. Event queries on correlated probabilistic streams. pages 715–728, 2008.

[11] S. Saltenis, C. S. Jensen, S. T. Leutenegger, and M. A. Lopez. Indexing the positions of continuously moving objects. pages 331–342, 2000.

[12] Y. Tao, C. Faloutsos, D. Papadias, and B. Liu. Prediction and indexing of moving objects with unknown motion patterns. pages 611–622, 2004.

[13] Y. Tao and D. Papadias. Time-parameterized queries in spatio-temporal databases. pages 334–345, 2002.

[14] Y. Tao, D. Papadias, and Q. Shen. Continuous nearest neighbor search. pages 287–298, 2002.

[15] Y. Tao, D. Papadias, and J. Sun. The tpr*tree: An optimized spatio-temporal access method for predictive queries. pages 790–801, 2003.

[16] G. Trajcevski, A. N. Choudhary, O. Wolfson, L. Ye, and G. Li. Uncertain range queries for necklaces. In *11th International Conference on Mobile Data Management (MDM 2010), Kansas City, Missouri*, pages 199–208, 2010.

[17] G. Trajcevski, R. Tamassia, H. Ding, P. Scheuermann, and I. F. Cruz. Continuous probabilistic nearest-neighbor queries for uncertain trajectories. pages 874–885, 2009.

[18] G. Trajcevski, O. Wolfson, K. Hinrichs, and S. Chamberlain. Managing uncertainty in moving objects databases. *ACM Trans. Database Syst.*, 29(3):463–507, 2004.

[19] G. Trajcevski, O. Wolfson, F. Zhang, and S. Chamberlain. The geometry of uncertainty in moving objects databases. pages 233–250, 2002.

[20] M.-Y. Yeh, K.-L. Wu, P. S. Yu, and M. Chen. Proud: A probabilistic approach to processing similarity queries over uncertain data streams. pages 684–695, 2009.

[21] J. Yuan, Y. Zheng, X. Xie, and G. Sun. Driving with knowledge from the physical world. pages 316–324, 2011.

[22] J. Yuan, Y. Zheng, C. Zhang, W. Xie, X. Xie, and Y. Huang. T-drive: Driving directions based on taxi trajectories. In *Proc. ACM GIS*, 2010.

[23] Y. Zheng and X. Zhou. *Computing with Spatial Trajectories*. Springer, 2011.

[24] A. Züfle. *Similarity search and mining in uncertain spatial and spatio-temporal databases*. PhD thesis, lmu, 2013.