

Probabilistic map-matching using particle filters

Kira Kempinska^{*1}, Toby Davies^{†1} and John Shawe-Taylor^{‡2}

¹Department of Security and Crime Science, University College London

²Department of Computer Science, University College London

January 8, 2016

Summary

Increasing availability of vehicle GPS data has created potentially transformative opportunities for traffic management, route planning and other location-based services. Critical to the utility of the data is their accuracy. Map-matching is the process of improving the accuracy by aligning GPS data with the road network. In this paper, we propose a purely probabilistic approach to map-matching based on a sequential Monte Carlo algorithm known as particle filters. The approach performs map-matching by producing a range of candidate solutions, each with an associated probability score. We outline implementation details and thoroughly validate the technique on GPS data of varied quality.

KEYWORDS: map-matching, GPS data, particle filter, probabilistic modelling

1 Introduction

Over the last years we have witnessed a rapid increase in the availability of GPS-receiving devices, such as smart phones or car navigation systems. The devices generate vast amounts of temporal positioning data that have been proven invaluable in various applications, from traffic management (Kühne et al., 2003) and route planning (Gonzalez et al., 2007; Li et al., 2011; Kowalska et al., 2015) to inferring personal movement signatures (Liao et al., 2006).

Critical to the utility of GPS data is their accuracy. The data suffer from measurement errors caused by technical limitations of GPS receivers and sampling errors caused by their receiving rates. When digital maps are available, it is common practice to improve the accuracy of the data by aligning GPS points with the road network. The process is known as map-matching.

Most map-matching algorithms align GPS trajectories with the road network by considering *positions* of each GPS point, either in isolation or in relation to other GPS points in the same trajectory.

*kira.kowalska.13@ucl.ac.uk

†toby.davies@ucl.ac.uk

‡j.shawe-taylor@cs.ucl.ac.uk

The techniques, although computationally efficient, are not very accurate in cases when the sampling rate is low or the street network complexity is high. More advanced map-matching techniques utilise both *timestamps* and *positions* of GPS points in order to achieve a higher degree of accuracy. They would typically use temporal information to infer speed and then assign GPS points to roads that are in their proximity and which speed profiles best match the inferred speed. A prominent example of a spatio-temporal algorithm is ST-Matching (Lou et al., 2009). It has been shown to outperform purely spatial map-matching approaches, especially when the sampling rate is low.

The major limitation of both spatial and spatio-temporal approaches is their deterministic nature. They would always snap a GPS trajectory to a road network, regardless if it even came from the road network in the first place. The lack of confidence scores associated with their outputs might lead to very misleading results, especially when the data quality is low.

In this paper, we address the issue of certainty by proposing a purely probabilistic spatio-temporal map-matching approach. It is based on a sequential Monte Carlo algorithm known as particle filters. The algorithm originates from the field of robotics (Thrun, 2002), where it has been widely applied in robot localisation problems. In the context of map-matching, it uses both spatial and temporal information to iteratively align a GPS trajectory with the road network; hence it can be used for both tracking and offline map-matching. It outputs the most likely road sequence that the GPS data came from together with the associated likelihood.

2 Problem Statement

In this section, we define the problem of probabilistic map-matching.

Definition 1 (*GPS trajectory*): A sequence of GPS points, where each GPS point contains latitude, longitude, bearing and timestamp.

Definition 2 (*Road network*): A directed graph with vertices representing road intersections and edges representing road segments. Bidirectional road segments are represented by two edges, each corresponding to a single direction of flow. Roads and intersections can be uniquely identified using their IDs.

Definition 3 (*Path*): A connected sequences of street segments in the road network.

Given a road network and a GPS trajectory, the goal of probabilistic map-matching is to find most probable paths that the GPS trajectory was generated from, together with their associated probability values.

3 Methodology

Our map-matching framework is based on particle filters. The algorithm computes candidate paths and their probabilistic values given a GPS trajectory. The most probable candidate path can then



Figure 1: Exemplary road network with a GPS trajectory to be map-matched.

be selected as the map-matching outcome. The framework is evaluated using cross-validation.

3.1 Particle Filter

Particle filter is a sequential Monte Carlo technique that approximately infers true states of a dynamical system given its noisy observations. In our case, the dynamical system is a vehicle following a path along the road network, noisy observations are GPS points and the true states that we want to infer are actual locations of the vehicle at different timestamps.

The algorithm is based on the assumption that the dynamical system can be modeled as a first-order Markov chain with unobserved (hidden) states (see Figure 2). That is, it assumes that the state of the system x_t at time t solely depends on the state at time $t - 1$ through the so-called *transition probability* $p(x_t | x_{t-1}, u_t)$, where u_t is the control giving information about the change of the system in the time interval $(t - 1; t]$. It adds that any measurements of the system are noisy descriptions of the unobserved true states, where the noise is modelled by the *measurement probability* $p(y_t | x_t)$. The goal of particle filters is to infer x_t given all available measurements $y_{1:t}$. The algorithm approximates the solution by recursively sampling from the posterior distribution (Bishop, 2006):

$$p(x_t | y^t, u^t) = \text{const.} \cdot p(y_t | x_t) \int p(x_t | x_{t-1}, u_t) \cdot p(x_{t-1} | y^{t-1}, u^{t-1}) \quad (1)$$

under the initial condition $p(x_0 | y^0, u^0) = p(x_0)$ where $p(x_0)$ is the so-called *initialisation distribution*. The samples are represented by *particles*, i.e. possible states of the system given measurements.

Definition 4 (*Particle*): A point on the road network containing unique road segment identifier, distance along the segment and direction of travel (defined by from-to endpoints of the segment).

The most basic version of particle filters is given by the following algorithm.




Figure 2: Graphical representation of a first-order Markov chain with hidden states $x_{1:N}$, measurements $y_{1:N}$ and controls $u_{1:N-1}$ at times $t = 1 : N$.

- **Initialisation:** At time $t = 0$, draw M particles according to $p(x_0)$. Call this set of particles X_0 .
- **Recursion:** At time $t > 0$, generate a particle x_t for each particle in X_{t-1} by sampling from the transition probability $p(x_t | x_{t-1}, u_t)$. Call the resulting set \bar{X}_t . Subsequently, draw M particles (with replacement) with a probability proportional to the measurement probability $p(y_t | x_t)$. The resulting set of particles is X_t .

When the recursion reaches the last measurement at $t = N$, the particles stored in X_N are approximate samples from the desired distribution $p(x_N | y_{1:N}, u_{2:N})$. In our context, they represent possible paths taken by a vehicle given the GPS trajectory. The certainty associated with each path is proportional to the fraction of particles that it is represented by.

3.2 Method Validation

The easiest way to validate the accuracy of our map-matching approach would be to compare predicted paths with actual paths taken by a vehicle. Unfortunately, the ground truth is not available in our case study and we need validation techniques that overcome this limitation.

We propose a validation framework based on the well-established technique of cross-validation (Barber, 2012). We remove 10% of GPS points from each available GPS trajectory (see Figure 3). We then align the incomplete trajectories with the road network. Finally, we measure the distance between each removed point and the corresponding aligned path. The average distance across all removed points is our estimate of map-matching error.

4 Results

4.1 Data

The data motivating the project is a complete GPS trajectory of a police patrol vehicle during its night shift (9pm to 7am) in the London Borough of Camden on February 9th 2015. The dataset



Figure 3: Exemplary GPS trajectory with points split into training and test sets.

contains 4,800 GPS points that were emitted roughly every second when moving. It was acquired for research purposes as part of the "Crime, Policing and Citizenship" project.¹

4.2 Implementation

A Initialisation

The initialisation probability distribution $p(x_0)$ is defined as a Gaussian centred at the position and bearing of the first GPS point. Particles initialised from the distribution are required to be positioned on the road network, hence their positions are first sampled (see Figure 4a) and then either kept or discarded depending on whether they coincide with the road network or not (see Figure 4b). Their direction of travel is inferred from the sampled bearing.

B Transition probability

The transition probability $p(x_t | x_{t-1}, u_t)$ is set as a linear estimate equal to the Cartesian distance between GPS points x_{t-1} and x_t (the control u_t) plus an additive Gaussian noise.

In the recursive step of particle filter, particles move along the road network by a distance sampled from $p(x_t | x_{t-1}, u_t)$. When they encounter a road intersection, they randomly choose which road to follow.

C Measurement probability

¹UCL Crime Policing and Citizenship: <http://www.ucl.ac.uk/cpc/>.




Figure 4: Initialisation of particles around the first GPS point in a trajectory.

Finally, the measurement noise $p(y_t | x_t)$ is also modelled as a Gaussian distribution, i.e. it is expected that GPS points are normally distributed around the true vehicle locations.

4.3 Performance Evaluation

In the first instance, the proposed algorithm is applied to the police vehicle data. An exemplary output of the algorithm is shown in Figure 5. The median cross-validation error is 4.9 meters, i.e. the inferred paths tend to be 4.9 meters away from GPS points not included in the map-matching. The error approximately equals the measurement noise of the GPS data themselves, therefore the results seem to be accurate.

The applicability of the algorithm to other datasets is then tested by artificially reducing the sampling rate of the data (removing some GPS points) and by increasing the noise of the data (perturbing GPS points). The algorithm shows good robustness against variation of the measurement noise (Figure 6a) that might in reality be due to high buildings, weather, etc.. However, it performs poorly on datasets with low sampling rates (Figure 6b). The decreased performance can be explained by the fact that low sampling rates largely increase the number of possible paths that the vehicle could have taken between subsequent GPS measurements (too many to cover with a fixed number of particles). The decrease in the algorithm’s performance is particularly apparent when compared to the relatively good performance of the state-of-the-art deterministic approach, the ST-Matching algorithm.

Further work is already being undertaken to bring together strengths of the two algorithms into a highly accurate, yet fully probabilistic, map-matching algorithm.

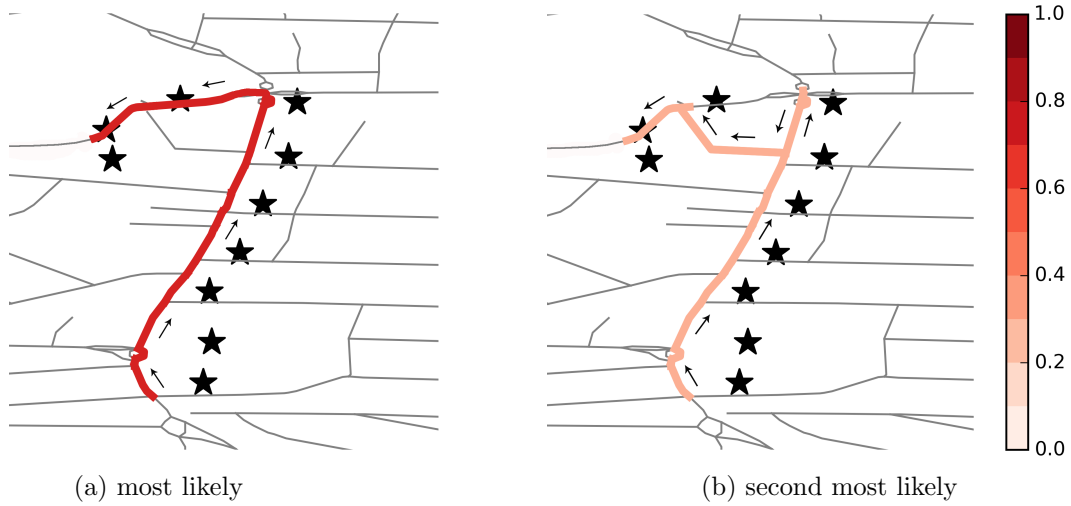


Figure 5: Exemplary map-matching outcome with colour-coded probability scores for the two most probable paths.

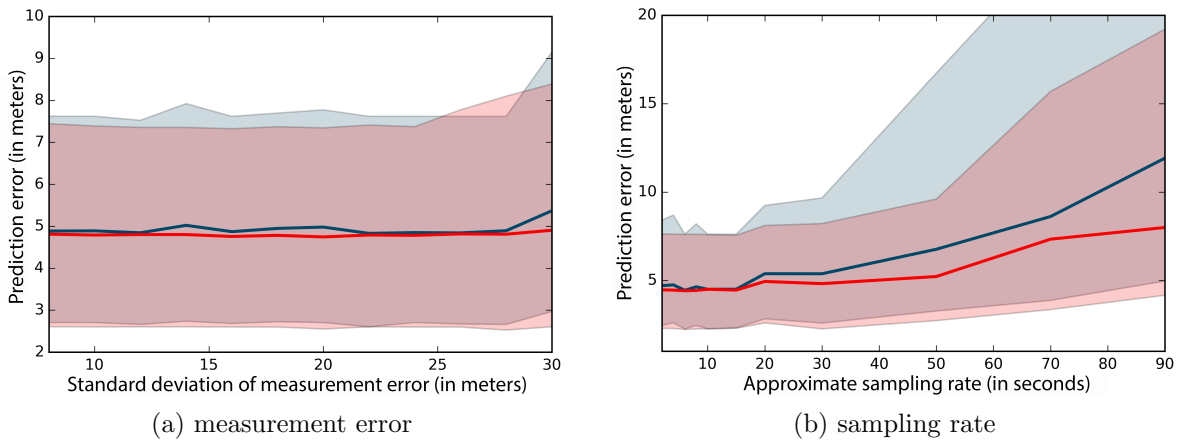


Figure 6: Sensitivity of Particle Filter (blue) and ST-Matching (red) to GPS measurement error and sampling rate represented as 25th, 50th and 75th percentiles of map-matching errors.

5 Acknowledgements

This work is part of the project - Crime, Policing and Citizenship (CPC): Space-Time Interactions of Dynamic Networks (www.ucl.ac.uk/cpc), supported by the UK Engineering and Physical Sciences Research Council (EP/J004197/1). The data provided by Metropolitan Police Service (London) is greatly appreciated.

6 Biography

Kira Kempinska is a PhD student in the Jill Dando Institute of Crime and Security Sciences at University College London. Her main research interests lie in the area of probabilistic machine learning and network analysis, particularly in application to crime and security issues.

Toby Davies is a Research Associate working on the Crime, Policing and Citizenship (CPC) project at UCL. His background is in mathematics, and his work concerns the application of mathematical techniques in the analysis and modelling of crime. His research interest include networks and the analysis of spatio-temporal patterns.

John Shawe-Taylor is a professor at University College London (UK) where he is the Head of the Department of Computer Science. His main research area is Statistical Learning Theory, but his contributions range from Neural Networks, to Machine Learning, to Graph Theory.

References

- Barber, D. (2012). *Bayesian Reasoning and Machine Learning*. Cambridge University Press.
- Bishop, C. M. (2006). *Pattern recognition and machine learning*. Springer.
- Gonzalez, H., Han, J., Li, X., Myslinska, M., and Sondag, J. P. (2007). Adaptive fastest path computation on a road network: a traffic mining approach. In *Proceedings of the 33rd International Conference on Very Large Data Bases*, pages 794–805. VLDB Endowment.
- Kowalska, K., Shawe-Taylor, J., and Longley, P. (2015). Data-driven modelling of police route choice. In *Proceedings of the 23rd GIS Research UK conference*.
- Kühne, R., Schäfer, R.-P., Mikat, J., and Lorkowski, S. (2003). New Approaches for Traffic Management in Metropolitan Areas. In *Proceedings of the 10th Symposium on Control in Transportation Systems*, Tokyo.
- Li, Q., Zeng, Z., Zhang, T., Li, J., and Wu, Z. (2011). Path-finding through flexible hierarchical road networks: An experiential approach using taxi trajectory data. *International Journal of Applied Earth Observation and Geoinformation*, 13(1):110–119.
- Liao, L., Patterson, D. J., Fox, D., and Kautz, H. (2006). Building personal maps from GPS data. *Annals of the New York Academy of Sciences*, 1093:249–65.

- Lou, Y., Zhang, C., Zheng, Y., Xie, X., Wang, W., and Huang, Y. (2009). Map-matching for low-sampling-rate GPS trajectories. In *Proceedings of the 17th ACM SIGSPATIAL International Conference on Advances in Geographic Information Systems - GIS '09*, page 352, New York, New York, USA. ACM Press.
- Thrun, S. (2002). Particle filters in robotics. In *Proceedings of Uncertainty in AI*, pages 511–518. Morgan Kaufmann Publishers Inc.