Improved Particle Filters for Vehicle Localisation

Kira Kempinska
Department of Security and Crime Science
University College London
London, WC1E 6BT
kira.kowalska.13@ucl.ac.uk

John Shawe-Taylor
Department of Computer Science
University College London
London, WC1E 6BT
j.shawe-taylor@ucl.ac.uk

Abstract

The ability to track a moving vehicle is of crucial importance in numerous applications. The task has been often approached by the importance sampling technique of particle filters due to its ability to model non-linear and non-Gaussian dynamics, of which a vehicle travelling on a road network is a good example. The performance of the particle filters method is strongly dependent on the choice of the proposal distribution. In this paper, we introduce a proposal distribution that samples around the most recent sensor observation. The proposal leads to an order of magnitude improvement in accuracy and efficiency over conventional particle filters, especially when the sensor noise is low.

1 Introduction

The problem of tracking the position of a moving vehicle has long been studied in the field of robotics. It requires estimating the true position of a vehicle in real time given its noisy sensor measurements and, optionally, the map of its environment. The problem has been successfully solved by particle filters, which, in contrast to the widely used Kalman filters, are able to approximate a wide range of non-linear and non-Gaussian distributions. It is an importance sampling technique which approximates the target distribution by sampling from a series of intermediate proposal distributions.

Critically, in common with any important sampling method, the performance of particle filters is strongly dependent on the choice of the proposal distribution. If the proposal distribution does not closely match the target distribution, the algorithm fails to converge or converges at a very slow rate. This problem is relatively common in vehicle localisation due to high complexity of the underlying road network that makes any predictions about current position from historical data difficult.

To alleviate the problem, several authors modified the proposal distribution to sample according to the most recent sensor measurement [5, 2], or a mixture of the most recent and historic measurements [7, 3, 6] or even future observations [4]. The modifications successfully reduced the divergence between the proposal and the posterior distributions, however, at the expense of relatively complex sampling schemes, including two-stage sampling and kd-trees for density estimation [2] or permutations at each timestamp [5]. The construction of good, but also computationally efficient proposal distributions is still an open research question.

In this paper, we propose an improved particle sampling scheme that is both computationally efficient and mathematically robust. The approach generates proposals based on the current sensor observation only, leading to good alignment between the proposal and the target distribution even with a small sample size. It converges to the desired target distribution at faster rates than standard particle filters, especially when observations are infrequent but low-noise. It is easy to implement and avoids the computational and analytical complexity of the previously discussed alternatives, including those using the same proposal distribution [5, 2].

30th Conference on Neural Information Processing Systems (NIPS 2016), Barcelona, Spain.

2 Problem statement

The key idea of particle filters is to estimate the marginal posterior distribution $p(x_t \mid z_{0:t})$ where x_t is the state of the system at time t and $z_{0:t} = \{z_0, \dots, z_t\}$ is a sequence of measurements collected up to time step t. We call the posterior the *belief* and use the following notation

$$Bel(x_t) = p(x_t \mid z_{0:t}) \tag{1}$$

In the context of vehicle tracking, the belief is our estimate of the vehicle position at time t given all measurements collected until then. Particle filters estimate $Bel(x_t)$ recursively. The underlying assumption of particle filters is that the system follows the *Markov assumption*, that is, measurements z_t are conditionally independent of past measurements given knowledge of the state x_t. We can use the Markov assumption and Bayes rule to arrive at a recursive form of $Bel(x_t)$:

$$Bel(x_t) = \frac{p(z_t \mid x_t)p(x_t \mid x_{t-1})Bel(x_{t-1})}{p(z_t \mid z_{0:t-1})} \tag{2}$$

The recursive equation is the basis for particle filters and the improved particle filters that we propose in this paper.

3 Particle Filters

The particle filters method is an importance sampling scheme. It estimates the belief $Bel(x)$ by a set of m weighted samples distributed according to $Bel(x)$:

$$Bel(x) = \{x^{(i)}, w^{(i)}\}_{i=1,\dots,m}$$

where each $x^{(i)}$ is a sample (a state) and $w^{(i)}$ are non-negative weights called *importance factors* that determine the importance of each sample.

Particle filters generate samples $x^{(i)}$ from a proposal distribution given by

$$Q = p(x_t \mid x_{t-1})Bel(x_{t-1}) \tag{3}$$

Consequently, the importance factors are calculated from the quotient

$$\begin{aligned} \frac{Bel}{Q} &= [p(x_t \mid x_{t-1})Bel(x_{t-1})]^{-1} \frac{p(z_t \mid x_t)p(x_t \mid x_{t-1})Bel(x_{t-1})}{p(z_t \mid z_{0:t-1})} \\ &\propto p(z_t \mid x_t) \end{aligned} \tag{4}$$

4 Improved Particle Filters

We propose an improved particle sampling scheme in which x_t are sampled directly around the most recent observation z_t according to the proposal distribution:

$$Q_{new} = \frac{p(z_t \mid x_t)}{\pi(z_t)} \qquad \text{with} \qquad \pi(z_t) = \int p(z_t \mid x_t) dx_t \tag{5}$$

This new proposal distribution possesses orthogonal strengths to the one in Equation 3, in that it generates samples that are highly consistent with the most recent sensor measurement but ignorant of past measurements and controls.

The importance factors for these samples are again calculated by the quotient:

$$\frac{Bel}{Q_{new}} = \left[\frac{p(z_t \mid x_t)}{\pi(z_t)}\right]^{-1} \frac{p(z_t \mid x_t)p(x_t \mid x_{t-1})Bel(x_{t-1})}{p(z_t \mid z_{0:t-1})}$$

$$= \frac{p(x_t \mid x_{t-1})Bel(x_{t-1})\pi(z_t)}{p(z_t \mid z_{0:t-1})} \quad (6)$$

$$\propto p(x_t \mid x_{t-1})Bel(x_{t-1})$$

Since $Bel(x_{t-1})$ is represented by a set of samples $x_{t-1}^{(i)}$ weighted by importance factors $w_{t-1}^{(i)}$, the (non-normalised) importance factor for any sample $x_t^{(j)}$ can be approximated by

$$\sum_{i=1}^{m} p(x_t^{(j)} \mid x_{t-1}^{(i)}) w_{t-1}^{(i)} \quad (7)$$

The proposed sampling scheme changes how measurements are used in particle filters: the current measurement is now used for sampling (instead of weighing); past measurements are used for calculating importance factors (instead of sampling).

5 Application to Vehicle Tracking

5.1 Data

We tested the improved particle filters method on a GPS trajectory of a police patrol vehicle during its night shift (9am to 7am) in the London Borough of Camden on February 9^{th} 2015. The dataset 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].

5.2 Results

A series of tests was conducted to elucidate the difference between the standard and the proposed particle filters. We found that the modified proposal distribution consistently outperforms conventional particle filters in terms of accuracy. Largest gains in accuracy are observed on datasets with long sampling intervals, e.g. one minute and more. Figure 1 plots the prediction error (in meters) of both algorithms for different sampling intervals and levels of sensor noise, using $m = 10$ samples only. It shows that the proposed method has lower *median* error across all examined sampling rates and sensor noise levels, as well as much lower error *variation*.

We tested the sensitivity of the proposed method to the number of samples used. Figure 2 shows comparative results on GPS data with the sampling interval of 70 seconds. The proposed method yields significantly better results, both in terms of accuracy and robustness to failure. Failure means that a method loses track of the position of the vehicle, i.e. all positions that it proposes are completely unlikely given sensor data. When only $m = 10$ samples are used, the proposed method reduces the estimation error by almost 10 meters and the percentage of failure by as much as 68%. The performance is further improved when more samples are used, but the gain is small compared to the conventional particle filters. Therefore, it can be reliably used with a small number of samples, making it highly computationally efficient.

6 Conclusions

This paper describes a modified particle filters method that shows uniformly superior accuracy and robustness to failure to the conventional particle filters. The improved algorithm utilizes a different proposal distribution which uses only the most recent observation in the position prediction process.

[1] UCL Crime Policing and Citizenship: `http://www.ucl.ac.uk/cpc/`

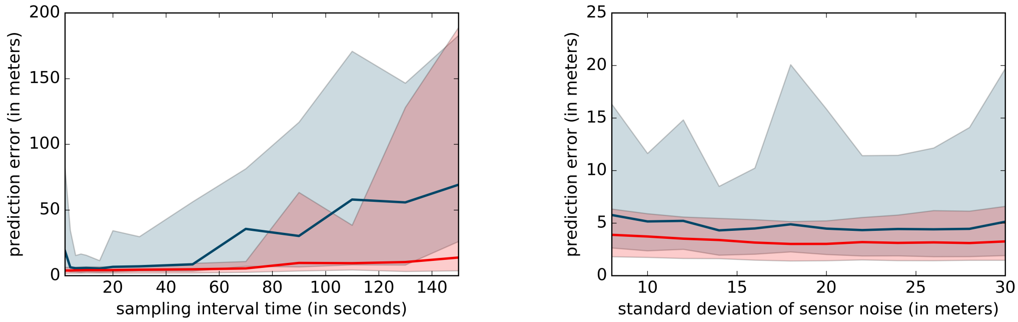

(a) sampling rate

(b) sensor noise

Figure 1: Accuracy of the improved particle filters (red) and the standard particle filters (blue) on GPS data with varied sampling rate and sensor noise, represented as 25th, 50th and 75th percentiles of prediction errors.

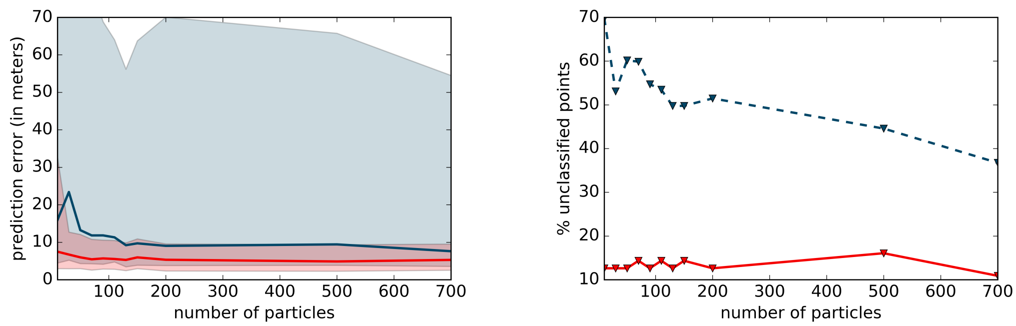

(a) accuracy

(b) failure rate

Figure 2: Accuracy and robustness to failure of the improved particle filters (red) and the conventional particle filters (blue) as a function of the number of samples used. Accuracy is shown as the 25th, 50th and 75th percentiles of prediction errors.

In doing so, it makes more efficient use of the particles, particularly in situations in which the sensor measurements are infrequent but lo-noise. We believe that our results illustrate that particle filters can be radically improved if one carefully chooses a proposal distribution, such that it extracts the most information from the available data.

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.

References

[1] Barber, D. (2012). *Bayesian Reasoning and Machine Learning*. Cambridge University Press.

[2] Fox, D., Thrun, S., Burgard, W., and Dellaert, F. (2001). Particle Filters for Mobile Robot Localization. In *Sequential Monte Carlo Methods in Practice*, pages 401–428. Springer New York, New York, NY.

[3] Kong, A., Liu, J. S., and Hung Wong, W. (1994). Sequential Imputations and Bayesian Missing Data Problems. *Journal of the American Statistical Association*, 89(425):278–288.

[4] Lin, M., Chen, R., and Liu, J. S. (2013). Lookahead Strategies for Sequential Monte Carlo. *Statistical Science*, 28(1):69–94.

[5] Lin, M. T., Zhang, J. L., Cheng, Q., and Chen, R. (2005). Independent Particle Filters. *Journal of the American Statistical Association*, 100(472):1412–1421.

[6] Liu, J. S. and Chen, R. (1998). Sequential Monte Carlo Methods for Dynamic Systems. *Journal of the American Statistical Association*, 93(443):1032.

[7] Montemerlo, M. and Thrun, S. (2007). FastSLAM 2.0. In *FastSLAM*, pages 63–90. Springer Berlin Heidelberg, Berlin, Heidelberg.