Improved Particle Filters for Vehicle Localisation

Kira Kempinska, Prof John Shawe-Taylor

kira.kowalska.13@ucl.ac.uk

Overview

Vehicle tracking has often been approached by the importance sampling technique of particle filters. The technique is able to model non-linear and non-Gaussian dynamics, of which a vehicle travelling on a road network is a good example. However, the performance of particle filters is significantly decreased when sensor observations are highly informative (e.g. infrequent but low-noise).

We introduce improved particle filters that sample around the most recent sensor observation. The proposed method shows an order of magnitude improvement in accuracy and efficiency over conventional particle filters, especially when sensor observations are infrequent.

Problem Statement

The key idea of vehicle tracking using particle filters is to estimate the belief

$$Bel(x_t) = p(x_t|z_{0:t})$$

where x_t is the vehicle position at time t and $z_{0:t}$ is a sequence of sensor measurements collected up to time step t. The belief represents our estimate of the vehicle position at time t given all measurements collected until then.

Particle filters solve $Bel(x_t)$ recursively, using the form:

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

Results

A series of tests was conducted to elucidate the difference between the standard and the proposed "improved" particle filters. We tested both methods on a GPS trajectory emitted by a police patrol vehicle over a 9-hour patrol shift on February 9th 2015 in the London Borough of Camden.

We found that the proposed method consistently outperforms conventional particle filters in terms of tracking accuracy in scenarios with varied sensor sampling rates (Fig 1a), varied sensor errors (Fig 1b) and varied numbers of proposal samples used (Fig 1c). As expected, largest gains in accuracy are observed on datasets with highly informative observations, e.g. long sensor sampling intervals of one minute and more (see Fig 1a). Overall, the improved particle filter leads to lower *median* tracking error as well as much lower error *variation*.

Particle Filters

Particle filter is an importance sampling scheme. It approximates the belief $Bel(x_t)$ by a set of m weighted samples

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

where each $x^{(i)}$ is a sample (a vehicle position) and $w^{(i)}$ are weights that determine the importance of each sample.

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

$$Q(x_t) = p(x_t|x_{t-1}, u_{t-1})Bel(x_{t-1})$$

Consequently, the importance factors are calculated from the quotient

$$\frac{Bel(x_t)}{Q(x_t)} \propto p(z_t|x_t)$$

Improved Particle Filters

We propose an improved particle sampling scheme in which $x^{(i)}$ are sampled directly around the most recent observation z_t according to the proposal distribution:

$$Q_{new}(x_t) = \frac{p(z_t|x_t)}{\pi(z_t)} \quad \text{with} \quad \pi(z_t) = \int p(z_t|x_t)dx_t$$

The importance factors are then calculated by the quotient

$$\frac{Bel(x_t)}{Q_{new}(x_t)} \propto p(x_t|x_{t-1})Bel(x_{t-1})$$

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

$$\sum_{i=1}^{m} p(x_t^{(j)}|x_{t-1}^{(i)})\, w_{t-1}^{(i)}$$

The proposed method possesses orthogonal strength to the conventional particle filters, in that it generates samples that are highly consistent with the most recent sensor measurement but ignorant of the past. As such, it should outperform particle filters when observations are very informative.

Future Work

Future research will explore the idea of using independent learning modules (e.g. neural networks) to learn even better representations of the proposal distribution from available sensor data.

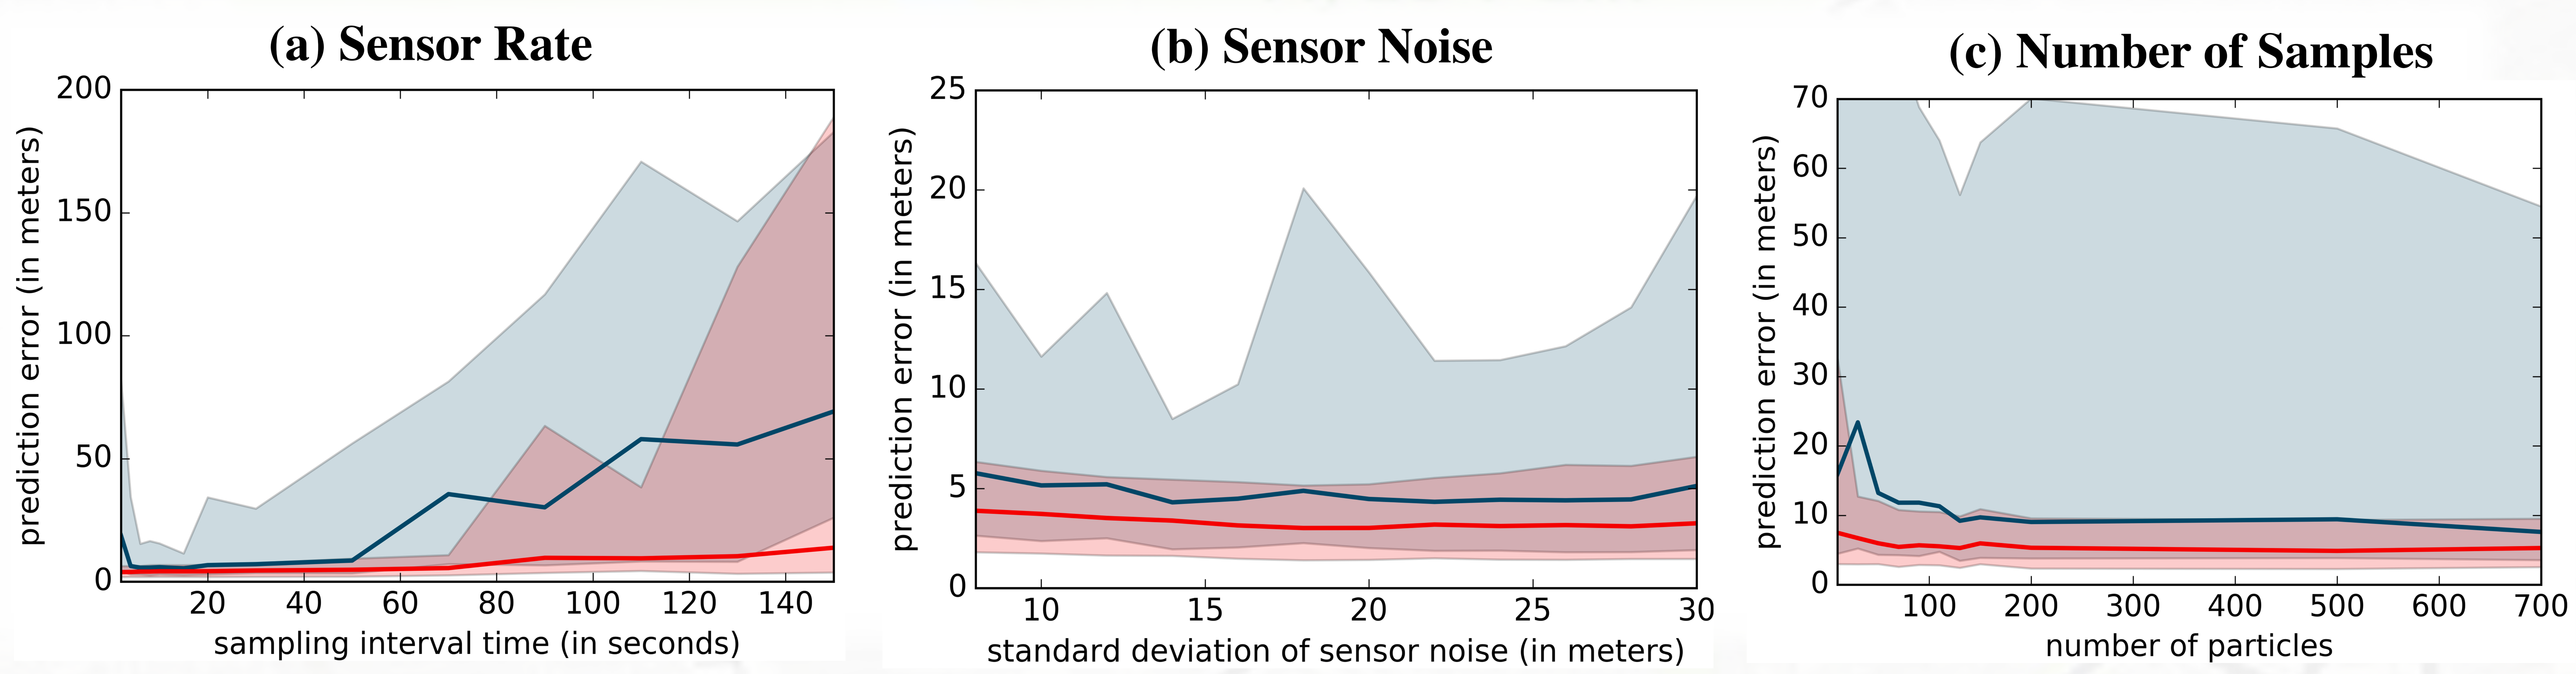

Figure 1. Accuracy of the improved particle filters (red) and the standard particle filters (blue) on GPS data with varied sensor sampling rate (a), varied sensor noise (b) and varied number of proposal samples used (c), represented as 25th, 50th and 75th percentiles of prediction errors.