Bayesian Optimisation for Parameter Inference

Johan Dahlin[†] and Fredrik Lindsten[†]

AUTOMATIC CONTROL
REGLERTEKNIK
LINKÖPINGS UNIVERSITET

Summary

- We propose a new derivative-free algorithm based on Bayesian optimisation and particle filters.
- Enables estimation of parameters in general state space models.
- Parameter estimates close to the true values are obtained using only 150 samples from the log-likelihood.

Frequentistic parameter inference

We are interested in solving the **parameter inference** problem in **nonlinear state space models**

$$x_{t+1}|x_t \sim f_\theta(x_{t+1}|x_t),$$
$$y_t|x_t \sim h_\theta(y_t|x_t),$$

given a set of observations $\mathcal{D}_T = \{y_t\}_{t=1}^T$ and where $\theta \in \Theta \subseteq \mathbb{R}^d$ denotes static parameters. The **maximum likelihood estimate** is given by

$$\widehat{\theta}_{\mathrm{ML}} = \operatorname*{argmax}_\theta \log p(\mathcal{D}_T|\theta),$$

where $\log p(\mathcal{D}_T|\theta)$ denotes the (often) intractable log-likelihood function.

Particle Bayesian optimisation

In Bayesian optimisation, we iteratively optimise a **surrogate function** $f(\theta)$ modelled as a Gaussian process by a three step procedure.

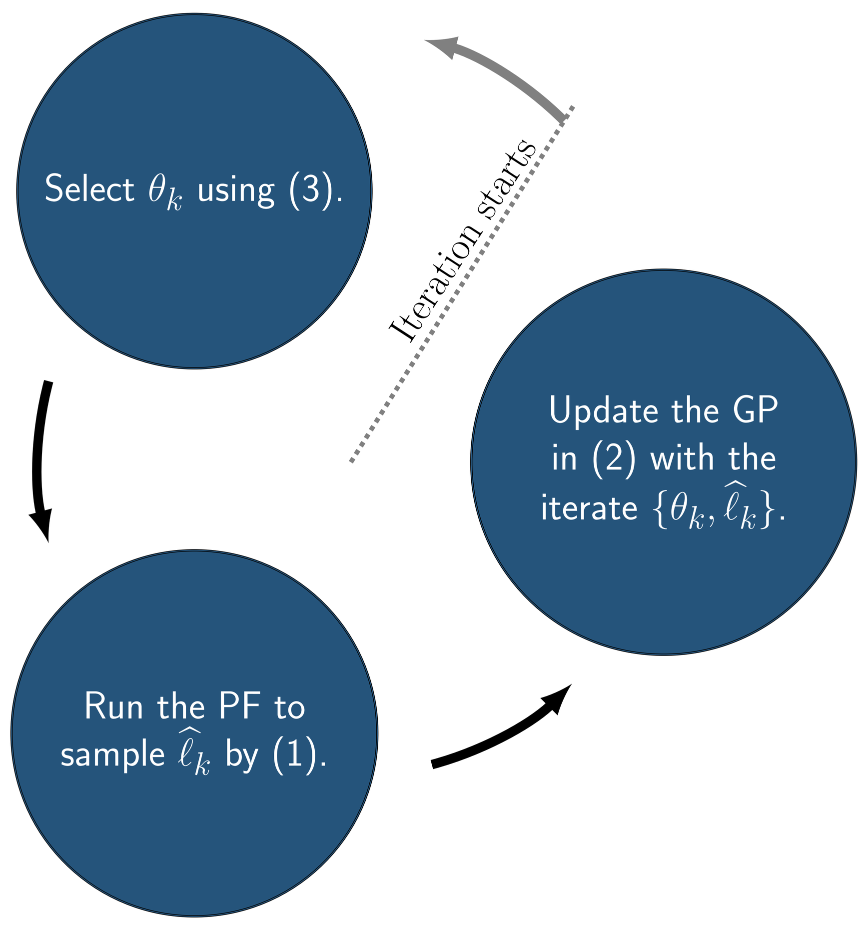

Figure: An iteration of the Particle Bayesian optimisation algorithm.

Main idea

Explore the likelihood landscape using a combination of particle filtering and Gaussian process models. New parameters are sampled according to the expected improvement of the model.

Estimating the likelihood

We run a **particle filter** (PF) targeting $p_{\theta_k}(x_t|\mathcal{D}_t)$, which returns the unnormalised particle system $\{x_t^{(i)}, w_{t|t}^{(i)}\}_{i=1}^N$. The log-likelihood can then be estimated using

$$\widehat{\ell}_k = \log \widehat{p}(\mathcal{D}_t|\theta_k) = \sum_{t=1}^{T} \log \left[\sum_{i=1}^{N} w_{t|t}^{(i)} \right]. \tag{1}$$

The resulting pair $\{\theta_k, \widehat{\ell}_k\}_{k=1}^m$ denotes the iterates of the algorithm.

Gaussian process model

The surrogate function in this optimisation is modelled as

$$f(\theta) \sim \mathcal{GP}\left(m(\theta), k(\theta, \theta')\right), \tag{2}$$

with a constant mean function, $m(\theta)$, and the Matérn covariance function, $k(\theta, \theta')$, with $\nu = 3/2$. The mean and variance of the model

$$\mu(\theta) = \mathbb{E}[f(\theta)|\{\theta_k, \widehat{\ell}_k\}_{k=1}^m],$$
$$\sigma^2(\theta) = \mathbb{V}[f(\theta)|\{\theta_k, \widehat{\ell}_k\}_{k=1}^m],$$

are updated recursively using standard results.

Acquisition rule

The next point in which to sample $p(\mathcal{D}_T|\theta)$ is determined by the maximising argument of **the expected improvement** defined as

$$\mathbb{EI}(\theta) = \left[\mu(\theta) - \max_\theta \mu(\theta) - \xi\right] \Phi(Z) + \sigma(\theta)\phi(Z), \text{ with} \tag{3}$$
$$Z = \frac{1}{\sigma(\theta)} \left[\mu(\theta) - \max_\theta \mu(\theta) - \xi\right],$$

where ξ denotes a coefficient that balances exploration and exploitation. Here, Φ and ϕ denote the CDF and PDF of the Gaussian distribution.

Example: Linear Gaussian model

$$x_{t+1}|x_t \sim \mathcal{N}\left(x_{t+1}; \theta_1 x_t, \theta_2^2\right),$$
$$y_t|x_t \sim \mathcal{N}\left(y_t; x_t, 0.1^2\right),$$

with true parameters $\theta^\star = \{\theta_1^\star, \theta_2^\star\} = \{0.5, 1.0\}$. We use $T = 2\,000$ time steps, $N = 2\,000$ particles and $M = 150$ iterations.

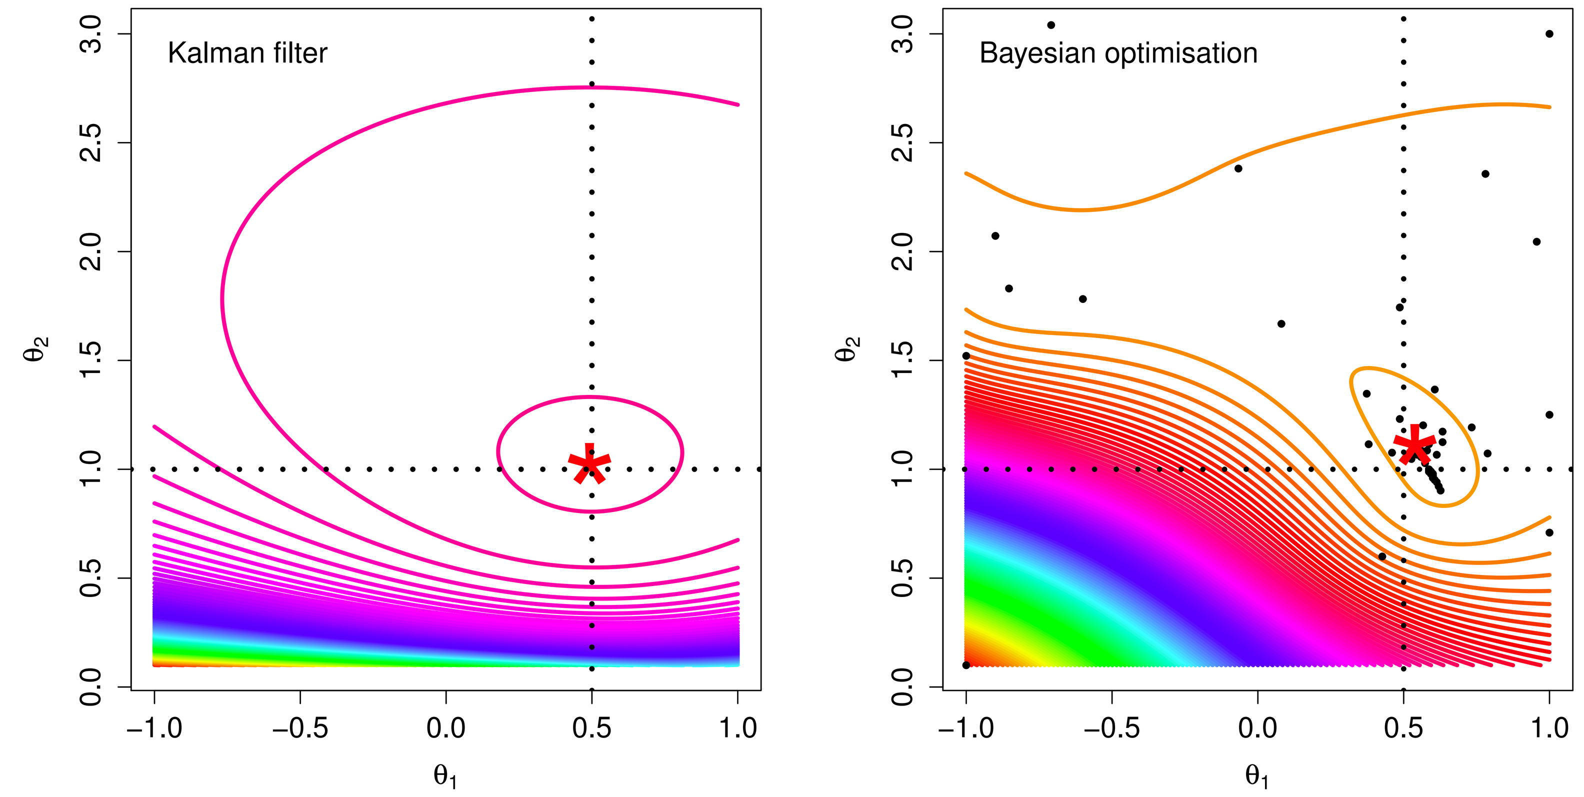

Example: Stochastic volatility model

$$x_{t+1}|x_t \sim \mathcal{N}\left(x_{t+1}; \theta_1 x_t, \theta_2^2\right),$$
$$y_t|x_t \sim \mathcal{N}\left(y_t; 0, 0.65^2 \exp(x_t)\right),$$

with $\theta^\star = \{\theta_1^\star, \theta_2^\star\} = \{0.98, 0.16\}$ and the same settings as before.

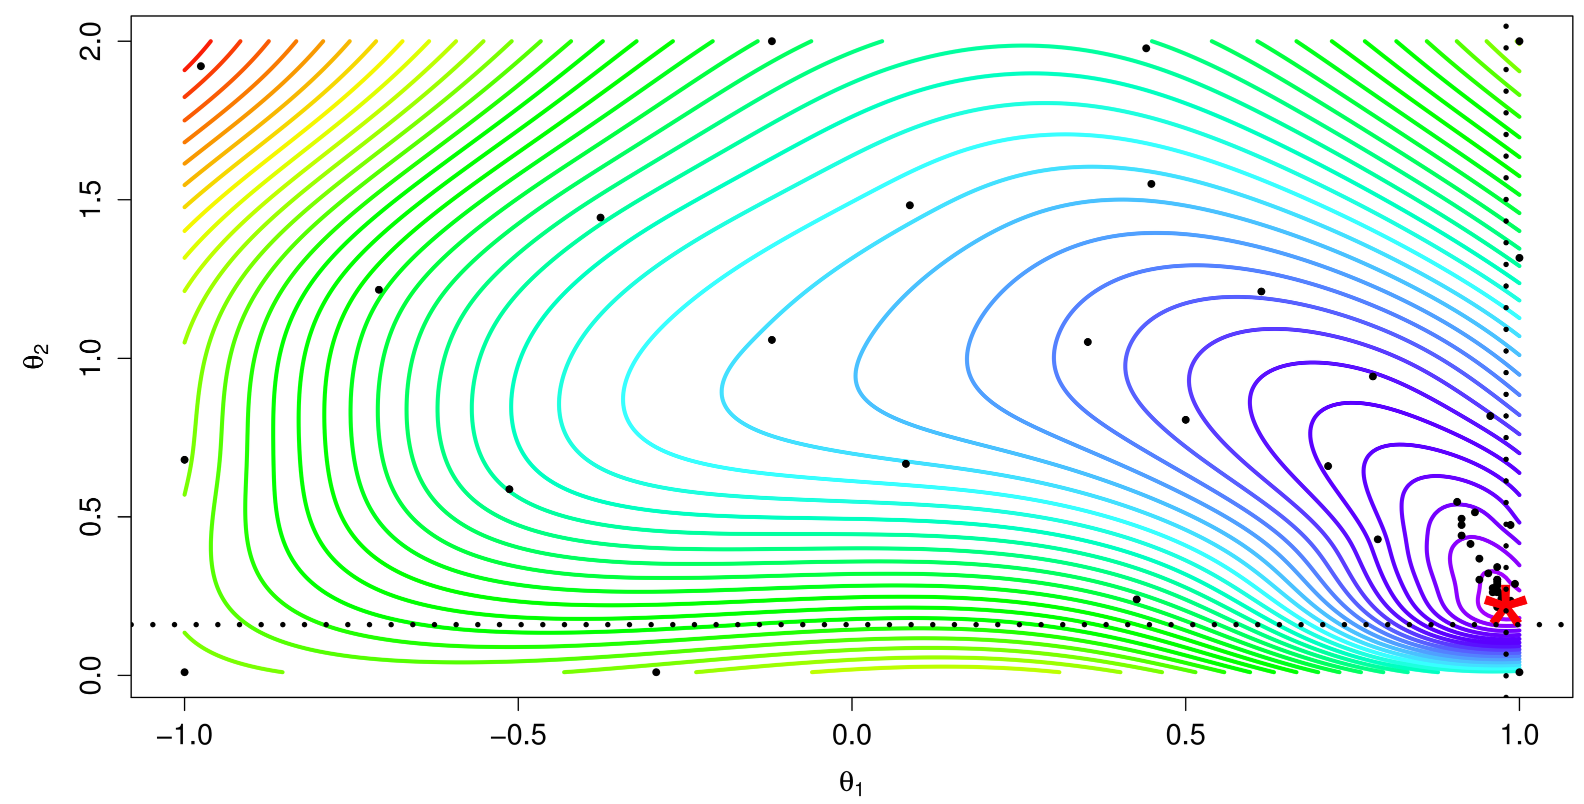

More information and source code

http://users.isy.liu.se/rt/johda87/