Estimating Bayesian state space models using simulated data meta-learning

This paper presents a simple algorithm for the estimation Bayesian state space models that is based on the principles of meta-learning literature. In contrast to previous works on simulation algorithms, we focus on high-dimensional state space estimation (not on low-dimensional model parameter estimation problems), which is much more important for practical macroeconomic modeling problems.

The algorithm consists of two main steps: artificial data generation and fitting a neural network to the variables of interest. In the first step, an artificial dataset is created as a set of samples from the joint distribution of parameters (generated from prior), states (generated from data generating process) and data (generated from data generating process). In the second step, the neural network is trained in a supervised manner to predict unobserved states (and parameters if necessary) on the previously generated dataset. The paper shows that the algorithm converges to the posterior mean or any other characteristic depending on the chosen loss function. It is shown that algorithm converges to the posterior mean or any other characteristic of posterior distribution depending on the loss function used in the training of neural network. The main advantage of the proposed method is that trained once it can be used for any dataset without additional training, so the inference of Bayesian model becomes almost instantaneous.

The properties of the algorithm are illustrated using three models are common in practice: Hodrick-Prescott filter with an estimated smoothness parameter, the stochastic volatility model, and the seasonal adjustment model. The first two models show the performance of the algorithm compared to the well-known MCMC and variation Bayes algorithms (on GDP data for various countries and GPBUSD and NYSE datasets). The third task demonstrates how problems that require a lot of effort to eliminate with exact algorithms (such as the presence of structural shifts in seasonality) can be easily solved by simulating a set of such situations at the first step of the proposed algorithm.

In all three cases, the algorithms are trained on 20 millions of artificial time series, which takes no more than 12 hours on GEFORCE RTX 2070 GPU, and then can be used on any data in less than a second (without using GPU). At the same time, algorithms show sufficient accuracy compared to MCMC and variational methods, which require from tens of minutes to days for the proposed tasks.