Robust State-Space Modelling

Supervisor: Dr. Mathieu Gerber (mathieu.gerber@bristol.ac.uk)

Background

State-space models: A state space model (SSM) is a model for a time series $(Y_t)_{t\geq 1}$ which assumes that the dynamic of the series is governed by an unobserved Markov chain $(X_t)_{t\geq 1}$. State-space models are popular in many fields, such as signal processing, robotics, economics, ecology, neuroscience and epidemiology. They also play a key role in navigation and positioning, and are for instance used within self-driving cars. In this application, X_t is the position of the car at time t and Y_t a multivariate observation containing, for instance, the GPS localisation of the car (known to have a precision of a few meters) and the distance of the car with respect to some nearby landmarks. As illustrated with this example, when using a SSM we are often interested in sequentially learning the latent variable X_t (the position of the car) from the available observations $Y_1, \ldots, Y_t$, a problem which is known as filtering.

State-space models and outliers: In practice, it is often the case that the observations are contaminated by some outliers, that is that they contain a few "extreme" data points. In the self-driving car example this problem arises due to sensor measurements faults. Outliers also commonly occur in indoor navigation using low-cost devices, such as mobile phones, where sensor measurement reliability poses a challenge. A SSM usually depends on a parameter θ that needs to be learnt from the data, and it is well-known that the presence of a few outliers can affect drastically the estimation of the model parameter. Developing parameter inference and filtering methods for SSMs which are robust to the presence of outliers is therefore crucial in many applications. However, existing robust procedures for inference in SSMs are essentially ad-hoc and/or problem specific and, somewhat surprisingly, usually focus exclusively on the filtering problem.

State-space models and misspecification: The specification of a statistical model is always a delicate task. On the one hand, the model needs to be flexible enough to capture well the true nature of the data but, on the other hand, the model has to be sufficiently simple so that it can be easily estimated and used e.g. to make predictions. The specification of a suitable model is particularly challenging in the context of SSMs for at least two reasons. Firstly, parameter inference in SSMs is a hard task from a computational point of view, preventing the use of SSMs having a large number of parameters. Secondly, in many applications where SSMs are used it is needed to make predictions with very high precision (such as in the above self-driving car example, where the position X_t of the car needs to be estimated with very high precision for obvious safety reasons), requiring the model to be nearly well-specified.

Project proposal

The first aim of this project is to propose a principled approach for parameter and state inference in SSMs in the presence of outliers. To do so, we will use a pseudo-Bayes approach which, informally speaking, amounts to performing parameter and state inference using a modified version of the original SSM. In this context, the key question that needs to be addressed is how to define the modified SSM. Notably, parameter inference based on the modified SSM must be (i) such that a single data point can only have a limited (i.e. bounded) impact on the estimated parameter value and (ii) such that we recover the true parameter value as the sample size increases if the model is well-specified. While there is a vast literature on the development of procedures satisfying (i) and (ii) for models for independent observations or for regression models, there currently exist no such methods for parameter inference in SSMs. In addition, state inference under the modified SSM must be as similar as possible to that of the original model when the data is free of outliers. Transforming the original SSM into a modified SSM satisfying the above constraints is the first main objective of the proposed project.

The second aim of this project is to develop techniques for transforming a parametric SSM in such a way that the resulting SSM is more robust to specification errors than the original model. Specifically, we will focus on situations where the practitioner believes that only some components of the model may be misspecified. For instance, in many applications (such as in epidemiology or as in the above self-driving car example) the ideal SSM is such that the dynamic of the Markov chain $(X_t)_{t\geq 1}$ is governed by a continuous time model. While the practitioner is confident in the specification of the continuous time model, the model needs to be discretized in order to define a SSM which can be deployed in practice. In this context, this part of the project may aim towards proposing a technique which can make the SSM robust to the resulting discretization errors. Similarly, in SSMs with high-dimensional latent states, it is often practically convenient to base inference on an informative, low-dimensional component of the state. In neglecting the dependence on these other components, some misspecification is induced, and so it is of interest to develop techniques which can robustly adjust for this discrepancy.