Improved unscented Kalman smoothing for stock volatility estimation

Onno Zoeter Alexander Ypma Tom Heskes

SNN, University of Nijmegen, The Netherlands

Abstract

In [1] we introduce a novel approximate inference algorithm for non-linear dynamical systems. We demonstrate the approach with an interesting inference problem in stochastic stock volatility models.

1 Stochastic volatility models

One of the most basic models for individual stocks is the geometric Brownian motion. For a stock with daily closing prices S_t this model implies that daily log returns are independently, identically and normally distributed:

$$y_t = \log \frac{S_t}{S_{t-1}} \sim \mathcal{N}(\mu, \sigma^2) \ .$$

In the above t ranges over exchange closing times and $\mathcal{N}(\mu, \sigma^2)$ denotes the normal distribution with mean μ and standard deviation σ. This standard deviation is often referred to as the stock's *volatility*. In practice the geometric Brownian motion does not seem to hold: the stock's volatility is not constant, but is auto-correlated and mean reverting.

In our experiments we use a discrete time model inspired by the model from Hull and White. With x_t the log of the variance at time t the model reads

$$x_t = a(x_{t-1} - l) + l + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, q), \quad y_t = \mathcal{N}(\mu, \exp(x_t)) \ . \tag{1}$$

All disturbances ϵ_t are assumed to be independently drawn. Figure 1 shows an artificial dataset generated from this model.

2 The one step unscented Kalman filter

Our interest is in posteriors $p(x_t|y_{1:\tau})$. For linear Gaussian models with $\tau = t$ the well-known Kalman filter recursions give exact posteriors. However due to the non-linear observation model, computation of exact posteriors in (1) is infeasible.

The current state of the art deterministic approximation method, the unscented Kalman filter, computes approximate posteriors by approximating the joint distribution $p(x_t, y_t|y_{1:t-1})$ with a Gaussian. We show that in models where x_t and y_t are uncorrelated the unscented Kalman filter effectively never updates prior

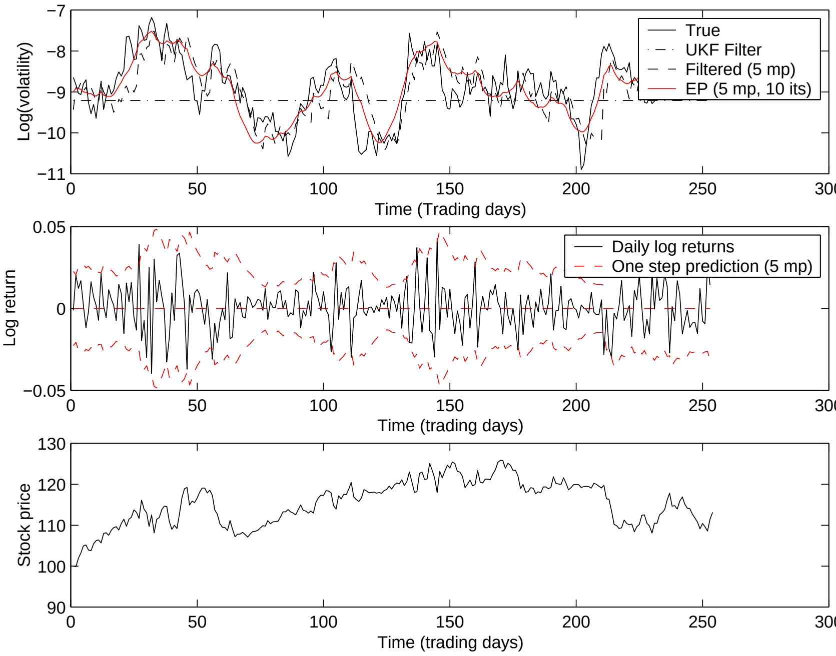

Figure 1: Results from inference on an artificial dataset. The bottom plot shows closing prices of an artificial stock. The middle plot shows the log-ratios of the closing prices (solid line) and one step ahead predictions (dashed line) and 2 standard deviations error bars of the proposed filter. The prediction is correctly constant at μ, the error bars show that the filter correctly captures the heteroskedasticity of the series. The top plot shows the artificially generated true volatilities (solid) and estimates from the unscented filter (dash-dotted), the proposed filter (dashed), and the propsed smoother (light solid). The traditional unscented filter only propagates the prior and gives nonsensical filtered results.

beliefs and thus fails on the model (1). We propose a new "one-step" filter that circumvents the linearization in the measurement update step of the unscented filter. Figure 1 shows an example of the novel filter on an artificial dataset.

The improvement in the filter also naturally leads to a symmetric smoothing pass and iteration scheme such that posteriors $p(x_t|y_{1:\tau})$ for $\tau > t$ can be computed.

References

[1] Onno Zoeter, Alexander Ypma, and Tom Heskes. Improved unscented Kalman smoothing for stock volatility estimation. In *Proceedings of the IEEE workshop on Machine Learning for Signal Processing*, 2004.