UQOP 2019
Uncertainty Quantification & Optimization Conference
18-20 March, Paris, France

A Heuristic Solution to Non-Linear Filtering with Imprecise Probabilities

Thomas Krak
Foundations Lab for Imprecise Probabilities, ELIS, Ghent University
thomas.krak@ugent.be

Cristian Greco
Aerospace Centre of Excellence, University of Strathclyde
c.greco@strath.ac.uk

Tathagata Basu
Department of Mathematical Sciences, Durham University
tathagata.basu@durham.ac.uk

Keywords: Non-Linear Filtering, State Estimation, Stochastic Process, Imprecise Probabilities

ABSTRACT

We consider the problem of estimating robust bounds on a functional of the unknown state of a dynamical system, using noisy measurements of the system's historical evolution, and given only partially specified probabilistic descriptions of the system's initial state and measurement (error) models. This problem is a general version of the well-known *filtering* problem, which finds applications throughout science and engineering. The generalisation studied in this work is based on the theory of *imprecise probabilities* [Augustin *et al.*, 2014].

Formally, we consider a dynamical system x_t, $t \geq 0$, taking values in a measurable space $(\mathcal{X}, \mathcal{A}_\mathcal{X})$, with $\mathcal{X} \subset \mathbb{R}^d$ and $d \geq 1$. The system's evolution is governed by the time-homogeneous non-linear vector differential equation

$$\mathrm{d}x_t = f(x_t)\,\mathrm{d}t\,. \tag{1}$$

With P_0, a probability measure on $(\mathcal{X}, \mathcal{A}_\mathcal{X})$, we assume that the initial state x_0 is an unknown realisation of a random variable $X_0 \sim P_0$. Together with (1), this induces random variables X_t, which describe the (uncertainty about the) state of the system at each time t. Explicitly, if for all $t \geq 0$ we define the map $F_t : \mathcal{X} \to \mathcal{X}$ as the point-wise solution to (1), $F_t(x_0) := x_0 + \int_0^t f(x_\tau)\,\mathrm{d}\tau$, then X_t is a random variable governed by the push-forward measure $P_t := P_0 \circ F_t^{-1}$ on $(\mathcal{X}, \mathcal{A}_\mathcal{X})$.

In the setting with *imprecise* probabilities, we additionally deal with (epistemic) uncertainty about the initial distribution P_0. We then say that P_0 is an element of some *set* $\mathcal{P}_0$ of probability measures on $(\mathcal{X}, \mathcal{A}_\mathcal{X})$. Notably, we do not consider any probability distribution over this set $\mathcal{P}_0$.

Finally, it is assumed that the system can only be observed through noisy measurements. The measurement y_t taken at time $t \geq 0$ is a realisation of a random variable Y_t, governed by a (stochastic) measurement model $P(Y_t \,|\, \cdot)$. We again do not assume full knowledge about $P(Y_t \,|\, \cdot)$; instead, we only know that $P(Y_t \,|\, \cdot) \in \mathcal{P}(Y_t \,|\, x_t)$ for some *set* of probability measures $\mathcal{P}(Y_t \,|\, x_t)$.

The filtering problem can now in general be described as computing the expected value of some function $h : \mathcal{X} \to \mathbb{R}$ on the uncertain state X_t, given observed measurements $y_{s_1}, \ldots, y_{s_n}$. When using imprecise probabilities, there are multiple precise probability models compatible with our set-valued assessments. The set of these compatible models is denoted $\mathcal{P}$. Any element $P \in \mathcal{P}$ is a probabilistic model, for which we can compute the quantity of interest using Bayes' rule. The aim is now to compute the *lower-* and *upper* expectations,

$$
\begin{aligned}
\underline{\mathbb{E}}\left[h(X_t)\,\middle|\,y_{s_1},\ldots,y_{s_n}\right] &:= \inf_{P\in\mathcal{P}} \mathbb{E}_P\left[h(X_t)\,\middle|\,y_{s_1},\ldots,y_{s_n}\right] \\
\overline{\mathbb{E}}\left[h(X_t)\,\middle|\,y_{s_1},\ldots,y_{s_n}\right] &:= \sup_{P\in\mathcal{P}} \mathbb{E}_P\left[h(X_t)\,\middle|\,y_{s_1},\ldots,y_{s_n}\right]
\end{aligned} \tag{2}
$$

which are the tightest possible bounds compatible with our imprecise assessments.

To the best of our knowledge, this is the first time that this filtering problem is considered at this level of generality. A non-exhaustive list of related work includes the imprecise Kalman filter [Benavoli *et al.*, 2011], which assumes the equations of motion to be linear; particle filters based on random set theory [Ristic, 2013], which deals with less general uncertainty structures; and filtering with imprecise hidden Markov chains [Krak *et al.*, 2017], which assumes $\mathcal{X}$ is finite.

In our current work, we take a heuristic approach to computing—approximately—the quantities of interest. To this end, we introduce a novel algorithm based on a problem that is equivalent to solving (2); this form is known as the *generalised Bayes' rule* in the imprecise probability literature. A known recursive decomposition thereof allows us to reduce the problem to computing (i) the n (independent) *lower-* and *upper likelihoods* of the states x_{s_i} given the observations y_{s_i}, (ii) the solution to $\inf_{P_0\in\mathcal{P}_0}\int_{\mathcal{X}} g(x)\,\mathrm{d}P_0(x)$ for any $g : \mathcal{X} \to \mathbb{R}$, and (iii) the map F_t for any t. We solve problem (i) trivially by parameterising the set $\mathcal{P}(Y\,|\,X)$ directly in terms of these lower- and upper likelihoods—a more general method would require solving the optimisation explicitly. The problem (ii) is just the lower expectation with respect to the initial model $\mathcal{P}_0$, which although it cannot really be simplified further, is relatively straightforward when compared to the original problem (2). We solve (iii) by replacing the maps F_t with surrogate model approximations that are both local in time (different surrogates for different values of t) and local in space (the approximations are only accurate on certain subsets of $\mathcal{X}$). We build these surrogates on a set of points quasi-randomly sampled from a sequential estimate of an imprecise $(1-\alpha)\times 100\%$ credible region of X_t, conditional on the sequential observations, and then propagated forward in time.

In this work, the developed method is applied to the case of robust state estimation of an Earth satellite orbiting in a strongly perturbed environment, and its performance compared against a traditional filtering approach based on precise probability distributions.

References

[Augustin *et al.*, 2014] Augustin, T., Coolen, F. P. A., De Cooman, G. and Troffaes, M. (2014). *Introduction to Imprecise Probabilities*, John Wiley & Sons.

[Benavoli *et al.*, 2011] Benavoli, A., Zaffalon, M. and Miranda, E. (2011). Robust filtering through coherent lower previsions. *IEEE Transactions on Automatic Control*, 56:1567–1581.

[Krak *et al.*, 2017] Krak, T., De Bock, J. and Siebes, A. (2017). Efficient Computation of Updated Lower Expectations for Imprecise Continuous-Time Hidden Markov Chains. *Proceedings of ISIPTA 2017*, 193–204.

[Ristic, 2013] Ristic, B. (2013). *Particle Filters for Random Set Models*, Springer.