Schedule-free variational message passing for Bayesian filtering

Wouter M. Kouw[*]

Bayesian Intelligent Autonomous Systems lab

Department of Electrical Engineering

Eindhoven University of Technology

Abstract

In Bayesian filtering, states and parameters of probabilistic state-space models are inferred in an online manner [1]. Using the Free Energy Principle [2], the state-space model is cast to a generative model p and the posterior distributions of interest are approximated using recognition distributions or beliefs q. The factorisation of state-space models into state transitions and observation likelihoods over time supports forming a factor graph and performing inference via message passing (see Fig. 1) [3, 4].

Tools for message passing on factor graphs typically employ a scheduling procedure [5, 6], in which a separate algorithm or compiler takes the model description and returns *which* nodes should pass messages *where* at *what* time. This can be sufficiently expensive to form a bottleneck. Moreover, it's not a biologically plausible mechanism for governing message passing. I explore the possibility of passing messages without a scheduler. A designated terminal node should pass an initial message, which will arrive at an initial variable. The corresponding belief is updated, a local Free Energy is computed and the belief is emitted to neighbouring factor nodes. From there on out, whenever an updated belief arrives at a factor node, the node fires messages to all other variables if the local Free Energy surpasses a threshold. If not, the node becomes silent.

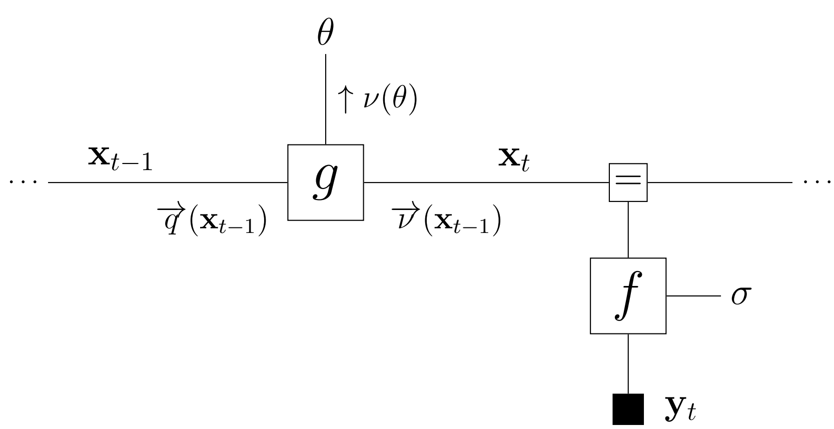

Figure 1: *Example of a Forney-style factor graph, in which variables reside on edges between factor nodes. Shown is the subgraph for time t. The function g represents a state transition from state x_{t-1} to x_t governed by parameters θ, while the function f is the likelihood of observing y_t given the current state x_t and parameters σ. An updated belief $q(\mathbf{x}_{t-1})$ arrives at node g, which decides to pass messages $\overrightarrow{\nu}(\mathbf{x}_{t-1})$ and $\uparrow \nu(\theta)$.*

[*] w.m.kouw@tue.nl

References

[1] S. Särkkä, *Bayesian filtering and smoothing*. Cambridge University Press, 2013, vol. 3.

[2] K. Friston, "The free-energy principle: a unified brain theory?" *Nature Reviews Neuroscience*, vol. 11, no. 2, p. 127, 2010.

[3] J. Winn and C. M. Bishop, "Variational message passing," *Journal of Machine Learning Research*, vol. 6, pp. 661–694, 2005.

[4] J. Dauwels, "On variational message passing on factor graphs," in *IEEE International Symposium on Information Theory*, 2007, pp. 2546–2550.

[5] T. van de Laar, M. Cox, I. Senoz, I. Bocharov, and B. de Vries, "Forneylab: a toolbox for biologically plausible free energy minimization in dynamic neural models," in *Conference on Complex Systems*, 2018.

[6] T. Minka, J. Winn, J. Guiver, Y. Zaykov, D. Fabian, and J. Bronskill, "/Infer.NET 0.3," 2018, Microsoft Research Cambridge. http://dotnet.github.io/infer.