Expectation-Propagation for large-scale imaging and computer vision problems

YOANN ALTMANN

Most imaging inverse problems or deep learning training procedures can be formulated as an optimization problem, relying on a cost function combining a data fidelity term and a regularisation term to compensate for the uncertainties associated with the measurements or training data. Very often though, the optimization of such a cost function only provides point estimates, e.g., single restored images or network parameters, without useful associated uncertainty measures. The Bayesian formalism allows for principled uncertainty management and quantification, but exact inference methods rarely scale well when the dimension of the problems increase. Thus, variational inference (VI) stands as a competitive alternative for uncertainty quantification at scale.

In this talk, we will discuss the application of Expectation-Propagation (EP) and variants, as an alternative to the more popular Variational Bayes (VB) methods used for approximate inference. In the context of image restoration, we will consider restoration from data corrupted by Gaussian and non-Gaussian noise and discuss how its modularity makes it a good candidate for a variety of inverse problems and priors/regularisation. We will then discuss how EP can also be adapted to train stochastic neural networks, where network weights and/or activation functions can be stochastic. We will finally review some of the current limitations of EP-based methods and conclude with exciting avenues combining EP within large inference schemes.