Reliable Variational Inference for Probabilistic Programming:

Several Directions

Supervisor: Luke Ong and Matthijs Vákár

The rise of big data has generated huge interest in statistical inference. The computation behind such systems boils down to the computation of high-dimensional integrals. For complex probabilistic models which do justice to the wealth of information present in big data, these integrals generally cannot be computed exactly. In practice, one often resorts to approximate algorithms like Markov Chain Monte Carlo (MCMC) methods. Recently, Variational Inference (VI) has been proposed, particularly in the machine learning community, as a (more) computationally tractable alternative to MCMC methods, enabling the scaling to even larger datasets.

While VI is proving to be very popular among machine learning practitioners for whom ad-hoc solutions may be good enough, a real question at the moment is whether VI is reliable and accurate enough for performing statistical inference, for instance on huge genomics datasets, and for safety critical machine learning applications, like self-driving cars or cancer detection. Indeed, from practical experience in the probabilistic programming language Stan, we know that we should not always trust the results of VI.

VI is lacking in at least two areas compared to MCMC:

1. The existence of a general and robust theory telling us when and how quickly we should expect the algorithm to converge to the correct answer

2. As a side effect of this theory, practical and easy to evaluate diagnostics that allow us to detect when not to trust our analyses.

We expect the candidate to begin by providing a literature review on variational inference for probabilistic programming. Thereafter several directions are possible.

Possible directions

- Prove that the VI implementation in Stan is correct with respect to its specification, to rule out this possible cause of incorrect results;

- Experimentally verify how and when VI currently fails;

- Develop diagnostics for determining when not to trusts the results of VI.

References

[1] Blei et al. Variational inference: A review for statisticians. *Journal of American Statistical Association*, 12:859–877, 2017.

[2] Alp Kucukelbir, Dustin Tran, Rajesh Ranganath, Andrew Gelman, and David M. Blei. Automatic differentiation variational inference. *Journal of Machine Learning Research*, 18:14:1–14:45, 2017.