Deep latent variable models: estimation and missing data imputation

Pierre-Alexandre Mattei National Research University, Moscow

Abstract

Deep latent variable models (DLVMs) combine the approximation abilities of deep neural networks and the statistical foundations of generative models. In this talk, we will first discuss how these models are estimated: variational methods are commonly used for inference; however, the exact likelihood of these models has been largely overlooked. We will discuss how to ensure the existence of maximum likelihood estimates, drawing connections with infinite mixture models. Then, we will present a simple variational method, called MIWAE, for training DLVMs when the training set contains missing-at-random data. Finally, we present Monte Carlo algorithms for missing data imputation using the exact conditional likelihood of DLVMs: a Metropolis-within-Gibbs sampler for DLVMs trained on complete datasets and an importance sampler for DLVMs trained on incomplete data sets. For complete training sets, our algorithm consistently and significantly outperforms the usual imputation scheme used for DLVMs. For incomplete training sets, we show that MIWAE-trained models provide accurate single and multiple imputations, and are highly competitive with state-of-the-art methods. This is joint work with Jes Frellsen, based on the following papers:

- `https://papers.nips.cc/paper/7642-leveraging-the-exact-likelihood-of-deep-latent-variable-models`
- `https://arxiv.org/abs/1812.02633`