Deep Bayesian Regression

Florian Frommlet, Medizinische Universität Wien

One of the most exciting recent developments in data analysis is deep learning. Multilayer networks have become extremely successful in performing prediction tasks and are successfully applied in many different areas. However, the resulting prediction models often difficult to interpret and potentially suffer from overfitting. The aim of this paper is to bring the ideas of deep learning into a statistical framework which yields more parsimonious models and allows to quantify model uncertainty. To this end we introduce the class of deep Bayesian regression models (DBRM) consisting of a generalized linear model combined with a comprehensive non-linear feature space, where non-linear features are generated just like in deep learning. DBRM can easily be extended to include latent Gaussian variables to model complex correlation structures between observations, which seems to be not easily possible with existing deep learning approaches. Two different algorithms based on MCMC are introduced to fit DBRM and to perform Bayesian inference. The predictive performance of these algorithms is compared with a large number of state of the art learning algorithms. Furthermore we illustrate how DBRM can be used for model inference in various applications