"Bayesian Models in Machine Learning"

Lukáš Burget. Brno University of Technology. República Checa

Summary:

Recently, Deep Neural Nets (DNN) have become a very popular machine learning tool successful in many areas (e.g. computer vision or speech processing). However, in their basic form, the discriminatively trained DNNs excel only when sufficient amount of well annotated training data is available. In contrast, Bayesian models (BMs) allow us to make robust decisions even in cases of scarce training data as they take into account the uncertainty in the model parameter estimates. The concept of latent variables makes BMs modular (i.e. more complex models can be built out of simpler ones) and well suitable for cases with missing data (e.g. unsupervised learning when annotations are missing). In this course, basic skills and intuitions about the BMs will be developed and more advanced topics will be introduced such as: approximate inference methods necessary for more complex models, infinite mixture models based on non-parametric Bms or the recently proposed Auto-Encoding Variational Bayes.

Program:

Monday
- Probability theory and probability distributions
- Generative models (Gaussian Mixture Model, Expectation-Maximization algorithm, ...)

Tuesday
- Bayesian models, (priors, uncertainty of the parameter estimates, predictive probability, ...)
- Inference in Bayesian models with conjugate priors

Wednesday
- Approximate inference in Bayesian models (Gibbs sampling, Variational Bayes, Stochastic Variational Bayes, ...)

Thursday
- Non-parametric Bayesian Models (Gaussian Process, Dirichlet Process, infinite mixture models, ...)

Friday
- Auto-Encoding Variational Bayes

For each topic, there will be uploaded and distributed some examples and demonstrations that the students will be able to download and experiment with.

Bibliography:
C. Bishop: Pattern Recognition and Machine Learning, Springer, 2006
S. J. Gershman and D.M. Blei: A tutorial on Bayesian nonparametric models, Journal of Mathematical Psychology, 2012.
P Orbanz: Tutorials on Bayesian Nonparametrics: http://stat.columbia.edu/~porbanz/npb-tutorial.html
D.P. Kingma, M. Welling: Auto-Encoding Variational Bayes, ICLR, Banff, 2014