

Exam questions.

«Neurobayesian Models», spring 2019

1. Natural gradients, local and global variables in probabilistic models. Stochastic variational inference.
2. Stochastic variational inference for the LDA model.
3. Non-conjugate variational approximations, doubly-stochastic variational inference. Unbiased gradient estimates using the log-derivative trick. Variance reduction techniques.
4. The scalable modification of Relevance Vector Machine.
5. Probabilistic PCA. The variational Auto-Encoder.
6. Stochastic gradient estimates using the reparameterization trick. The formula for the density of a random variable after a bijective differentiable transformation.
7. Implicit probabilistic models. Generative Adversarial Networks, the optimal discriminator configuration. Implicit probabilistic models as variational approximations, the adversarial auto-encoder.
8. F-divergences, lower bounds for f-divergences. f-GAN
9. Binary and Gaussian dropout. Bayesian neural networks. The log-uniform distribution. Variational dropout. The local reparameterization trick.
10. Bayesian sparsification of neural networks. Connection to the automatic relevance determination effect.
11. Stochastic Gradient Estimation in Discrete Latent Variable Models. Comparison of the REINFORCE algorithm, the reparameterization trick, and the Gumbel-Softmax trick. REBAR and hyperparameter tuning.
12. Semi-implicit probabilistic models. Hierarchical variational inference, semi-implicit variational inference, generalizations for semi-implicit prior distributions. Applications.

The last question is not mandatory for YSDA students because the lecture on this subject was canceled.

Essential knowledge, “theoretical minimum”

The following questions cover the essential concepts and algorithms from the course. To pass the exam you must be able to answer any of these questions during the exam (if you can't answer any of these questions you get 0 out of 10 for the exam).

1. The exponential family of distributions.
2. The REINFORCE algorithm.
3. The reparameterization trick.
4. The normal distribution and the multivariate normal distribution. Their properties.
5. The Kullback-Leibler divergence, approximating distributions using the Kullback-Leibler divergence
6. The general form of EM-algorithm.
7. The general form of variational inference for approximate Bayesian inference.
8. The variational Auto-Encoder: the model and how to train it.
9. Bayesian neural networks: the model and how to train it.