


Reading list: books, tutorials & reviews

(If you need to refresh your math skills)

Deisenroth, Faisal, Ong
“Mathematics for Machine Learning”

<https://mml-book.github.io/>




Reading list: books, tutorials & reviews

Chris Bishop's PRML book

<https://www.microsoft.com/en-us/research/people/cmbishop/prml-book/>

- Chapters 2, 4 (basics in probabilistic models)
- Chapters 8-11 (graphical models, approximate inference, sampling methods)



Reading list: books, tutorials & reviews

Wainwright and Jordan (2008)

“Graphical Models, Exponential Families, and Variational Inference”


Foundations and Trends in Machine Learning

https://people.eecs.berkeley.edu/~wainwrig/Papers/WaiJor08_FTML.pdf

An excellent tutorial on graphical model inference

(many topics introduced from convex optimisation point of view)

- Chapters 2-3 (basics in graphical models and exponential families)
- Chapters 4-5 (VI and message passing on graphs)
- Chapter 6 (variational methods for parameter estimations)



Reading list: books, tutorials & reviews

- Approximate inference methods:
 - Gradient estimation:
 - Mohamed et al. (2019) “Monte Carlo Gradient Estimation in Machine Learning”
 - <https://arxiv.org/abs/1906.10652>
 - Normalising flow:
 - Papamakarios et al. (2019) “Normalizing Flows for Probabilistic Modeling and Inference”
 - <https://arxiv.org/abs/1912.02762>
 - Variational auto-encoders:
 - Kingma & Welling (2019) “An Introduction to Variational Autoencoders”
 - <https://arxiv.org/abs/1906.02691>
 - EP roadmap (Tom Minka)
 - <https://tminka.github.io/papers/ep/roadmap.html>

Reading list: books, tutorials & reviews

- Sampling methods:
 - Importance sampling:
 - Art Owen's Monte Carlo book, Chapter 9 "Importance sampling"
 - <https://statweb.stanford.edu/~owen/mc/Ch-var-is.pdf>
 - Sequential Monte Carlo:
 - Doucet & Johansen (2008) "A Tutorial on Particle Filtering and Smoothing: Fifteen years later"
 - https://www.stats.ox.ac.uk/~doucet/doucet_johansen_tutorialPF2011.pdf
 - Naesseth et al. (2019) "Elements of Sequential Monte Carlo"
 - <https://www.nowpublishers.com/article/Details/MAL-074>
 - MCMC:
 - Andrieu et al. (2003) "An Introduction to MCMC for Machine Learning"
 - https://www.cs.ubc.ca/~arnaud/andrieu_defreitas_doucet_jordan_intromontecarlomachinelearning.pdf

Other tutorials to watch online

- Approximate Inference tutorial by Tom Minka (MLSS 2009, mainly message passing)
 - http://videolectures.net/mlss09uk_minka_ai/
- MCMC tutorial by Iain Murray (MLSS 2019)
 - <https://www.youtube.com/watch?v=v4Eb09qp7Q>
- Variational Inference by Blei, Ranganath, Mohamed (NeurIPS 2016)
 - https://www.youtube.com/watch?v=ogdv_6dbvVQ
- Scalable MCMC by David Dunson (NeurIPS 2018)
 - <https://www.youtube.com/watch?v=6zOTL5S9H-Y>

Other tutorials to watch online

- Variational Inference tutorial by Tamara Broderick (ICML 2018)
 - <https://www.youtube.com/watch?v=Moo4-KR5qNg>
- Bayesian deep learning tutorial by Emti Khan (NeurIPS 2019)
 - <https://slideslive.com/38923183/deep-learning-with-bayesian-principles>
- Divergences and IPMs by Gretton, Sutherland, Jitkrittum (NeurIPS 2019)
 - <https://slideslive.com/38923184/interpretable-comparison-of-distributions-and-models>
- Uncertainty estimation by Tran, Snoek, Lakshminarayanan (NeurIPS 2020)
 - Including non-Bayesian methods

Shameless Plug

- Yingzhen Li's PhD Thesis
 - "Approximate Inference: New Visions" (PhD Thesis, 2018)
 - http://yingzhenli.net/home/pdf/phd_thesis.pdf
- Cheng Zhang's review paper on variational inference
 - Zhang et al. (2019) "Advances in Variational Inference"
 - <https://ieeexplore.ieee.org/abstract/document/8588399>