10417/10617 Intermediate Deep Learning: Fall2019

Russ Salakhutdinov

Machine Learning Department
rsalakhu@cs.cmu.edu

https://deeplearning-cmu-10417.github.io/

Final Review

Neural Networks

- Neural Networks and Regularization

- Backpropagation, drop-out, L2 regularization, batch-normalization
    - ReLUs, sigmoid activations, tanh activations

- Cross-entropy error loss, squared-loss, least squared estimation

- Convolutional networks, properties, backpropagation, and optimization

Graphical Models, Unsupervised Learning

- Directed vs. Undirected Graphical models
- Conditional Independence properties for both types of graphical models
- RBMs, DBNs, VAEs as graphical models, their definition, optimization, learning rules
- Sparse Coding models, its definition, and training
- Autoencoders, definition, training objective, backprop
- Contrastive autoencoders, denoising autoencoders

RNNs, Language and Sequence Modeling

- Definition of RNNs, LSTMs
- Forward and Backward propagation in RNNs
- Neural Language Models, its properties, Markov assumption, computing probability of a sentence, training objective and optimization
- Sequence-to-sequence models, definition of attention mechanism, its advantages, its applications in machine translation, beam search
- Skip-thought models, definition and optimization
- Auto-regressive Models

GANs and VAEs, RBMs, DBNs

- Definition of GANs and VAEs, its optimization objectives
- Two-player game, discriminator and generator, WGANs
- Cycle GANs, the role of discriminator and reconstruction loss
- VAEs as deep directed graphical models
- Variational lower-bound, its definition, connection to KL
- Reparameterization trick for training VAEs, estimating gradients.
- IWAE, its definition and optimization
- Definition of RBMs and DBNs, its optimization objectives