MIDTERM EXAM GUIDE

The exam is open book, open note, open Internet and focuses on material covered in the lectures, assignments, and additional readings. The exam questions will require you to demonstrate a good understanding of the key concepts, implementation details, etc. and the ability to judge pros and cons of the approaches.

Material Covered: The class concentrates on the following three modules:

1. Deep Learning Basics,
2. Deep Generative Models,
3. Self-Supervised Learning

Hence, the midterm exam will cover all materials contained in Lectures 2-13. Topics covered in the lectures are listed in detail below:

Deep Learning Basics

- Lecture 2: Neural Building Blocks I: Spatial Processing with CNNs
  deep learning, computation in a neural net, optimization, backpropagation, convolutional neural networks, residual connections, training tricks
- Lecture 3: Neural Building Blocks II: Sequential Processing with RNNs
  sequence modeling, recurrent neural networks (RNNs), RNN applications, vanilla RNN, training RNNs, long short-term memory (LSTM), LSTM variants, gated recurrent unit (GRU)
- Lecture 4: Neural Building Blocks III: Attention and Transformers
  content-based attention, location-based attention, soft vs. hard attention, self-attention, transformer networks

Deep Generative Models

- Lecture 5: Autoregressive Models
  histograms as simple generative models, parameterized distributions and maximum likelihood, RNN-based autoregressive models, masking-based autoregressive models
- Lecture 6: Normalizing Flow Models
  1-D flows, change of variables, autoregressive flows, inverse autoregressive flows, affine flows, RealNVP, Glow, Flow++, FFJORD, multi-scale flows, dequantization
- Lecture 7: Variational Autoencoders
  latent variable models, variational autoencoders (VAEs), importance weighted autoencoders, variational lower bound/evidence lower bound, likelihood ratio gradients vs. reparameterization trick gradients, Beta-VAE, vector quantization VAEs (VQ-VAE), variational dequantization
- Lecture 8-9: Generative Adversarial Networks I
  implicit models, generative adversarial networks (GANs), evaluation metrics, theory behind GANs, GAN architectures, conditional GANs, cycle-consistent adversarial networks, representation learning in GANs, applications

- Lecture 10: Diffusion Models
  nergy based models, score-based generative models, langevin dynamics, denoising diffusion models
- Lecture 11: Strengths and Weaknesses of Current Models
  a critique of autoregressive models, flow-based models, latent variable models, implicit models, and diffusion models

Self-Supervised Learning

- Lecture 12: Self-Supervised Learning
  denoising autoencoder, in-painting, colorization, split-brain autoencoder, proxy tasks in computer vision: relative patch prediction, jigjaw puzzles, rotations, contrastive learning: word2vec, contrastive predictive coding, instance discrimination, current instance discrimination models
- Lecture 13: Pretraining Language Models
  RNN-based language models, contextualized word embeddings, scaling up generative pretraining (GPT-1, GPT-2, GPT-3) models, masked language modeling and BERT-based models