Final review

K-means clustering

Intended learning outcomes:

- Describe the differences between distance metrics and similarity functions.
- Describe the steps of the K-means algorithm, and understand how it is based on coordinate descent.
- Explain why it is important to run the algorithm several times with various starting points.
- Recall the elbow method for estimating the optimal number for K given the dataset.
- Understand what Voronoi regions are and how they are related to their corresponding clusters.

Information theory

Intended learning outcomes:

- Memorize the formulas for entropy, cross-entropy and KL-divergence, and be able to use them for simple computations.
- Understand these quantities in the context of optimal coding theory.
- Know what their properties are; e.g. non-symmetry, KL-divergence is always non-negative and is zero when $p = q$, etc.

EM algorithm

Intended learning outcomes:

- Be familiar with the distribution of the Gaussian mixture model.
- Describe the EM algorithm, and understand how it differs from K-means when used for clustering.
- Know what are the potential problems that may arise when using EM.
- Know what Bayesian Information Criterion is, and how it is used for model selection.
- Understand the EM algorithm in the general case, and how optimizing the lower bound is the same as optimizing the log-likelihood.

PCA

Intended learning outcomes:

- Understand that PCA is a form of dimensionality reduction.
- Know the steps of PCA, and the role singular value decomposition (SVD), $X = U\Sigma V^T$, plays in it.
- Understand the significance of the transformation $T = XV = U\Sigma$.

Hidden Markov Models

Intended learning outcomes:

- Write down the joint-probability distribution of HMM in terms of the initial, transition and emission probabilities.
- Know that EM algorithm is used to learn the parameters.
- Understand the forward-backward algorithm, and how it is used to solve the evaluation problem in HMMs, as well as to compute the conditional probabilities in the E-step when learning the parameters.
- Know that the Viterbi algorithm is used to solve the decoding problem in HMMs.
- Understand dynamic programming in general, and how it is used in the forward-backward and Viterbi algorithms.

Multi-armed bandit problem

Intended learning outcomes:

- Understand the exploration-exploitation trade-off in reinforcement learning.
- Know what are action-values, and how to incrementally update their estimates.
- Know what are the greedy, ϵ-greedy and UCB policies.
- Understand how Monte-Carlo tree search (MCTS) works (each of its four phases per iteration), and how it compares to minimax search.
- Understand how the exponential recency-weighted average update allows one to handle non-stationary cases.

Markov decision processes

Intended learning outcomes:

- Understand what a Markov decision process is.
- Know what policy, state-value and action-value functions are.
- Understand the Bellman equations for the value functions under a policy, and how it can be solved directly as a linear system of equations for small state spaces.
- Understand the optimal Bellman equations which the optimal value function must satisfy, and how they are non-linear.
- Understand how to iteratively converge to the solution of the Bellman equation for policy evaluation, and how to use this approximate solution for policy improvement.
- Understand policy iteration and value iteration.