

Final exam

This is a *90 minutes* open-book exam. Please read the following instructions carefully.

- The exam consists of 7 pages including this cover page. Please make sure you have all the pages before you start.
- You are allowed to use your cheat sheet (one double-sided A4 paper).
- You are *not* allowed to use any electronic devices (*e.g.*, mobile phones, tablets, laptops) during the exam.
- Please write your answers in *English* using an *indelible pen* in the provided *answer sheets*, and make sure to write your name and matriculation number on *each* answer sheet before the exam starts.
- Please write your solutions following the problem numbering on the answer sheets. If you need more space, please ask for additional answer sheets and clearly indicate where the continuation is.
- Please submit *only* your answer sheets after the exam, and you may keep the exam paper (for, *e.g.*, the exam review). This means that you can use the back (or in theory, any blank part) of the exam paper for rough work, but we will not grade anything written there.
- Please write clearly and structure your answers well. We require complete clarity as to what you are arguing or doing in your solutions. We may not be able to give you full credit for answers that are hard to read or follow.

1. *Multiple choice (20 points)*. Please mark each listed statement as true (\checkmark) or false (\times). For each answer that is correctly marked as either true or false you will get 1 point, resulting in 4 points per question. Note, that it can also happen that all answers are true or false.

- (a) Which of the following are true about probabilistic classification problems?
- Bayesian classifiers aim at finding the posterior distribution over the class labels given sample features.
 - All variants of Bayesian classifiers can only accept discrete sample features.
 - One can integrate expert knowledge into a Bayesian classifier via the prior on class labels.
 - The classification performance of chain classifiers is not dependent on the order of classes in the chain.
- (b) Which of the following statements are true regarding Metropolis-Hastings (MH) algorithm?
- The target distribution has not necessarily to be in the normalized form.
 - We should have access to the analytical expression of the target distribution when applying MH.
 - A valid proposal $q(x' | x)$ for MH should satisfy $\text{supp}(p^*) \subseteq \cup_x \text{supp}(q(\cdot | x))$.
 - The acceptance probability $A = \min \left\{ 1, \frac{p^*(x')}{p^*(x)} \right\}$ holds for all types of valid proposal $q(x' | x)$.
- (c) Which of the following statements are true about control as probabilistic inference?
- Under deterministic environment, the exact inference process formulates a soft version of the Bellman optimality equations.
 - Under stochastic environment, the optimal policy from the exact inference tends to have a risk-seeking behavior.
 - The objective of maximum entropy control is to maximize the expected cumulative reward as well as the entropy of the policy.
 - The optimal policy from maximum entropy control still has a risk-seeking behavior.
- (d) Which of the following statements are true regarding Hamiltonian Monte Carlo methods?
- HMC is a purely random-walk-based method that does not use deterministic steps.
 - The step size in HMC must always remain constant throughout the entire simulation.
 - HMC methods use the gradient of the target distribution to propose new states.
 - HMC can be more efficient than traditional Metropolis-Hastings in high dimensional problems.
- (e) Which of the following statements are true regarding Markov decision processes (MDP) and dynamical programming (DP)?
- In an MDP, the Markov property ensures that the future state depends only on the current state and action.

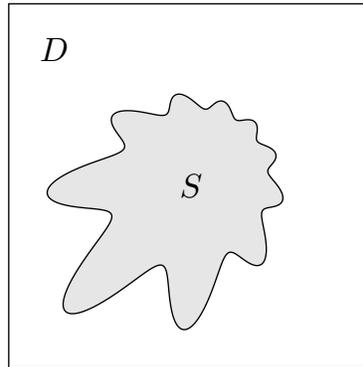
- Dynamic programming methods, such as value iteration and policy iteration, require a model of the transition probabilities and rewards.
- The Bellman equation provides a recursive relationship for the optimal value function in an MDP.
- Policy iteration is guaranteed to converge to the globally optimal policy in an MDP with finite states and actions.

2. *Bayesian networks (10 points)*. Consider a set of discrete random variables $\{A, B, C, D, E\}$. Their conditional independence relationships and parent-child dependencies are given by:

- $\text{pa}(A) = \emptyset$,
- $\text{pa}(B) = \{A\}$,
- $\text{pa}(C) = \{A\}$,
- $\text{pa}(D) = \{B, C\}$,
- $\text{pa}(E) = \{C\}$.

- (a) Draw the *directed acyclic graph* for this Bayesian network. Based on the structure, write down the factorization of the joint probability distribution $\mathbf{P}(A, B, C, D, E)$ using the chain rule.
- (b) Using the ideas of *d-separation*, determine whether the following statements are true or false for any distribution compatible with the graph. Justify your answer by identifying the type of path, *e.g.*, chain, fork, or inverted fork (collider).
- i. $(B \perp\!\!\!\perp C \mid A)$.
 - ii. $(B \perp\!\!\!\perp C \mid D)$.
- (c) For this Bayesian network, what is the Markov blanket of each variable?

3. Area estimation via Monte Carlo methods (10 points).

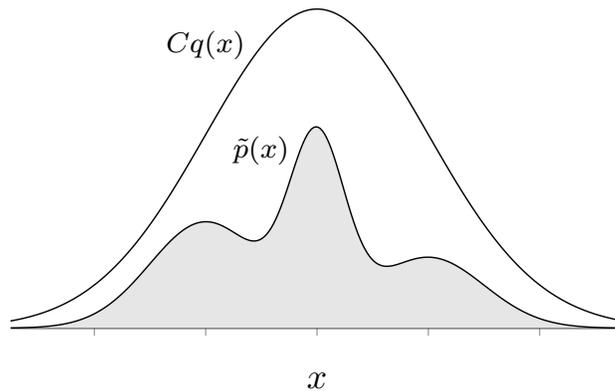


Consider two sets $S \subseteq D \subseteq \mathbf{R}^2$ in the 2-dimensional plane as shown in the figure above. Suppose we are given an oracle $f: D \rightarrow \{0, 1\}$ that can determine whether a given point $x \in D$ lies in S , *i.e.*,

$$f(x) = \begin{cases} 1 & x \in S \\ 0 & \text{otherwise.} \end{cases}$$

- (a) To estimate the area of S using Monte Carlo methods, is the information provided above sufficient? If not, what else do we need?
- (b) Based on your answer to (a), provide a detailed procedure of estimating the area of S via Monte Carlo methods, and express the estimation results mathematically.

4. *Rejection sampling (10 points).*



Consider a univariate probability density function $p(x) = \tilde{p}(x)/Z_p$ as shown in the figure above, where $\tilde{p}(x)$ is the unnormalized density and $Z_p = \int \tilde{p}(x) dx$ is the unknown normalizing constant. What we only know about p is that we can evaluate $\tilde{p}(x)$ for any given x . To sample from the target distribution p , suppose we are given a proposal distribution q that is easy to sample from, and a constant $C > 0$ such that $Cq(x) \geq \tilde{p}(x)$ for all x .

- (a) Describe the rejection sampling algorithm to obtain samples from the unknown distribution p based on the information provided above.
- (b) Prove that the samples obtained from rejection sampling are indeed distributed according to the target distribution p .

5. *Linear output hidden Markov models (10 points)*. Suppose we are given a dataset

$$(x(t), y(t)), \quad t = 1, \dots, m,$$

where $x(t) \in \mathbf{R}^n$ is the feature vector and $y(t) \in \mathbf{R}$ is the corresponding output, observed from a linear output hidden Markov model with K hidden states. It is assumed that the output $y(t)$ at each time step is generated from a linear measurement model, given by

$$y(t) = \theta_{z(t)}^T x(t) + \epsilon(t),$$

where $z(t) \in \{1, \dots, K\}$ is the hidden state at time t , $\theta_{z(t)} \in \mathbf{R}^n$ is the coefficient vector associated with state $z(t)$, and $\epsilon(t)$ is standard Gaussian noise.

We would like to estimate the parameters of the model, including the state transition probabilities P and the linear measurement coefficients $\theta_1, \dots, \theta_K$, based on the observed dataset. We consider a two-step procedure, which first estimates the coefficients θ_i together with the hidden states $z(t)$, so that in the second step the state transition probabilities P can be estimated based on the inferred hidden states.

- (a) The first step consists in minimizing the total squared error between the observed outputs and the outputs predicted by the linear output hidden Markov model. For simplicity of notation, we may transform the hidden state labels into one-hot vectors, *i.e.*, let $z(t) \in \{e_1, \dots, e_K\} \subseteq \mathbf{R}^K$, where e_i is the i -th standard basis vector in \mathbf{R}^K , which has 1 in the i -th entry and 0 elsewhere. Then we can formulate the first step of the procedure as an optimization problem, which is given by

$$\begin{aligned} \text{minimize} \quad & \sum_{t=1}^m z(t)^T \begin{bmatrix} (y(t) - \theta_1^T x(t))^2 \\ \vdots \\ (y(t) - \theta_K^T x(t))^2 \end{bmatrix} \\ \text{subject to} \quad & z(t) \in \{e_1, \dots, e_K\}, \quad t = 1, \dots, m \end{aligned} \tag{1}$$

where the optimization variables are $\theta_1, \dots, \theta_K$ and $z(1), \dots, z(m)$.

- i. What does each term inside the summation of the objective in the problem (1) represent? What is the role of the one-hot vector $z(t)$?
- ii. The problem (1) is a combinatorial optimization problem in the two groups of variables $\{\theta_i\}_{i=1}^K$ and $\{z(t)\}_{t=1}^m$ which can be very difficult to solve directly. As a simple heuristic for approximately solving (1), we can split the problem into two subproblems so that we can alternatively optimize over each group of variables while keeping the other group fixed. Based on this idea, please describe an algorithm for approximately solving the optimization problem (1). In particular, you should explain what is the objective of each subproblem in your formulation.

Hint. You might get some inspiration from the *expectation-maximization* algorithm or the *k-means clustering*.

- (b) Now suppose we have obtained an estimate of the hidden states $\tilde{z}(t) \in \{1, \dots, K\}$ for all $t = 1, \dots, m$. Explain how to estimate the state transition probabilities P based on the inferred hidden states. You can either express your answer in words or mathematically.