

STATS 305B: Practice Final

Write your name here:

Instructions:

- Write on the exam. We will scan it. If you use the extra pages at the back, label clearly.
- You can bring handwritten notes on **two sides** of an 8.5x11" piece of paper.
- Unless otherwise specified, you can write your answers using the "named distribution PDF shorthand," e.g. write the pdf of a Gaussian distribution with mean μ and variance σ^2 as $\mathcal{N}(x; \mu, \sigma^2)$.

Some tips:

1. It's usually a good idea to look through the whole exam before taking it to make sure there aren't missing pages; and so that you roughly know what you are up against
2. It's usually a good idea to skip questions if you are stuck and circle back.

Stanford Honor Code

1. The Honor Code is an undertaking of the students, individually and collectively:
 - that they will not give or receive aid in examinations; that they will not give or receive unpermitted aid in class work, in the preparation of reports, or in any other work that is to be used by the instructor as the basis of grading;
 - that they will do their share and take an active part in seeing to it that others as well as themselves uphold the spirit and letter of the Honor Code.
2. The faculty on its part manifests its confidence in the honor of its students by refraining from proctoring examinations and from taking unusual and unreasonable precautions to prevent the forms of dishonesty mentioned above. The faculty will also avoid, as far as practicable, academic procedures that create temptations to violate the Honor Code.
3. While the faculty alone has the right and obligation to set academic requirements, the students and faculty will work together to establish optimal conditions for honorable academic work.

Problem 1: Mixture Models

Consider a problem of predicting $y \in \mathbb{R}$ given inputs $\mathbf{x} \in \mathbb{R}^D$. Suppose we believe the targets are a linear function of the input on some region of input space and a different linear function on another region. We can encode these beliefs using a latent variable. The resulting model is called a *mixture of experts* model.

In this model, an “expert” $z \in \{0, 1\}$ is selected based on the input \mathbf{x} , then the target y is generated as a linear function of \mathbf{x} , where the linear parameters depend on which expert was chosen. More formally, we use the following probabilistic model:

$$p(z_n = 1 \mid \mathbf{x}_n; \boldsymbol{\theta}) = \sigma(\mathbf{c}^\top \mathbf{x}_n)$$

$$p(y_n \mid \mathbf{x}_n, z_n; \boldsymbol{\theta}) = \begin{cases} \mathcal{N}(y_n \mid \mathbf{w}_0^\top \mathbf{x}_n, \sigma_0^2), & \text{if } z_n = 0 \\ \mathcal{N}(y_n \mid \mathbf{w}_1^\top \mathbf{x}_n, \sigma_1^2), & \text{if } z_n = 1 \end{cases}$$

The parameters of this model are $\boldsymbol{\theta} = \{\mathbf{c}, \mathbf{w}_0, \mathbf{w}_1, \sigma_0, \sigma_1\}$. Suppose we observe data $\{(\mathbf{x}_n, y_n)\}_{n=1}^N$. Let $\mathbf{X} \in \mathbb{R}^{N \times D}$ denote the matrix of inputs with rows \mathbf{x}_n^\top , let $\mathbf{y} \in \mathbb{R}^N$ denote the vector of targets, and let $\mathbf{z} \in \{0, 1\}^N$ denote the vector of latent variables. In this question, you will derive the EM updates for this model.

- (a) As a warm-up, is the following statement true or false? *When applying the EM updates to the parameters, they may get trapped in a local optima.*
- (b) Write the joint log-probability $\log p(\mathbf{y}, \mathbf{z} \mid \mathbf{X}; \boldsymbol{\theta})$ for this model. Do not replace the sigmoid with its definition $\sigma(x) = \frac{1}{1+\exp(-x)}$ or substitute the pdf of the normal distribution.

Hint: Use that:

$$p(z_n \mid \mathbf{x}_n; \boldsymbol{\theta}) = [\sigma(\mathbf{c}^\top \mathbf{x}_n)]^{z_n} [(1 - \sigma(\mathbf{c}^\top \mathbf{x}_n))]^{1-z_n}$$

and:

$$p(y_n \mid \mathbf{x}_n, z_n; \boldsymbol{\theta}) = [\mathcal{N}(y_n \mid \mathbf{w}_1^\top \mathbf{x}_n, \sigma_1^2)]^{z_n} [\mathcal{N}(y_n \mid \mathbf{w}_0^\top \mathbf{x}_n, \sigma_0^2)]^{1-z_n}$$

- (c) In the E-step, we compute the posterior over latent variables given the current parameter values. Give an expression for the posterior probability $p(z_n = 1 \mid \mathbf{x}_n, y_n; \boldsymbol{\theta})$. Do not replace the sigmoid with its definition $\sigma(x) = \frac{1}{1+\exp(-x)}$ or substitute the pdf of the normal distribution.
- (d) Let $r_n = p(z_n = 1 \mid \mathbf{x}_n, y_n; \boldsymbol{\theta}^{\text{old}})$. Recall the objective for the M-step is given by:

$$\arg \max_{\boldsymbol{\theta}} \mathbb{E}_{p(\mathbf{z} \mid \mathbf{y}, \mathbf{X}; \boldsymbol{\theta}^{\text{old}})} [\log p(\mathbf{y}, \mathbf{z} \mid \mathbf{X}; \boldsymbol{\theta})]$$

Simplify this expectation, substituting r_n where appropriate. Do not replace the sigmoid with its definition $\sigma(x) = \frac{1}{1+\exp(-x)}$ or substitute the pdf of the normal distribution.

- (e) In the M-step, we maximize the expected log-joint with respect to the model parameters. For this model, it is difficult to jointly maximize with respect to $\{\mathbf{w}_0, \mathbf{w}_1\}$ and $\{\sigma_0^2, \sigma_1^2\}$. Instead, we will first maximize with respect to $\{\mathbf{w}_0, \mathbf{w}_1\}$ and then update $\{\sigma_0^2, \sigma_1^2\}$ using the updated $\{\mathbf{w}_0, \mathbf{w}_1\}$.

Using the objective above, find the M-step update for w_1 . Here, you will need to replace \mathcal{N} with the normal pdf.

Hint: You only need to consider the parts of the expected log-joint which contain w_1 .

By analogy, write down the M-step update for w_0 . You do not need to show your work for this update.

- (f) Find the M-step update for σ_1^2 . Again, you will need to replace \mathcal{N} with the normal pdf.

Hint: Maximize with respect to σ_1^2 and not σ_1 .

By analogy, write down the M-step update for σ_0^2 . You do not need to show your work for this update.

- (g) *[Bonus]* Consider the M-step update for c . Rewrite the objective from above, dropping additive constants which do not depend on c . You should recognize this as (almost) the maximum likelihood objective for a familiar statistical model. Which model is this, and what makes the M-step objective for c slightly different? Is there a closed form update for c ? If so, write this closed form update. If not, explain what could be done instead to update c .

Problem 2: Variational Autoencoders

Consider a VAE with non-negative latent variables,

$$\begin{aligned} z_{n,h} &\sim \text{Exp}(\lambda) && \text{for } h = 1, \dots, H \text{ and } n = 1, \dots, N \\ \mathbf{x}_n &\sim \mathcal{N}(g(\mathbf{z}_n; \theta), \sigma^2 I) && \text{for } n = 1, \dots, N \end{aligned}$$

where $\mathbf{x}_n \in \mathbb{R}^D$ are the observed data points, $\mathbf{z}_n = (z_{n,1}, \dots, z_{n,H})^\top \in \mathbb{R}_+^H$ is the latent “encoding” of \mathbf{x}_n , and $g : \mathbb{R}_+^H \mapsto \mathbb{R}^D$ is the decoder with parameters θ . Assume a fixed form for the approximate variational posteriors,

$$q(\mathbf{z}_n; \mathbf{x}_n, \phi) = \prod_{h=1}^H \text{Ga}(z_{n,h} \mid 2, f_h(\mathbf{x}_n; \phi)),$$

where $f_h : \mathbb{R}^D \mapsto \mathbb{R}_+$ is an encoder with parameters ϕ that outputs the rate parameter for the gamma posterior on the h -th coordinate of the encoding. Note that we are assuming the shape parameter of the gamma posterior is fixed to 2.

To learn the parameters of the parameters of the encoder (ϕ) and decoder (θ), we maximize the ELBO,

$$\mathcal{L}(\theta, \phi) = \sum_{n=1}^N \mathbb{E}_{q(\mathbf{z}_n; \mathbf{x}_n, \phi)} [\log p(\mathbf{x}_n \mid \mathbf{z}_n; \theta, \sigma^2)] - D_{\text{KL}}(q(\mathbf{z}_n; \mathbf{x}_n, \phi) \parallel p(\mathbf{z}_n; \lambda)),$$

where $p(\mathbf{z}_n; \lambda)$ denotes the independent exponential prior from above.

(a) (1.5 pts) Is the following statement true or false? Explain your answer.

$$\nabla_{\theta} \mathcal{L}(\theta, \phi) = \sum_{n=1}^N \mathbb{E}_{q(\mathbf{z}_n; \mathbf{x}_n, \phi)} [\nabla_{\theta} \log p(\mathbf{x}_n \mid \mathbf{z}_n; \theta, \sigma^2)].$$

(b) (1.5 pts) Is the following statement true or false? Explain your answer.

$$\nabla_{\phi} \mathcal{L}(\theta, \phi) = \sum_{n=1}^N \mathbb{E}_{q(\mathbf{z}_n; \mathbf{x}_n, \phi)} [\nabla_{\phi} \log p(\mathbf{x}_n \mid \mathbf{z}_n; \theta, \sigma^2)] - \nabla_{\phi} D_{\text{KL}}(q(\mathbf{z}_n; \mathbf{x}_n, \phi) \parallel p(\mathbf{z}_n; \lambda)).$$

- (c) (5 pts) Derive a reparameterization trick for the gamma posterior. Specifically, write a sample $z_{n,h} \sim \text{Ga}(2, f_h(\mathbf{x}_n; \phi))$ as a transformation of two uniform random variables, $u_1, u_2 \stackrel{\text{iid}}{\sim} \text{Unif}([0, 1])$.

- (d) (5 pts) Compute a closed form expression for the KL divergence in the ELBO.

$$D_{\text{KL}}(q(\mathbf{z}_n; \mathbf{x}_n, \phi) \parallel p(\mathbf{z}_n; \lambda)) = \mathbb{E}_{q(\mathbf{z}_n; \mathbf{x}_n, \phi)} [\log q(\mathbf{z}_n; \mathbf{x}_n, \phi) - \log p(\mathbf{z}_n; \lambda)]$$

(e) (2 pts) In 1-2 sentences, explain how you would use the results from (c) and (d) to obtain an unbiased estimator of $\nabla_{\phi} \mathcal{L}(\theta, \phi)$.

Problem 3: Hidden Markov Models

Consider a hidden Markov model (HMM),

$$\begin{aligned} z_1 &\sim \text{Cat}(\boldsymbol{\pi}_0) \\ z_t &\sim \text{Cat}(\boldsymbol{\pi}_{z_{t-1}}) && \text{for } t = 2, \dots, T \\ \mathbf{x}_t &\sim p(\mathbf{x} \mid \boldsymbol{\theta}_{z_t}) && \text{for } t = 1, \dots, T \end{aligned}$$

where $z_t \in \{1, \dots, K\}$ denotes the discrete latent state at time t . The HMM parameters consist of the initial state probabilities, $\boldsymbol{\pi}_0$, the transition matrix $\boldsymbol{\Pi}$ with rows $\boldsymbol{\pi}_k \in \Delta_K$ for $k = 1, \dots, K$, and the emission parameters $\{\boldsymbol{\theta}_k\}_{k=1}^K$.

We can represent the latent states in terms of their states at the times of change points, $c_n \in \{1, \dots, K\}$, and the corresponding durations between change points, $d_n \in \mathbb{N}_+$ (these are positive integers). Let $c_1 = z_1$ and let d_1 denote the number of time steps before the state changes. Then let $c_2 = z_{1+d_1}$ and let d_2 be the number of time steps before the state changes again. For example, this state sequence would be represented as follows,

$$\mathbf{z}_{1:T} = \underbrace{3, 3, 3, 3, 3}_{c_1=3, d_1=5}, \underbrace{1, 1, 1, 1}_{c_2=1, d_2=4}, \underbrace{2, 2, 2, 2, 2, 2}_{c_3=2, d_3=6}, \underbrace{1, 1, 1, 1, 1}_{c_4=1, d_4=5}.$$

(a) (5 pts) Derive the probability mass function $p(d_n \mid c_n = k)$.

(b) (5 pts) Derive the probability mass function $p(c_{n+1} \mid c_n = k)$.

- (c) (5 pts) Suppose we want to model a system in which the emissions are conditionally independent given the latent state, as above, but where durations are uniformly distributed between 1 and 3 time steps. That is, for all k ,

$$p(d_n = i \mid c_n = k) = \begin{cases} \frac{1}{3} & \text{if } i = 1, 2, 3 \\ 0 & \text{otherwise} \end{cases}$$

How could model such a system using an HMM on an extended state space? Specifically, let $z'_t \in \{1, \dots, K'\}$ denote the states of your HMM; you may have $K' > K$. What are the transition and emission probabilities?

- (d) () What is the computational complexity of the forward-backward algorithm for this HMM?

Problem 4: *Recurrent Neural Networks*

Consider a stochastic RNN with linear hidden state dynamics

$$h_{t+1} \sim N(ah_t, 1),$$

and initial distribution $h_0 \sim N(0, 1)$. This is called a Gaussian linear dynamical system (LDS). It's also similar to the noising process we studied in the context of denoising diffusion models.

(a) Derive the conditional distribution, $p(h_T | h_0)$?

(b) Suppose we receive one observation, $y_T \sim N(h_T, 1)$. Compute the posterior, $p(h_0 | y_T)$.

(c) In this case the posterior has a closed form, but suppose we didn't know that and we tried to solve for the posterior mode $h_0^{(\text{MAP})}$ using gradient ascent. Compute the gradient $\nabla_{h_0} \log p(h_0, y_T)$.

(d) How does the gradient scale with T , and why could that be problematic for gradient based learning?