EECS 182 Deep Neural Networks Spring 2023 Anant Sahai

This homework is due on Friday, April 28, 2022, at 10:59PM.

1. Coding Question: Summarization (Part II)

Please follow the instructions in this notebook. You will learn how to efficiently enable the Transformer encoder-decoder model to generate sequences. Then, you will fine-tune another Transformer encoder-decoder model based on the pretrained language model T5 for the same task.

Note: this notebook takes **45 min to 1 hour** to run after you implemented everything correctly. Please start early.

- Download submission_log.json and submit it to "Homework 11 (Code) (Summarization)" in Gradescope.
- Answer the following questions in your submission of the written assignment:
- (a) What are the ROUGE-1, ROUGE-2, and ROUGE-L scores of the model trained from scratch on the summarization task?
- (b) In the T5 paper, which optimizer is employed for training T5, and what are the peak learning rates for pretraining and finetuning?
- (c) What are the ROUGE-1, ROUGE-2, and ROUGE-L scores of the finetuned model, and how do these scores compare to those of the model trained from scratch?
- (d) Provide appropriate prompts for the summarization task, apply the prompt to the example documents, and fed them into an off-the-shelf large language model (e.g., ChatGPT). Specify the model used, and share the generated outputs. Compare the outputs with your finetuned model's outputs qualitatively and identify the main reason for any differences.

2. Comparing Distributions

Divergence metrics provide a principled measure of difference between a pair of distributions (P, Q). One such example is the Kullback-Leibler Divergence, that is defined as

$$D_{\mathrm{KL}}(P||Q) = \mathbb{E}_{x \sim P(x)} \left[\log \frac{P(x)}{Q(x)}\right]$$

- (a) Technically D_{KL} is not a true distance since it is asymmetric, i.e. generally $D_{\text{KL}}(P||Q) \neq D_{\text{KL}}(Q||P)$. Give an example of univariate distributions P and Q where $D_{\text{KL}}(P||Q) \neq \infty$, $D_{\text{KL}}(Q||P) = \infty$.
- (b) For a fixed target distribution P, we call $D_{\text{KL}}(P||Q)$ the *forward-KL*, while calling $D_{\text{KL}}(Q||P)$ the *reverse-KL*. Due to the asymmetric nature of KL, distributions Q that minimize $D_{\text{KL}}(P||Q)$ can be different from those minimizing $D_{\text{KL}}(Q||P)$.

From the following plots, identify which of (A, B) correspond to minimizing forward vs. reverse **KL**. Give brief reasoning. Here, only the mean and standard deviation of Q is allowed to vary during the minimization.



3. Continual Learning (Optional)

Run this notebook and answer the questions. We will explore some strategies that we can mitigate catastrophic forgetting when our neural network model sequentially learns the new tasks. Let's try and compare three classic methods: 1) naive 2) Elastic Weight Consolidation (EWC) and 3) Rehearsal.

- (a) Naive approach
 - i. What do you observe? How much does the network forget from the previous tasks? Why do you think this happens?
 - ii. (Open-ended question) We are using CNN. Does MLP perform better or worse than CNN? Try it out and report your results.
- (b) Elastic Weight Consolidation
 - i. Hyperparameter is underexplored in this assignment. Try different values of λ and report your results.
 - ii. What is the role of λ ? What happens if λ is too small or too large? Explain the results with plasticity and stability of the network.
- (c) Rehearsal
 - i. What would be the pros and cons of rehearsal?

4. Variational AutoEncoders

(Parts of this problem are adapted from Deep Generative Models, Stanford University)

For this problem we will be using PyTorch to implement the variational autoencoder (VAE) and learn a probabilistic model of the MNIST dataset of handwritten digits. Formally, we observe a sequence of binary pixels $\mathbf{x} \in \{0, 1\}^d$ and let $\mathbf{z} \in \mathbb{R}^k$ denote a set of latent variables. Our goal is to learn a latent variable model $p_{\theta}(\mathbf{x})$ of the high-dimensional data distribution $p_{data}(\mathbf{x})$.

The VAE is a latent variable model with a specific parameterization $p_{\theta}(\mathbf{x}) = \int p_{\theta}(\mathbf{x}, \mathbf{z}) d\mathbf{z} = \int p(\mathbf{z}) p_{\theta}(\mathbf{x}|\mathbf{z}) d\mathbf{z}$ Specifically, VAE is defined by the following generative process (often called **reparameterization trick**):

 $p(\mathbf{z}) = \mathcal{N}(\mathbf{z}|0, I)$ (sample noise from standard Gaussian) $p_{\theta}(\mathbf{x}|\mathbf{z}) = \text{Bern}(\mathbf{x}|f_{\theta}(\mathbf{z}))$ (decode noise to generate sample from real-distribution)

That is, we assume that the latent variables z are sampled from a unit Gaussian distribution $\mathcal{N}(z|0, I)$. The latent z are then passed through a neural network decoder $f_{\theta}(\cdot)$ to obtain the parameters of the *d* Bernoulli random variables that model the pixels in each image.

To learn the parameterized distibution we would like to maximize the marginal likelihood $p_{\theta}(\mathbf{x})$. However computing $p_{\theta}(\mathbf{x}) = \int p(\mathbf{z})p_{\theta}(\mathbf{x}|\mathbf{z})d\mathbf{z}$ is generally intractable since this requires integrating over all possible values of $\mathbf{z} \in \mathbb{R}$. Instead, we consider a variational approximation to the true posterior

$$q_{\phi}(\mathbf{z}|\mathbf{x}) = \mathcal{N}\left(\mathbf{z}|\mu_{\phi}(\mathbf{x}), \operatorname{diag}(\sigma_{\phi}^{2}(\mathbf{x}))\right)$$

In particular, we pass each image \mathbf{x} through a neural network that outputs mean μ_{ϕ} and diagonal covariance $\operatorname{diag}(\sigma_{\phi}^2(\mathbf{x}))$ of the multivariate Gaussian distribution that approximates the distribution over the latent variables \mathbf{z} given \mathbf{x} . The high level intuition for training parameters (θ, ϕ) requires considering two expressions:

- Decoding Latents : Sample latents from $q_{\phi}(\mathbf{z})$, maximize likelihood of generating samples $\mathbf{x} \sim p_{data}$
- Matching Prior : A Kullback-Leibler (KL) term to constraint $q_{\phi}(\mathbf{z})$ to be close to the $p(\mathbf{z})$

Putting these terms together, gives us a lower-bound of the true marginal log-likehood, called the **evidence lower bound** (ELBO):

$$\log p_{\theta}(\mathbf{x}) \geq \text{ELBO}(\mathbf{x}; \theta, \phi) = \underbrace{\mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})}[\log p_{\theta}(\mathbf{x}|\mathbf{z})]}_{\text{Decoding Latents}} - \underbrace{D_{\text{KL}}(q_{\phi}(\mathbf{z}|\mathbf{x})||p(\mathbf{z})))}_{\text{Matching Prior}}$$

In this notebook, implement the reparameterization trick in the function sample_gaussian. Specifically, your answer will take in the mean m and variance v of the Gaussian and return a sample $\mathbf{x} \sim \mathcal{N}(m, diag(v))$

Then, implement negative_elbo_bound loss function.

Note: We ask for the negative ELBO, as PyTorch optimizers minimze the loss function. Furthere, since we are computing the negative ELBO over a mini-batch of data $\{x^{(i)}\}_{i=1}^{n}$, make sure to compute the average of per-sample ELBO. Finally, note that the ELBO itself cannot be computed exactly since computation of the reconstruction term is intractable. Instead, you should estaimate the reconstruction term via Monte-Carlo sampling

$$-\mathbb{E}_{q_{\phi}(\mathbf{z}|\mathbf{x})}[\log p_{\theta}(\mathbf{x}|\mathbf{z})] \approx -\log p_{\theta}(\mathbf{x}|\mathbf{z}^{(1)})$$

where $\mathbf{z}^{(1)} \sim q_{\phi}(\mathbf{z}|\mathbf{x})$ denotes a single sample from the learned posterior.

The negative_elbo_bound expects as output three quantities: *average* negative ELBO, reconstruction loss, KL divergence.

Answer these questions:

(a) Test your implementation by training VAE with

python experiment.py --model vae

Once the run is complete (10000 iterations), report the following numbers : the average

- negative ELBO
- KL-Divergence term
- reconstruction loss

Since we're using stochastic optimization, you may wish to run the model multipple times and **report** each metric's mean and corresponding standard error (*Hint: the negative ELBO on the test subset* should be ~ 100)

(b) Visualize 200 digits (generate a single image tiled in a grid of 10×20 digits) sampled from $p_{\theta}(\mathbf{x})$

5. Generative Adversarial Networks (Optional)

(Parts of this problem are adapted from Deep Generative Models, Stanford University)

Unlike VAEs, that explicitly model data distributions with likelihood-based training, Generative Adversarial Networks (GANs) belong to the family of implicit generative models.

To model high-dimensional data distributions $p_{\text{data}}(\mathbf{x})$ (with $\mathbf{x} \in \mathbb{R}^n$), define

- a generator $G_{\theta} : \mathbb{R}^k \to \mathbb{R}^n$
- a discriminator $D_{\phi} : \mathbb{R}^n \to (0, 1)$

To obtain samples from the generator, we first sample a k-dimensional random vector $\mathbf{z} \sim \mathcal{N}(0, 1)$ and return $G_{\theta}(\mathbf{z}) \in \mathbb{R}^n$. The discriminator is effectively a classifier that judges how realistic the *fake* image $G_{\theta}(\mathbf{z})$ are, compared to *real* samples from the data distribution $x \sim p_{\text{data}}(\mathbf{x})$. Because its output is intended to be interpreted as a probability, the last layer of the discriminator is frequently the **sigmoid** function,

$$\sigma(x) = \frac{1}{1 + e^{-x}}$$

such that $\sigma(x) \in (0,1)$. Therefore, for logits $h_{\phi}(\mathbf{x})$, discriminator output is $D_{\phi}(\mathbf{x}) = \sigma(h_{\phi}(\mathbf{x}))$.

For training GANs we define learning objectives $L_{discriminator}(\phi; \theta)$ and $L_{generator}(\theta; \phi)$ that are optimized iteratively in two-stages with gradient descent. In particular, we take a gradient step to minimize $L_{discriminator}(\phi; \theta)$ w.r.t discriminator parameters ϕ , followed by gradient step to minimize $L_{generator}(\theta; \phi)$ w.r.t. generator parameters θ . In lecture we've considered following versions of the losses:

$$L_{\text{discriminator}}(\phi;\theta) = -\underbrace{\mathbb{E}_{\mathbf{x}\sim p_{\text{data}}}\left[\log D_{\phi}(\mathbf{x})\right]}_{\text{Real Data}} - \underbrace{\mathbb{E}_{\mathbf{z}\sim\mathcal{N}(0,I)}\left[\log\left(1 - D_{\phi}(G_{\theta}(\mathbf{z}))\right)\right]}_{\text{Generated Data}}$$

$$L_{\text{generator}}^{\text{minimax}}(\theta;\phi) = \mathbb{E}_{\mathbf{z}\sim\mathcal{N}(0,I)}\left[\log\left(1 - D_{\phi}(G_{\theta}(\mathbf{z}))\right)\right]$$

Training a GAN can be viewed as solving the following minimax optimization problem, for generator G_{θ} and discriminator D_{ϕ} :

$$\min_{G} \max_{D} V(G, D) \equiv \mathbb{E}_{\mathbf{x} \sim p_{\text{data}}} \Big[\log D_{\phi}(\mathbf{x}) \Big] + \mathbb{E}_{\mathbf{z} \sim \mathcal{N}(0, I)} \Big[\log \big(1 - D_{\phi}(G_{\theta}(\mathbf{z})) \big) \Big]$$

(a) Vanishing Gradient with Minimax Objective

Rewriting the above loss in terms of discriminator logits, sigmoid we have

$$L_{\text{generator}}^{\text{minimax}}(\theta;\phi) = \mathbb{E}_{\mathbf{z}\sim\mathcal{N}(0,I)} \Big[\log \big(1 - \sigma(h_{\phi}(G_{\theta}(\mathbf{z})))\big) \Big]$$

Show that $\nabla_{\theta} L_{\text{generator}}^{\text{minimax}}(\theta; \phi) \to 0$ when discriminator output $D_{\phi}(G_{\theta}(\mathbf{z})) \approx 0$. Why is this problematic for training the generator when the discriminator is well-trained in identifying fake samples?

(b) GANs as Divergence Minimization

To build intuition about the training objective, consider the distribution $p_{\theta}(\mathbf{x})$ corresponding to:

$$\mathbf{x} = G_{\theta}(\mathbf{z})$$
 where $\mathbf{z} \sim \mathcal{N}(0, I)$

• Optimal Discriminator

The discriminator minimizes the loss

$$L_{\text{discriminator}}(\phi;\theta) = -\mathbb{E}_{\mathbf{x}\sim p_{\text{data}}} \left[\log D_{\phi}(\mathbf{x})\right] - \mathbb{E}_{\mathbf{x}\sim p_{\theta}(\mathbf{x})} \left[\log \left(1 - D_{\phi}(\mathbf{x})\right)\right]$$

For a fixed generator θ , show that the discriminator loss is minimized when $D_{\phi}^* = \frac{p_{\text{data}}(\mathbf{x})}{p_{\theta}(\mathbf{x}) + p_{\text{data}}(\mathbf{x})}$.

Generator Loss

For a fixed generator θ , and corresponding optimal discriminator D_{ϕ}^* , show that the minimax objective $V(G, D^*)$ satisfies

$$V(G, D^*) = -\log 4 + 2D_{\text{JSD}}(p_{\text{data}}||p_{\theta})$$

where $D_{\text{JSD}}(p||q)$ is the Jenson-Shannon Divergence.

Note: A divergence measures the distance between two distributions p, q. In particular, for distributions p, q with common support \mathcal{X} , typically used divergence metrics include

$$D_{\mathrm{KL}}(p||q) = \mathbb{E}_{\mathbf{x}\sim p} \left[\log \frac{p(x)}{q(x)} \right]$$
(Kullback-Leibler Divergence)
$$D_{\mathrm{JSD}}(p||q) = \frac{1}{2} D_{\mathrm{KL}} \left(p||\frac{p+q}{2} \right) + \frac{1}{2} D_{\mathrm{KL}} \left(q||\frac{p+q}{2} \right)$$
(Jensen-Shannon Divergence)

(c) Training GANs on MNIST

To mitigate vanishing gradients during training, ? propose the non-saturating loss

$$L_{\text{generator}}^{\text{ns}}(\theta;\phi) = -\mathbb{E}_{\mathbf{z}\sim\mathcal{N}(0,I)} \Big[\log D_{\phi}(G_{\theta}(\mathbf{z}))\Big]$$

For mini-batch approximation, we use Monte-Carlo estimates of the learning objectives, such that

$$\begin{split} L_{\text{discriminator}}(\phi;\theta) &\approx -\frac{1}{m} \sum_{i=1}^{m} \log D_{\phi}(\mathbf{x}^{(i)}) - \frac{1}{m} \sum_{i=1}^{m} \log \left(1 - D_{\phi}(G_{\theta}(\mathbf{z}^{(i)}))\right) \\ L_{\text{generator}}^{\text{ns}}(\phi;\theta) &\approx -\frac{1}{m} \sum_{i=1}^{m} \log D_{\phi}(G_{\theta}(\mathbf{z}^{(i)})) \end{split}$$

for batch-size *m*, and batches of *real-data* $\mathbf{x}^{(i)} \sim p_{\text{data}}(\mathbf{x})$ and *fake-data* $\mathbf{z}^{(i)} \sim \mathcal{N}(0, I)$. Following these details, implement training for GANs with above learning objectives by filling relevant snippets in gan.py. Test your implementation by running

python experiment.py --model gan

Visualize 200 digits (generate a single image tiled in a grid of 10×20 digits) sampled from $p_{\theta}(x)$

6. 0th Order Optimization - Policy Gradient

We will now talk about 0^{th} order optimization, also known as Policy Gradient in a Reinforcement Learning context. Although this method is primarily used in an RL context we will be adapting this method to do 0^{th} order optimization on a Neural Network.

 k^{th} order optimization means that in the optimization, we use a k^{th} order derivative $(\frac{\delta L^k}{\delta^k w})$ to do the opti-

mization. So we can see that gradient descent is a first order optimization method, while Newton's method is a second order optimization method.

Polciy gradient is a 0^{th} order optimization method - which means that you use no derivative for the optimization. This is used in contexts in which the loss is a **blackboxed** function, hence propogating a gradient through it is impossible.

Policy gradient at a high level approximates the gradient and then does gradient descent using this approximated gradient.

(a) **Prove** that

$$p_{\theta}(x)\nabla_{\theta}\log(p_{\theta}(x)) = \nabla_{\theta}p_{\theta}(x)$$

(b) Let's say we have a neural network f(x) which takes in a x and uses the weights(w) to output 2 logits (P = [P(y = 0), P(y = 1)]).

Let p(x, y) be the joint distribution of the input and output data according to **our model**. Hence $p_w(x, y) = p(x)p_w(y|x)$, where p(x) is the ground distribution of x, while $p_w(y|x) = f(x)[y]$ is what our model predicts.

Similarly we have a **blackboxed** loss function L(x, f(x)) which outputs a loss. For example if I wanted to learn to classify y = 1 if x > 5 and y = 0 otherwise, L(4, (0.1, 0.9)) would be small while L(4, (0.9, 0.1)) would be very high. As we already discussed, since this loss is blackboxed we can't take the derivative through it.

We want to optimize the following objective function

$$w^* = argmin_w J(w)$$

where

$$J(w) = E_{(x,f(x)) \sim p_w(x,y)}[L(x,f(x))].$$

To do this optimization we want to approximate $\nabla_w J(w)$ so that we could use an optimization method like gradient descent to find w^*

Prove that $\nabla_w J(w)$ can be approximated as $\frac{1}{N} \sum_{i=1}^{i=N} (\nabla_w \log(p_w(y_i|x_i)) L(x_i, f(x_i)))$

Hints:

- Try creating a $\tau = (x, f(x))$
- $E[X] = \int_x x P(X = x) dx$
- Use the result from part a which was $p_{\theta}(x)\nabla_{\theta}\log(p_{\theta}(x)) = \nabla_{\theta}p_{\theta}(x)$
- $p_w(x,y) = p(x)p_w(y|x)$
- (c) The following two parts are based on this notebook, where you need to implement policy gradient according to your derivation before answering these questions.

Include the screenshot of the accuracy plot here. With a correct implementation, you should observe a test accuracy of approximately 75% after the final iteration.

(d) Compare the policy gradient and supervised learning approaches for this classification task, focusing on their convergence speed, stability, and final performance. Explain any observed differences.

7. Homework Process and Study Group

Citing sources and collaborators are an important part of life, including being a student! We also want to understand what resources you find helpful and how much time homework is taking, so we can change things in the future if possible.

- (a) What sources (if any) did you use as you worked through the homework?
- (b) If you worked with someone on this homework, who did you work with? List names and student ID's. (In case of homework party, you can also just describe the group.)
- (c) Roughly how many total hours did you work on this homework? Write it down here where you'll need to remember it for the self-grade form.

Contributors:

- Linyuan Gong.
- Kumar Krishna Agrawal.
- Suhong Moon.
- Dhruv Shah.
- Anant Sahai.
- Aditya Grover.
- Stefano Ermon.
- Yashish Mohnot.