

EECS 182 Deep Neural Networks
 Fall 2025 Anant Sahai and Gireeja Ranade

Homework 3

This homework is due on Sep. 26th, 2025, at 10:59PM.

1. Maximal Update Parameterization

During lecture, we learned how SGD and Adam can be seen as constrained optimization problems, defined by some norm over parameter space. In this problem (which builds directly on the last discussion section), we will explore the payoff for defining these norms over the *output features* of a dense layer and then bringing it back to the matrices involved.

- (a) Consider a dense layer with input $\mathbf{x} \in \mathbb{R}^{d_1}$, weights $W \in \mathbb{R}^{d_2 \times d_1}$, where W is initialized with i.i.d. standard Gaussian entries, and entries of \mathbf{x} are sampled from an i.i.d. unit Gaussian.

What is the expected squared RMS norm of the output features $\mathbf{y} = W\mathbf{x}$? How does this scale with d_1 or d_2 ? What constant should we multiply W by to ensure that the expected squared RMS norm of $W\mathbf{x}$ is 1, regardless of d_1 and d_2 ?

Hint: Consider a simplified dense layer with a single output feature, $W \in \mathbb{R}^{1 \times d_1}$. What is the distribution of the scalar $y = W\mathbf{x}$? What is its variance?

- (b) We will now consider how to ensure that RMS norms of features update at a constant rate *during training*, again regardless of layer width. Assume we are using the SignGD optimizer (which is a simplified version of Adam). Unlike at initialization, where we assume weights and inputs are independent, the updates to weights during training are very much correlated with the inputs. For simplicity, assume that the minibatch training input $\mathbf{x}_i \in \mathbb{R}^{d_1}$ are sampled from an i.i.d. unit Gaussian, and the raw gradient $\nabla_W f(W)$ is an outer product of the input and a constant backpropagated vector $\mathbf{g}_i \in \mathbb{R}^{d_2}$.

$$W_{t+1} \leftarrow W_t + \eta \text{sign}(\mathbf{x}_i \mathbf{g}_i^T).$$

What is the expected RMS norm squared of the (unscaled) change in features $\mathbf{y}' = \text{sign}(\mathbf{x}_i \mathbf{g}_i^T) \mathbf{x}_i$? How does this scale with d_1 or d_2 ? What constant should we multiply the update by to ensure that the expected RMS norm of \mathbf{y}' is 1, regardless of d_1 and d_2 ?

- (c) You may notice that the above update rule only depends on d_1 . **Why is this the case?**

2. Visualizing Maximal Update Parameterization (Coding Question)

In this question, using simple MLPs, you will explore how the maximal update parameterization preserves feature learning as the width of the network changes. Please follow the instructions in [this Google Colab notebook](#).¹

3. Maximal Update Parameterization Research

This homework problem talks about the research papers behind “maximal update parameterization,” also called μP in the community. This idea is discussed in detail in the paper [Tensor Programs V: Tuning Large](#)

¹This notebook is very newly developed and has not been as extensively tested as previously assigned coding problems. Nonetheless, we believe it is a great aid in learning these timely concepts. If you encounter any issues, please post on Ed or notify the course staff in their office hours.

Neural Networks via Zero-Shot Hyperparameter Transfer. Understanding the full paper is out of scope of the class and requires sophistication that is not a prerequisite, but we want to highlight some points.

- Look at Figure 1 of the paper. **What are the two side-by-side figures plotting? What is the observation you make about the width of a matrix?**
- Look at Table 3 in the paper. In lecture, we gave an elementary argument to support the $1/\text{fan}_{\text{in}}$ scaling for Adam LR for the hidden layers, the rightmost entry in the bottom row of the paper. **What does the table say is the corresponding standard parameterization?**
- The paper **A Spectral Condition for Feature Learning** presents the same scaling using elementary linear algebra. While understanding the entirety of the paper is out of scope of the class, we presented some simplified arguments in lecture.

Look at Desideratum 1. **Why is this behavior desired for feature learning? Can you rewrite this Desideratum in terms of the RMS norm? Similarly, rewrite Condition 1 in terms of the RMS norm as well. How does Condition 1 imply that $\|h_\ell(x)\|_2 \leq \Theta(\sqrt{n_\ell})$ and that $\|\Delta h_\ell(x)\|_2 \leq \Theta(\sqrt{n_\ell})$? What is the key assumption that allows us to also get the lower bounds on $\|h_\ell(x)\|_2$ and $\|\Delta h_\ell(x)\|_2$?**

4. Policy Gradient and the Reparameterization Gradient Estimator

In this question, you are going to derive two gradient estimators for the following objective function:

$$\mathcal{F}(\theta) = \mathbb{E}_{\mathbf{x} \sim p_\theta} [f(\mathbf{x})], \quad (1)$$

where \mathbf{x} is a random variable that follows the probability distribution of $p_\theta : \mathcal{X} \mapsto \Delta(\mathcal{X})$ that is parameterized by $\theta \in \mathbb{R}^P$, and $f : \mathcal{X} \mapsto \mathbb{R}$ is a function.

- Let \mathbf{x} be a k -D multivariate Gaussian random variable that is parameterized by the mean $\mu \in \mathbb{R}^K$ under the distribution,

$$p_\mu(\mathbf{x}) = (2\pi\sigma)^{-n/2} \exp\left(-\|\mathbf{x} - \mu\|_2^2 / (2\sigma^2)\right), \quad (2)$$

where $\sigma \in \mathbb{R}$ is a scalar constant. **Express $\nabla_\mu \mathcal{F}(\mu)$ as an expectation under $p_\mu(\mathbf{x})$** (i.e., find $g(\mathbf{x})$ in $\nabla \mathcal{F}(\mu) = \mathbb{E}_{\mathbf{x} \sim p_\mu} [g(\mathbf{x})]$). (Hint: $\mathbb{E}_{\mathbf{x} \sim p_\theta} [f(\mathbf{x})] = \int p_\theta(\mathbf{x}) f(\mathbf{x}) d\mathbf{x}$)

- Rewrite the expression in $\mathcal{F}(\mu)$ with an expectation with the distribution over a standard normal $\mathcal{N}(0, \mathbf{I})$?** (Hint: if \mathbf{x} follows the distribution of $p_\mu(\mathbf{x})$, then $\mathbf{x} - \mu$ follows the distribution of $\mathcal{N}(0, \sigma^2 \mathbf{I})$)
- Using the expression you obtained from the previous part, **can you express $\nabla_\mu \mathcal{F}(\mu)$ in a similar way** (i.e., with an expectation with the distribution over $p_{\mu=0}(\mathbf{x})$)? (Hint: ∇_μ can be safely moved inside the expectation because the expectation no longer follows a distribution that depends on μ and expectations are linear.)

For the following two parts, we are going to generalize our findings from the specific Gaussian case to arbitrary p_θ 's (you may assume p_θ has well-defined gradient $\nabla_\theta p_\theta$ that expresses what change in parameters would make a particular realization more likely).

- In the general case, **can you still write $\nabla \mathcal{F}(\theta)$ as an expectation under $p_\theta(\mathbf{x})$?** (Hint: look at your answer in Part (a) and see how the term inside your expectation relates to $\log p_\theta(\mathbf{x})$).
- Assume there exists a function $g(\mathbf{z}, \theta) : \mathcal{Z} \times \mathbb{R}^P \mapsto \mathcal{X}$ and a distribution over \mathbf{z} , $p(\mathbf{z})$, such that $g(\mathbf{z}, \theta)$ has the same distribution as $\mathbf{x} \sim p_\theta(\mathbf{x})$. **Prove that $\nabla_\theta \mathcal{F}(\theta) = \mathbb{E}_{\mathbf{z} \sim \mathcal{Z}} [\nabla_\theta g(\mathbf{z})^\top \nabla_{\mathbf{x}} f|_{\mathbf{x}=g(\mathbf{z}, \theta)}]$.** In the

previous part (b) and (c), we actually prove a special case of this. **Can you determine what g and $p(z)$ are for the special case?**

These kinds of tricks are practically useful because the spirit of stochastic gradient descent plays nicely whenever we have gradients that we can view as expectations.

5. Tensor Rematerialization

You want to train a neural network on a new chip designed at UC Berkeley. Your model is a 10 layer network, where each layer has the same fixed input and output size of s . The chip your model will be trained on is heavily specialized for model evaluation. It can run forward passes through a layer very fast. However, it is severely memory constrained, and can only fit in memory the following items (slightly more than twice of the data necessary for performing a forward pass):

- (a) the inputs;
- (b) $2s$ activations in memory;
- (c) optimizer states necessary for performing the forward pass through the current layer.

To train despite this memory limitation, your friend suggests using a training method called **tensor rematerialization**. She proposes using SGD with a batch size of 1, and only storing the activations of every 5th layer during an initial forward pass to evaluate the model. During backpropagation, she suggests recomputing activations on-the-fly for each layer by loading the relevant last stored activation from memory, and rerunning forward through layers up till the current layer.

Figure 1 illustrates this approach. Activations for Layer 5 and Layer 10 are stored in memory from an initial forward pass through all the layers. Consider when weights in layer 7 are to be updated during backpropagation. To get the activations for layer 7, we would load the activations of layer 5 from memory, and then run them through layer 6 and layer 7 to get the activations for layer 7. These activations can then be used (together with the gradients from upstream) to compute the gradients to update the parameters of Layer 7, as well as get ready to next deal with layer 6.




Figure 1: Tensor rematerialization in action - Layer 5 and Layer 10 activations are stored in memory along with the inputs. Activations for other layers are recomputed on-demand from stored activations and inputs.

- (a) Assume a forward pass of a single layer is called a `fwd` operation. **How many `fwd` operations are invoked when running a single backward pass through the entire network?** Do not count the initial forward passes required to compute the loss, and don't worry about any extra computation beyond activations to actually backprop gradients.
- (b) Assume that each memory access to fetch activations or inputs is called a `loadmem` operation. **How many `loadmem` operations are invoked when running a single backward pass?**
- (c) Say you have access to a local disk which offers practically infinite storage for activations and a `loaddisk` operation for loading activations. You decide to not use tensor rematerialization and

instead store all activations on this disk, loading each activation when required. Assuming each `fwd` operation takes 20ns and each `loadmem` operation (which loads from memory, not local disk) takes 10ns for tensor rematerialization, **how fast (in ns) should each `loaddisk` operation be to take the same time for one backward pass as tensor rematerialization?** Assume activations are directly loaded to the processor registers from disk (i.e., they do not have to go to memory first), only one operation can be run at a time, ignore any caching and assume latency of any other related operations is negligible.

6. 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?**

Contributors:

- Kevin Frans.
- Anant Sahai.
- Gireeja Ranade.
- Qiyang Li.
- Romil Bhardwaj.