EECS 182 Deep Neural Networks
Fall 2026 Gireeja Ranade

Homework 10

This homework is due on Wednesday, November 11, 2026, at 10:59PM.

1. Ridge-Attention

In lecture, you saw how the standard softmax-attention mechanism can be viewed as a softened version of something akin to a nearest-neighbor model in which the value returned for a query reflects a weighted combination of the values that correspond to the keys closest to the query. In this view, the (key, value) pairs in the memory represent a kind of in-context "training data" and the query is a test input for which we want to predict the right output given that data.

(a) To start, let's think about why it is possible to efficiently update simple averaging. Let $m = \frac{1}{n}\sum_{i=1}^{n} x_i$ be the average of n points. **Use m, x_{n+1}, n and simple arithmetic operations to compute $m' = \frac{1}{n+1}\sum_{i=1}^{n+1} x_i$ — the average of all points including the new point x_{n+1}.**
(HINT: Start by multiplying m by n.)

(b) Let us now shift to thinking about traditional ridge-regression with n training pairs $(\mathbf{x}_i, y_i)$ where $\mathbf{x}_i$ are d-dimensional vectors and y_i are scalars. Let the matrix $A = \begin{bmatrix} \mathbf{x}_1^\top \\ \mathbf{x}_2^\top \\ \vdots \\ \mathbf{x}_n^\top \end{bmatrix}$ and vector $\mathbf{y} = \begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_n \end{bmatrix}$ so that we can find the familiar closed-form solution

$$\mathbf{w}_* = (A^\top A + \lambda I)^{-1} A^\top \mathbf{y} \tag{1}$$

that allows us to make scalar predictions on any new test input $\mathbf{x}$ by computing $\mathbf{w}_*^\top \mathbf{x}$.
First, **write the two terms $(A^\top A + \lambda I)$ and $A^\top \mathbf{y}$ as sums involving the $\mathbf{x}_i$ and y_i.**
i.e. Complete:
$(A^\top A + \lambda I) = \lambda I + \sum_{i=1}^{n}$
$A^\top \mathbf{y} = \sum_{i=1}^{n}$

(c) Suppose we wanted to do ridge-self-attention (non-causal – "encoder-style") with a context length of n and d-dimensional query, key, and value vectors. Recall that this is the style of attention where each of the n queries is applied to the same pool of n (key, value) pairs. The goal is to calculate multi-dimensional ridge-regression predictions, after "training" on the pool of (key, value) pairs, and given the query as a kind of "test" input. (Note: the keys are playing the role of the A matrix in ridge-regression, each query is like the $\mathbf{x}$ we are testing on, and in place of the scalar y_i, we have an entire value vector for a multi-dimensional ridge-regression problem so the $\mathbf{y}$ vector is replaced with a matrix that has a row for each value vector.)

Assume that the cost of inverting a $d \times d$ matrix is $O(d^3)$ and the cost of multiplying two such matrices is also $O(d^3)$. Assume that a $d \times d$ matrix times either a d-dimensional row or column vector costs d^2 operations. You should assume $d < n$.

What is the computational cost of a non-causal ridge self-attention layer?

 All Rights Reserved. This may not be publicly shared without explicit permission.

- ◯ $O(d^4)$
- ◯ $O(nd^2)$
- ◯ $O(n^2d^3)$
- ◯ $O(n^2d^2)$
- ◯ $O(n^2)$
- ◯ $O(1)$

(HINT: Do not forget that for a single d-dimensional query vector $\mathbf{q}$, attention needs to return a d-dimensional result.)

(d) Assume that a ridge self-attention layer is used in a Transformer architecture and there is a downstream loss. **For which of these will backprop successfully pass gradients if we use ridge self-attention?**

- ☐ The ridge λ viewed as a learnable parameter for the self-attention layer.
- ☐ The keys
- ☐ The values
- ☐ The queries

(e) Now step back. There is a nice trick (called the Sherman–Morrison Formula) by which one can update the inverse of an invertible matrix to which you make a rank-1 update. Let M be an invertible square $d \times d$ matrix and let $\mathbf{u}, \mathbf{v}$ be two d-dimensional vectors. Then:

$$(M + \mathbf{u}\mathbf{v}^\top)^{-1} = M^{-1} - \frac{1}{1 + \mathbf{v}^\top M^{-1}\mathbf{u}}(M^{-1}\mathbf{u})(\mathbf{v}^\top M^{-1}) \tag{2}$$

Assume that a $d \times d$ matrix times either a d-dimensional row or column vector costs d^2 operations, and so does the evaluation of a dyad $\mathbf{u}\mathbf{v}^\top$. Assume that computing a Euclidean inner-product costs d operations. **Assuming that you already had M^{-1} in hand, what is the computational cost of one application of (2)?**

- ◯ $O(d^4)$
- ◯ $O(d^3)$
- ◯ $O(d^2)$
- ◯ $O(d)$
- ◯ $O(1)$

(f) Consider implementing causal ridge-self-attention with a context length of n but where the pool of (key, value) vectors that one is querying at position t consists only of the t (key, value) pairs so far.

Describe explicitly how you would compute causal ridge-self-attention in a computationally efficient manner. Leverage your decomposition of the ridge-regression formula in part (b) of this problem together with the Sherman-Morrison formula from (2) to avoid having to do $O(n^2)$ computations while still calculating causal ridge-self-attention outputs correctly for all n positions in the context.

(HINT: Think recursively. What do you need to track from one time step to the next to avoid repeating work?)

(g) Many people consider important the ability to visualize the attention weights. For traditional softmax-attention, these are the outputs of the softmax for a given query vector. They tell you the exact amounts by which the attention outputs at this position are linear combinations of the values being fed in at this and other positions.

For ridge-attention and a given query vector, how would you compute the (possibly negative) weights associated to each of the value vectors in the context?

 All Rights Reserved. This may not be publicly shared without explicit permission.

2. Ordinary Softmax Multihead Attention Implementation

(a) Complete the incomplete lines in the implementation of multi-head attention (with a final linear layer on top) by filling in (on the following page) what (a),(b),(c), etc. should be in the code below.

```
class MultiHeadAttention(nn.Module):
    def __init__(self, d_model, num_heads):
        super(MultiHeadAttention, self).__init__()

        self.d_model = d_model
        self.num_heads = num_heads
        self.d_k = d_model // num_heads

        self.W_q = nn.Linear(d_model, d_model)
        self.W_k = nn.Linear(d_model, ___(a)___)
        self.W_v = nn.Linear(___(b)___, d_model)
        self.W_o = nn.Linear(d_model, ___(c)___)

    def scaled_dot_product_attention(self, Q, K, V, mask=None):
        attn_scores = torch.matmul(___(d)___, K.transpose(-2, -1))
        attn_scores = attn_scores /  math.sqrt(self.d_k)
        if mask is not None:
            attn_scores = attn_scores.masked_fill(mask == 0, -1e9)
        attn_probs = torch.___(e)____(attn_scores, dim=-1)
        output = torch.matmul(attn_probs, ____(f)___)
        return output

    def split_heads(self, x):
        batch_size, seq_length, d_model = x.size()
        return x.view(batch_size, seq_length,
                      _____(g)_____, self.d_k).transpose(1, 2)

    def combine_heads(self, x):
        batch_size, _, seq_length, d_k = x.size()
        return x.transpose(1, 2).view(batch_size,
                                      seq_length, self.d_model)

    def forward(self, Q, K, V, mask=None):
        Q = self.split_heads(self.W_q(Q))
        K = self.split_heads(self.W_k(K))
        V = self.split_heads(self.W_v(V))

        attn_output = self.scaled_dot_product_attention(Q, K, V, mask)
        output = self._____(h)_____(self.combine_heads(attn_output))
        return output
```

Fill in the missing blanks. Some have been given to you so you can see what kind of answers we are looking for. You need to write in the boxes.

 All Rights Reserved. This may not be publicly shared without explicit permission.

(a)

(b)

(c) `d_model`

(d)

(Hint: which of `Q`*,* `K`*,* `V` *should go here?)*

(e)

(Hint: What operation do we need to use to go from scores to probabilities?)

(f)

(Hint: which of `Q`*,* `K`*,* `V` *should go here?)*

(g) `self.num_heads`

(h)

(Hint: which linear model do we still need to apply?)

(b) Suppose that you want to update the initialization `W_v` in this way:

```
self.W_v = nn.Linear(in_features=d_model, out_features=2 * d_model)
```

This increases the length of the value vectors that we get.

What other minimal modifications should you apply to the code to make the code run?

(Hint: Remember that in transformers, we typically use attention blocks with residual connections around them. What does that require?)

3. Transformer Decoding Optimization

Consider a multi-head attention (MHA) layer in a transformer decoder performing incremental decoding. The attention layer operates on a batch of b sequences, each with a dimension of d. For each of the h heads, the dimensions of query, key, and value projections are $q = \frac{d}{h}$, $k = \frac{d}{h}$, and $v = \frac{d}{h}$ respectively.

The following pseudocode performs the multi-head self-attention operation for a single token generation in the sequence (for each sequence in our batch).

```
def multihead_attention_incremental(x, prev_K, prev_V, W_q, W_k, W_v, W_o):
    """Multi-head Self-Attention (one step).
    Args:
        x: a tensor with shape [b, d] - current token embedding
        prev_K: tensor with shape [b, h, n, k] - cached keys from prev. tokens
        prev_V: tensor with shape [b, h, n, v] - cached values from prev.
            tokens
        W_q: a tensor with shape [h, d, q] - query projection weights
        W_k: a tensor with shape [h, d, k] - key projection weights
        W_v: a tensor with shape [h, d, v] - value projection weights
        W_o: a tensor with shape [h, v, d] - output projection weights
    Returns:
        y: a tensor with shape [b, d] - output embedding
```

 All Rights Reserved. This may not be publicly shared without explicit permission.

```
    """
    # Project inputs to query, key, and value
    q = torch.einsum("bd,hdq->bhq", x, W_q)  # [b, h, q]
    k = torch.einsum("bd,hdk->bhk", x, W_k)  # [b, h, k]
    v = torch.einsum("bd,hdv->bhv", x, W_v)  # [b, h, v]

    # Append new key and value to previous cached states
    new_K = torch.cat([prev_K, k.unsqueeze(2)], dim=2)  # [b, h, n+1, q]
    new_V = torch.cat([prev_V, v.unsqueeze(2)], dim=2)  # [b, h, n+1, v]

    # Compute attention scores, apply softmax, and get weighted values
    logits = torch.einsum("bhq,bhnq->bhn", q, new_K)  # [b, h, n+1]
    weights = torch.nn.functional.softmax(logits, dim=-1)  # [b, h, n+1]
    o = torch.einsum("bhn,bhnv->bhv", weights, new_V)  # [b, h, v]

    # Project back to output dimension
    y = torch.einsum("bhv,hvd->bd", o, W_o)  # [b, d]

    return y
```

Multi-Query Attention (MQA): Multi-query attention is a variant of multi-head attention where all heads share the same keys and values, reducing memory access and improving efficiency for incremental decoding.

```
def multiquery_attention_incremental(x, prev_K, prev_V, W_q, W_k, W_v, W_o):
    """Multi-query Self-Attention (one step).
    Args:
        x: a tensor with shape [b, d] - current token embedding
        prev_K: tensor with shape [b, n, k] - cached keys (shared across heads
            )
        prev_V: tensor with shape [b, n, v] - cached values (shared across
            heads)
        W_q: a tensor with shape _____A____ - query projection weights
        W_k: a tensor with shape _____B____ - key projection weights
        W_v: a tensor with shape [d, v] - value projection weights
        W_o: a tensor with shape [h, v, d] - output projection weights
    Returns:
        y: a tensor with shape [b, d] - output embedding
    """
    # Project input to queries (per head), keys and values (shared)
    q = _____C_____
    k = _____D_____
    v = _____E_____

    # Append new key and value to previous cached states
    new_K = torch.cat([prev_K, k.unsqueeze(1)], dim=1)  # [b, n+1, k]
    new_V = torch.cat([prev_V, v.unsqueeze(1)], dim=1)  # [b, n+1, v]

    # Compute attention scores, apply softmax, and get weighted values
    logits = torch.einsum("bhk,bnk->bhn", q, new_K)     # [b, h, n+1]
    weights = torch.nn.functional.softmax(logits, dim=-1)  # [b, h, n+1]
    o = torch.einsum("bhn,bnv->bhv", weights, new_V)    # [b, h, v]

    # Project back to output dimension
```

 All Rights Reserved. This may not be publicly shared without explicit permission.

```
    y = torch.einsum("bhv,hvd->bd", o, W_o)              # [b, d]

    return y
```

First, fill in the blanks in the code/comments above by picking one:

(a) (3pts) **What is A?**

◯ $[h, d, q]$ ◯ $[d, q]$ ◯ $[h, q]$

(b) (3pts) **What is B?**

◯ $[h, d, k]$ ◯ $[d, k]$ ◯ $[h, k]$

(c) (4pts) **What is C?**

◯ `torch.einsum("bd,dq->bhq", x, W_q)`
◯ `torch.einsum("bd,hdq->bhq", x, W_q)`
◯ `torch.einsum("bhd,dq->bhq", x, W_q)`

(d) (4pts) **What is D?**

◯ `torch.einsum("bd,dk->bk", x, W_k)`
◯ `torch.einsum("bd,hdk->bhk", x, W_k)`
◯ `torch.einsum("bd,dk->bhk", x, W_k)`

(e) (4pts) **What is E?**

◯ `torch.einsum("bd,hdv->bhv", x, W_v)`
◯ `torch.einsum("bd,dv->bhv", x, W_v)`
◯ `torch.einsum("bd,dv->bv", x, W_v)`

(f) (4pts) Recall that for generating a batch of single tokens in multi-head-attention with a cache of n previous tokens, the operations have computational complexity $O(bd^2 + bnd)$ since it is $O(bd^2)$ to do the projections (query, key, value, output) and $O(bnd)$ from attention (dot products with n cached keys/values for each head). The memory access complexity is $O(d^2 + bd + bnd)$ since it is $O(d^2)$ to read weights, $O(bd)$ to read/write current token embeddings (x, q, k, v, o, y), and $O(bnd)$ to read cached keys/values. The MHA arithmetic intensity is thus $\frac{O(bd^2+bnd)}{O(d^2+bd+bnd)}$.

- Both terms are kept; which one dominates depends on the relative size of n and d.
- If n is small, projections dominate; if n is large, memory access for cached keys/values dominates.

Now consider MQA. For generating a single token with a cache of n previous tokens, **what is the computational complexity (in Big-O notation)?**

(g) (4pts) For MQA generating a single token with a cache of n previous tokens, **what is the memory access complexity (in Big-O notation)?**

4. Coding Question: Visualizing Attention

Please run the cells in the Visualizing_BERT.ipynb notebook, then answer the questions below.

(a) Attention in GPT: Run part a of the notebook and generate the corresponding visualizations

i. **What similarities and differences do you notice in the visualizations between the examples in this part?** Explore the queries, keys, and values to identify any interesting patterns associated with the attention mechanism.

 All Rights Reserved. This may not be publicly shared without explicit permission.

ii. **How does attention differ between the different layers of the GPT model? Do you notice that the tokens are attending to different tokens as we go through the layers of the network?**

(b) BERT pays attention: Run part b of the notebook and generate the corresponding visualizations.

i. Look at different layers of the BERT model in the visualizations of part (b) and identify different patterns associated with the attention mechanism. Explore the queries, keys, and values to further inform your answer. **For instance, do you notice that any particular type of tokens are attended to at a given timestep?**

ii. **Do you spot any differences between how attention works in GPT vs. BERT? Think about how the model architectures are different.**

iii. For the example with syntactically similar but definitionally different sentences, look through the different layers of the two BERT networks associated with sentence a and sentence b, and take a look at the queries, keys, and values associated with the different tokens. **Do you notice any differences in the embeddings learned for the two sentences that are essentially identical in structure but different in meaning?**

iv. **For the pre-training related examples, do you notice BERT's bi-directionality in play? Do you think pre-training the BERT helped it learn better representations?**

(c) BERT has multiple heads!: Run part c of the notebook and generate the corresponding visualizations.

i. **Do you notice different features being learned throughout the different attention heads of BERT? Why do you think this might be?**

ii. **Can you identify any of the different features that the different attention heads are focusing on?**

(d) Visualizing untrained attention weights

i. **What differences do you notice in the attention patterns between the randomly initialized and trained BERT models?**

ii. **What are some words or tokens that you would expect strong attention between? What might you guess about the gradients of this attention head for those words?**

(e) **Were you able to identify interesting patterns in the visualizations?** If yes, please share some examples (describe in text or paste a screenshot). If not, feel free to use this space for your frustrations.

5. Kernelized Linear Attention (Part I)

The softmax attention is widely adopted in transformers, however the $\mathcal{O}\left(N^2\right)$ (N stands for the sequence length) complexity in memory and computation often makes it less desirable for processing long document like a book or a passage, where the N could be beyond thousands. There is a large body of the research studying how to resolve this [1].

Under this context, this question presents a formulation of attention via the lens of the kernel. A large portion of the context is adopted from [2]. In particular, attention can be seen as applying a kernel over the inputs with the kernel scores being the similarities between inputs. This formulation sheds light on individual components of the transformer's attention, and helps introduce some alternative attention mechanisms that replaces the "softmax" with linearized kernel functions, thus reducing the $\mathcal{O}\left(N^2\right)$ complexity in memory and computation.

[1] https://huggingface.co/blog/long-range-transformers

[2] Tsai, Yao-Hung Hubert, et al. "Transformer dissection: a unified understanding of transformer's attention via the lens of kernel" (2019).

 All Rights Reserved. This may not be publicly shared without explicit permission.

We first review the building block in the transformer. Let $x \in \mathbb{R}^{N \times F}$ denote a sequence of N feature vectors of dimensions F. A transformer[3] is a function $T : \mathbb{R}^{N \times F} \to \mathbb{R}^{N \times F}$ defined by the composition of L transformer layers $T_1(\cdot), \cdots, T_L(\cdot)$ as follows,

$$T_l(x) = f_l(A_l(x) + x). \tag{3}$$

The function $f_l(\cdot)$ transforms each feature independently of the others and is usually implemented with a small two-layer feedforward network. $A_l(\cdot)$ is the self attention function and is the only part of the transformer that acts across sequences.

We now focus on the the self attention module which involves softmax. The self attention function $A_l(\cdot)$ computes, for every position, a weighted average of the feature representations of all other positions with a weight proportional to a similarity score between the representations. Formally, the input sequence x is projected by three matrices $W_Q \in \mathbb{R}^{F \times D}$, $W_K \in \mathbb{R}^{F \times D}$ and $W_V \in \mathbb{R}^{F \times M}$ to corresponding representations Q, K and V. The output for all positions, $A_l(x) = V'$, is computed as follows,

$$\begin{aligned} Q = xW_Q, K = xW_K, V = xW_V, \\ A_l(x) = V' = \text{softmax}(\frac{QK^T}{\sqrt{D}})V. \end{aligned} \tag{4}$$

Note that in the previous equation, the softmax function is applied rowwise to QK^T. Following common terminology, the Q, K and V are referred to as the "queries", "keys" and "values" respectively.

Equation 4 implements a specific form of self-attention called softmax attention where the similarity score is the exponential of the dot product between a query and a key. Given that subscripting a matrix with $i\cdot$ returns the i-th row as a vector and subscripting with $\cdot j$ returns the j-th column, we can write a generalized attention equation for any similarity function as follows,

$$V'_{i\cdot} = \frac{\sum_{j=1}^{N} \text{sim}\left(Q_{i\cdot}, K_{j\cdot}\right) V_{j\cdot}}{\sum_{j=1}^{N} \text{sim}\left(Q_{i\cdot}, K_{j\cdot}\right)}. \tag{5}$$

Equation 5 is equivalent to equation 4 if we substitute the similarity function with $\text{sim}_{\text{softmax}}(q, k) = \exp(\frac{q^T k}{\sqrt{D}})$. This can lead to

$$V'_{i\cdot} = \frac{\sum_{j=1}^{N} \exp(\frac{Q_{i\cdot} K_{j\cdot}^T}{\sqrt{D}}) V_{j\cdot}}{\sum_{j=1}^{N} \exp(\frac{Q_{i\cdot} K_{j\cdot}^T}{\sqrt{D}})}. \tag{6}$$

For computing the resulting self-attended feature $A_l(x) = V'$, we need to compute all $V'_{i\cdot}$ $i \in 1, ..., N$ in equation 6.

(a) **Identify the conditions that needs to be met by the** sim **function to ensure that** $V_{i\cdot}$ **in Equation 5 remains finite (the denominator never reaches zero).**

(b) The definition of attention in equation 5 is generic and can be used to define several other attention implementations.

(i) One potential attention variant is the "polynomial kernel attention", where the similarity function

[3] Vaswani, Ashish, et al. "Attention is all you need" (2017).

 All Rights Reserved. This may not be publicly shared without explicit permission.

as sim (q, k) is measured by polynomial kernel $\mathcal{K}$ [4]. **Considering a special case for a "quadratic kernel attention" that the degree of "polynomial kernel attention" is set to be 2, derive the sim (q, k) for "quadratic kernel attention". (NOTE: any constant factor is set to be 1.)** .

(ii) One benefit of using kernelized attention is that we can represent a kernel using a feature map $\phi(\cdot)$ [5]. **Derive the corresponding feature map $\phi(\cdot)$ for the quadratic kernel.**

(iii) **Considering a general kernel attention, where the kernel can be represented using feature map that $\mathcal{K}(q, k) = (\phi(q)^T \phi(k))$, rewrite kernel attention of equation 5 with feature map $\phi(\cdot)$.**

(c) We can rewrite the softmax attention in terms of equation 5 as equation 6. **For all the $V_{i\cdot}'$ ($i \in \{1, ..., N\}$), derive the time complexity (asymptotic computational cost) and space complexity (asymptotic memory requirement) of the above softmax attention in terms of sequence length N, D and M.**

NOTE: for memory requirement, we need to store any intermediate results for backpropagation, including all Q, K, V

(d) Assume we have a kernel $\mathcal{K}$ as the similarity function and the kernel can be represented with a feature map $\phi(\cdot)$, we can rewrite equation 5 with $\text{sim}(x, y) = \mathcal{K}(x, y) = (\phi(Q_{i\cdot})^T \phi(K_{j\cdot}))$ in part (b). We can then further simplify it by making use of the associative property of matrix multiplication to

$$V_{i\cdot}' = \frac{\phi(Q_{i\cdot}) \sum_{j=1}^{N} \phi(K_{j\cdot})^T V_{j\cdot}}{\phi(Q_{i\cdot}) \sum_{j=1}^{N} \phi(K_{j\cdot})^T}. \tag{7}$$

Note that the feature map $\phi(\cdot)$ is applied row-wise to the matrices Q and K.

Considering using a linearized polynomial kernel $\phi(x)$ of degree 2, and assume $M \approx D$, derive the computational cost and memory requirement of this kernel attention as in (7).

6. Coding Question: Summarization

Please follow the instructions in this notebook. You will implement a Transformer using fundamental building blocks in PyTorch. You'll apply the Transformer encoder-decoder model to a sequence-to-sequence NLP task: document summarization. Refer to the Attention is All You Need paper for details on the model architecture.

Once you finished with the notebook, answer the following questions in your submission of the written assignment:

(a) **Please submit the screenshots of the training loss and the validation loss displayed on Tensorboard.**

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?**

[4] https://en.wikipedia.org/wiki/Polynomial_kernel
[5] https://en.wikipedia.org/wiki/Kernel_method

 All Rights Reserved. This may not be publicly shared without explicit permission.

(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: Matteo Guarrera, Mert Cemri, Sizhe Chen, Gireeja Ranade, Anant Sahai, Yossi Gandelsman, Naman Jain, Dhruv Shah, Shivam Singhal, Kevin Li, Olivia Watkins, Bryan Wu, Sheng Shen, Shaojie Bai, Angelos Katharopoulos, Hao Peng, Linyuan Gong, Hao Liu, Allie Gu, CS182 Staff from past semesters.

 All Rights Reserved. This may not be publicly shared without explicit permission.