

This week we will work through a set of exercises on attention tensor shapes, the structure of multi-head attention, hand-computed single-head and multi-head forward passes, and the block-matrix view used in practical MHA implementations.

1. Single-Head Attention Shapes

The figure below shows a vanilla single-head attention block. Assume the input tensor has shape $X \in \mathbb{R}^{B \times n \times d}$, where B is the batch size, n is the sequence length, and d is the token embedding dimensionality. Assume further that the dimension of the key vector d_k and dimension of the value vector d_v for the attention are both equal to d . This question asks you to determine the shapes of relevant tensors.

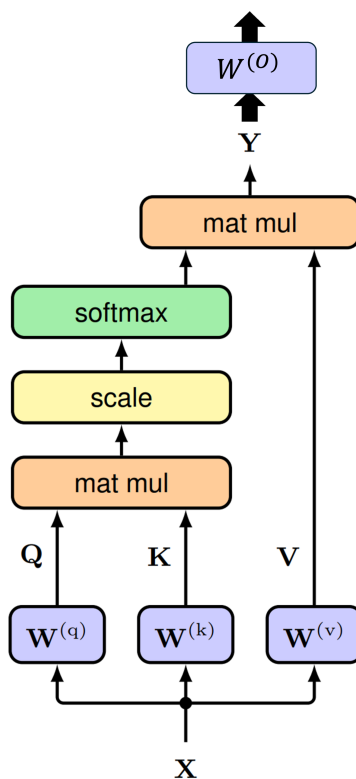


Figure 1: Single-head attention block.

- (a) Please provide the shapes for the query/key/value input projection weight matrices $W^{(q)}$, $W^{(k)}$, $W^{(v)}$.
- (b) Please provide the shapes for the query/key/value activations Q , K , V after the input projection.
- (c) Please provide the shapes for the result after the first matmul QK^\top , the scaled matmul result $\frac{QK^\top}{\sqrt{d}}$, and the softmax result (i.e. the attention weights) $\text{Softmax}\left(\frac{QK^\top}{\sqrt{d}}\right)$. Note that the formulae here are for unbatched inputs.
- (d) Please provide the shapes for the result of the second matmul Y , the output projection weight matrix $W^{(o)}$, and the final output of the attention block after the output projection $YW^{(o)}$.

2. Multi-Head Attention Mechanism

The figure below shows a sketch of the multi-head attention mechanism. In this question we will review how we can have multiple attention heads from a single input, and how each head's outputs are combined to produce a single MHA output tensor.

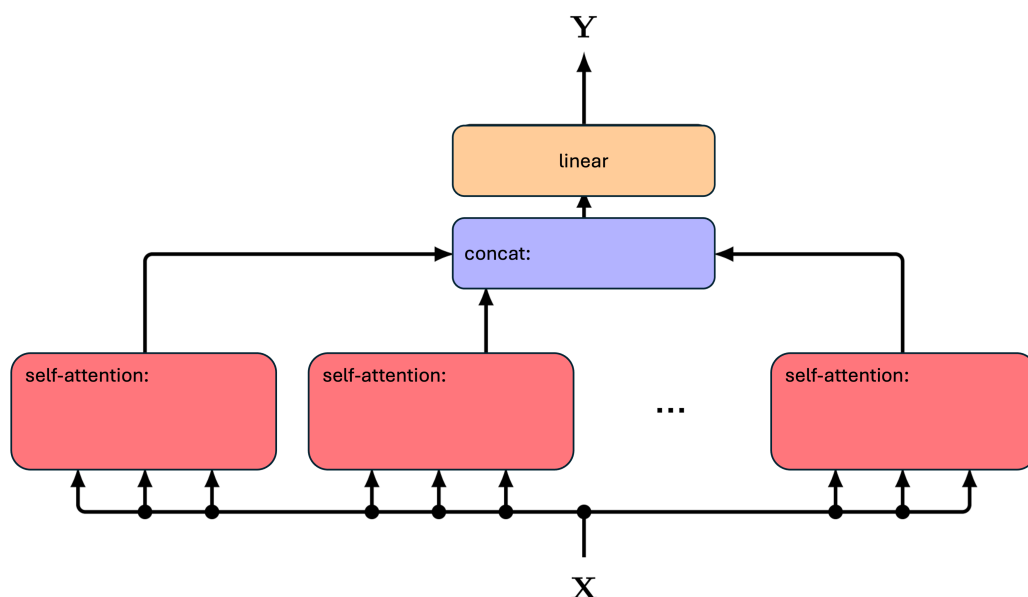


Figure 2: Multi-head attention block.

- Please draw on the figure where each query/key/value input projection weight matrices $W_h^{(q)}$, $W_h^{(k)}$, $W_h^{(v)}$ are located.
- Please draw on the figure where each query/key/value activation Q_h , K_h , V_h and the per-head attention output H_h are calculated.
- Please draw on the figure how the per-head attention outputs $H_1, H_2, \dots, H_{n_{\text{heads}}}$ are combined to produce a single tensor before the final output projection linear layer.

3. Single Head Attention Forward Pass

Recall the single head attention mechanism:

$$Q = XW^{(q)}$$

$$K = XW^{(k)}$$

$$V = XW^{(v)}$$

$$H = \text{Attention}(Q, K, V) = \text{Softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right)V$$

$$Y = HW^{(o)}$$

In this question, we consider a toy example where sequence length $n = 2$, hidden dimension $d = 2$, and the input is not batched. Assume the input $X \in \mathbb{R}^{n \times d}$, the input projection weights $W^{(q)}, W^{(k)}, W^{(v)} \in \mathbb{R}^{d \times d}$, and the output projection weight $W^{(o)} \in \mathbb{R}^{d \times d}$ bear the following values:

$$X = \begin{bmatrix} 1 & 1 \\ 0 & 1 \end{bmatrix}, \quad W^{(q)} = \begin{bmatrix} 3 & -1 \\ 1 & 1 \end{bmatrix}, \quad W^{(k)} = \begin{bmatrix} 2 & -2 \\ 1 & 3 \end{bmatrix}, \quad W^{(v)} = \begin{bmatrix} 2 & -2 \\ 0 & 2 \end{bmatrix}, \quad W^{(o)} = \begin{bmatrix} 2 & 0 \\ 0 & 1 \end{bmatrix}.$$

(a) Compute the matrices Q , K , and V .

(b) Given that the attention weights are approximately

$$\text{Softmax}\left(\frac{QK^\top}{\sqrt{d_k}}\right) \approx \begin{bmatrix} 1.00 & 0.00 \\ 0.50 & 0.50 \end{bmatrix}$$

Please compute the output Y of the single head attention.

4. MHA with Block Matrices For multi-head attention with number of heads $n_{\text{heads}} > 1$, in practice we often do not explicitly keep separate projection matrices $W_h^{(q)}$, $W_h^{(k)}$, $W_h^{(v)}$ for each head in the forward computation. Instead, we form single large projection matrices $W^{(q)}$, $W^{(k)}$, $W^{(v)}$ and compute

$$\begin{aligned} Q &= XW^{(q)} \\ K &= XW^{(k)} \\ V &= XW^{(v)}. \end{aligned}$$

Then each head operates on the appropriate feature chunk of Q , K , and V .

For example, if $n_{\text{heads}} = 2$ and hidden dimension $d = 64$, then each head has size $d_h = 32$. Head 0 operates on $Q[:, : 32]$, $K[:, : 32]$, $V[:, : 32]$, and head 1 operates on $Q[:, 32 :]$, $K[:, 32 :]$, $V[:, 32 :]$.

To show why this is valid, recall the block-matrix identity (here the square brackets denote matrix concatenate-

nation instead of indexing)

$$A[B_0 \mid B_1] = [AB_0 \mid AB_1].$$

Here B_0 and B_1 are a split of B along its column axis.

(a) Consider the two-head case $n_{\text{heads}} = 2$ and assume $d = 2d_h$. Write

$$W^{(q)} = [W_0^{(q)} \mid W_1^{(q)}],$$

where $W_0^{(q)}, W_1^{(q)} \in \mathbb{R}^{d \times d_h}$. Show that computing $Q = XW^{(q)}$ and then splitting Q along the feature dimension is equivalent to separately computing $Q_0 = XW_0^{(q)}$ and $Q_1 = XW_1^{(q)}$.

(b) Explain why the same argument applies to K and V , and conclude why chunking Q , K , and V after the large projection is equivalent to explicitly keeping track of separate per-head projection matrices.

5. Multi-Head Attention Forward Pass

Recall the multi-head attention equations:

$$\begin{aligned} Q_h &= XW_h^{(q)} \\ K_h &= XW_h^{(k)} \\ V_h &= XW_h^{(v)} \\ H_h &= \text{Attention}(Q_h, K_h, V_h) \\ Y &= \text{Concat}[H_0, \dots, H_{n_{\text{heads}}-1}]W^{(o)}. \end{aligned}$$

Adopting the block matrix convention introduced in the previous question, we can group together

$$\begin{aligned} Q &= [Q_0 \mid Q_1 \mid \dots \mid Q_{n_{\text{heads}}-1}] = X[W_0^{(q)} \mid W_1^{(q)} \mid \dots \mid W_{n_{\text{heads}}-1}^{(q)}] = XW^{(q)}, \\ K &= [K_0 \mid K_1 \mid \dots \mid K_{n_{\text{heads}}-1}] = X[W_0^{(k)} \mid W_1^{(k)} \mid \dots \mid W_{n_{\text{heads}}-1}^{(k)}] = XW^{(k)}, \end{aligned}$$

$$V = [V_0 \mid V_1 \mid \dots \mid V_{n_{\text{heads}}-1}] = X[W_0^{(v)} \mid W_1^{(v)} \mid \dots \mid W_{n_{\text{heads}}-1}^{(v)}] = XW^{(v)}.$$

In this question, we consider a toy example with $n_{\text{heads}} = 2$, sequence length $n = 2$, hidden dimension $d = 4$, and per-head hidden dimension $d_h = 2$. The input is not batched, so $X, Q, K, V \in \mathbb{R}^{2 \times 4}$.

(a) Suppose for a particular input X , the projected query, key, and value matrices are

$$Q = \begin{bmatrix} 1 & 2 & 3 & 4 \\ 5 & 6 & 7 & 8 \end{bmatrix}, \quad K = \begin{bmatrix} 2 & 0 & 1 & 1 \\ 1 & 3 & 0 & 2 \end{bmatrix}, \quad V = \begin{bmatrix} 1 & 0 & 2 & 1 \\ 3 & 1 & 0 & 2 \end{bmatrix}.$$

Compute the per-head query, key, and value matrices $Q_0, Q_1, K_0, K_1, V_0, V_1$ by splitting Q, K , and V along the feature dimension.

(b) Suppose for another input X , the attention outputs for the two heads are

$$H_0 = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}, \quad H_1 = \begin{bmatrix} 2 & 1 \\ 1 & 2 \end{bmatrix},$$

and suppose the output projection matrix is

$$W^{(o)} = \begin{bmatrix} 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 \\ 1 & 0 & 0 & 1 \\ 0 & 1 & 1 & 0 \end{bmatrix}.$$

Compute the concatenated matrix $H = \text{Concat}[H_0, H_1]$ and then compute the final output $Y =$

$HW^{(o)}$.