

CHAPTER

9

The Transformer

“The true art of memory is the art of attention ”

Samuel Johnson, *Idler #74*, September 1759

In this chapter we introduce the **transformer**, the standard architecture for building **large language models**. Transformer-based large language models have completely changed the field of speech and language processing. Indeed, every subsequent chapter in this textbook will make use of them. We’ll focus for now on left-to-right (sometimes called **causal** or autoregressive) language modeling, in which we are given a sequence of input tokens and predict output tokens one by one by conditioning on the prior context.

The transformer is a neural network with a specific structure that includes a mechanism called **self-attention** or **multi-head attention**.¹ Attention can be thought of as a way to build contextual representations of a token’s meaning by **attending to** and integrating information from surrounding tokens, helping the model learn how tokens relate to each other over large spans.

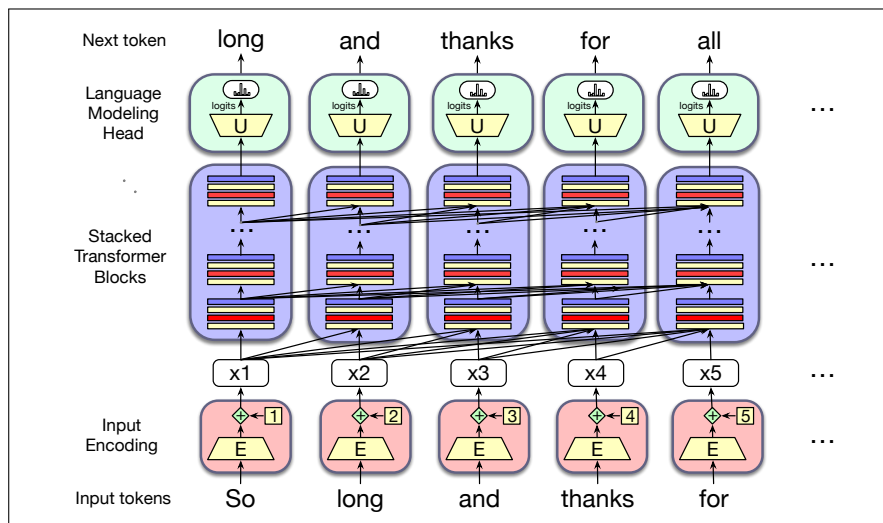


Figure 9.1 The architecture of a (left-to-right) transformer, showing how each input token get encoded, passed through a set of stacked transformer blocks, and then a language model head that predicts the next token.

Fig. 9.1 sketches the transformer architecture. A transformer has three major components. At the center are columns of transformer **blocks**. Each block is a multilayer network (a **multi-head attention** layer, feedforward networks and layer normalization steps) that maps an input vector \mathbf{x}_i in column i (corresponding to input

¹ Although multi-head attention developed historically from the **RNN attention** mechanism (Chapter 8), we’ll define attention from scratch here for readers who haven’t yet read Chapter 8.

token i) to an output vector \mathbf{h}_i . The set of n blocks maps an entire **context window** of input vectors $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ to a window of output vectors $(\mathbf{h}_1, \dots, \mathbf{h}_n)$ of the same length. A column might contain from 12 to 96 or more stacked blocks.

The column of blocks is preceded by the **input encoding** component, which processes an input token (like the word **thanks**) into a contextual vector representation, using an **embedding matrix E** and a mechanism for encoding token position. Each column is followed by a **language modeling head**, which takes the embedding output by the final transformer block, passes it through an **unembedding matrix U** and a softmax over the vocabulary to generate a single token for that column.

Transformer-based language models are complex, and so the details will unfold over the next 5 chapters. In the next sections we'll introduce multi-head attention, the rest of the transformer block, and the input encoding and language modeling head components. Chapter 10 discusses how language models are **pretrained**, and how tokens are generated via **sampling**. Chapter 11 introduces **masked language modeling** and the **BERT** family of bidirectional transformer encoder models. Chapter 12 shows how to **prompt** LLMs to perform NLP tasks by giving instructions and demonstrations, and how to **align** the model with human preferences. Chapter 13 will introduce machine translation with the **encoder-decoder** architecture.

9.1 Attention

Recall from Chapter 6 that for **word2vec** and other static embeddings, the representation of a word's meaning is always the same vector irrespective of the context: the word **chicken**, for example, is always represented by the same fixed vector. So a static vector for the word **it** might somehow encode that this is a pronoun used for animals and inanimate entities. But in context **it** has a much richer meaning. Consider **it** in one of these two sentences:

(9.1) **The chicken** didn't cross the road because **it** was too tired.

(9.2) **The chicken** didn't cross the road because **it** was too wide.

In (9.1) **it** is the chicken (i.e., the reader knows that the chicken was tired), while in (9.2) **it** is the road (and the reader knows that the road was wide).² That is, if we are to compute the meaning of this sentence, we'll need the meaning of **it** to be associated with **the chicken** in the first sentence and associated with **the road** in the second one, sensitive to the context.

Furthermore, consider reading left to right like a causal language model, processing the sentence up to the word **it**:

(9.3) **The chicken** didn't cross the road because **it**

At this point we don't yet know which thing **it** is going to end up referring to! So a representation of **it** at this point might have aspects of both **chicken** and **road** as the reader is trying to guess what happens next.

This fact that words have rich linguistic relationships with other words that may be far away pervades language. Consider two more examples:

(9.4) The **keys** to the cabinet **are** on the table.

(9.5) I walked along the **pond**, and noticed one of the trees along the **bank**.

² We say that in the first example **it** **corefers** with the chicken, and in the second **it** corefers with the road; we'll return to this in Chapter 23.

In (9.4), the phrase *The keys* is the subject of the sentence, and in English and many languages, must agree in grammatical number with the verb *are*; in this case both are plural. In English we can't use a singular verb like *is* with a plural subject like *keys* (we'll discuss agreement more in Chapter 18). In (9.5), we know that *bank* refers to the side of a pond or river and not a financial institution because of the context, including words like *pond*. (We'll discuss word senses more in Chapter 11.)

contextual
embeddings

The point of all these examples is that these contextual words that help us compute the meaning of words in context can be quite far away in the sentence or paragraph. Transformers can build contextual representations of word meaning, **contextual embeddings**, by integrating the meaning of these helpful contextual words. In a transformer, layer by layer, we build up richer and richer contextualized representations of the meanings of input tokens. At each layer, we compute the representation of a token i by combining information about i from the previous layer with information about the neighboring tokens to produce a contextualized representation for each word at each position.

Attention is the mechanism in the transformer that weighs and combines the representations from appropriate other tokens in the context from layer $k - 1$ to build the representation for tokens in layer k .

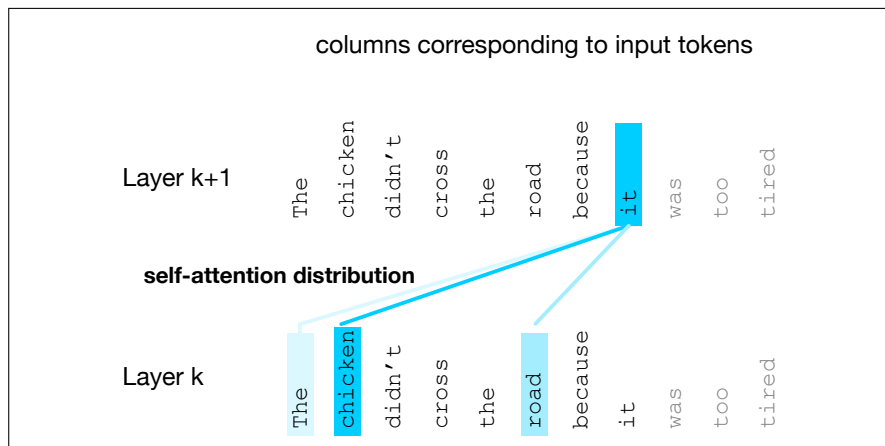


Figure 9.2 The self-attention weight distribution α that is part of the computation of the representation for the word *it* at layer $k + 1$. In computing the representation for *it*, we attend differently to the various words at layer l , with darker shades indicating higher self-attention values. Note that the transformer is attending highly to the columns corresponding to the tokens *chicken* and *road*, a sensible result, since at the point where *it* occurs, it could plausibly corefer with the chicken or the road, and hence we'd like the representation for *it* to draw on the representation for these earlier words. Figure adapted from [Uszkoreit \(2017\)](#).

Fig. 9.2 shows a schematic example simplified from a transformer ([Uszkoreit, 2017](#)). The figure describes the situation when the current token is *it* and we need to compute a contextual representation for this token at layer $k + 1$ of the transformer, drawing on the representations (from layer k) of every prior token. The figure uses color to represent the attention distribution over the contextual words: the tokens *chicken* and *road* both have a high attention weight, meaning that as we are computing the representation for *it*, we will draw most heavily on the representation for *chicken* and *road*. This will be useful in building the final representation for *it*, since *it* will end up coreferring with either *chicken* or *road*.

Let's now turn to how this attention distribution is represented and computed.

9.1.1 Attention more formally

As we've said, the attention computation is a way to compute a vector representation for a token at a particular layer of a transformer, by selectively attending to and integrating information from prior tokens at the previous layer. Attention takes an input representation \mathbf{x}_i corresponding to the input token at position i , and a context window of prior inputs $\mathbf{x}_1.. \mathbf{x}_{i-1}$, and produces an output \mathbf{a}_i .

In causal, left-to-right language models, the context is any of the prior words. That is, when processing \mathbf{x}_i , the model has access to \mathbf{x}_i as well as the representations of all the prior tokens in the context window (context windows consist of thousands of tokens) but no tokens after i . (By contrast, in Chapter 11 we'll generalize attention so it can also look ahead to future words.)

Fig. 9.3 illustrates this flow of information in an entire causal self-attention layer, in which this same attention computation happens in parallel at each token position i . Thus a self-attention layer maps input sequences $(\mathbf{x}_1, \dots, \mathbf{x}_n)$ to output sequences of the same length $(\mathbf{a}_1, \dots, \mathbf{a}_n)$.




Figure 9.3 Information flow in causal self-attention. When processing each input \mathbf{x}_i , the model attends to all the inputs up to, and including \mathbf{x}_i .

Simplified version of attention At its heart, attention is really just a weighted sum of context vectors, with a lot of complications added to how the weights are computed and what gets summed. For pedagogical purposes let's first describe a simplified intuition of attention, in which the attention output \mathbf{a}_i at token position i is simply the weighted sum of all the representations \mathbf{x}_j , for all $j \leq i$; we'll use α_{ij} to mean how much \mathbf{x}_j should contribute to \mathbf{a}_j :

$$\text{Simplified version: } \mathbf{a}_i = \sum_{j \leq i} \alpha_{ij} \mathbf{x}_j \quad (9.6)$$

Each α_{ij} is a scalar used for weighing the value of input \mathbf{x}_j when summing up the inputs to compute \mathbf{a}_i . How shall we compute this α weighting? In attention we weight each prior embedding proportionally to how **similar** it is to the current token i . So the output of attention is a sum of the embeddings of prior tokens weighted by their similarity with the current token embedding. We compute similarity scores via **dot product**, which maps two vectors into a scalar value ranging from $-\infty$ to ∞ . The larger the score, the more similar the vectors that are being compared. We'll normalize these scores with a softmax to create the vector of weights $\alpha_{ij}, j \leq i$.

$$\text{Simplified Version: } \text{score}(\mathbf{x}_i, \mathbf{x}_j) = \mathbf{x}_i \cdot \mathbf{x}_j \quad (9.7)$$

$$\alpha_{ij} = \text{softmax}(\text{score}(\mathbf{x}_i, \mathbf{x}_j)) \quad \forall j \leq i \quad (9.8)$$

Thus in Fig. 9.3 we compute \mathbf{a}_3 by computing three scores: $\mathbf{x}_3 \cdot \mathbf{x}_1$, $\mathbf{x}_3 \cdot \mathbf{x}_2$ and $\mathbf{x}_3 \cdot \mathbf{x}_3$, normalizing them by a softmax, and using the resulting probabilities as weights indicating each of their proportional relevance to the current position i . Of course,

the softmax weight will likely be highest for \mathbf{x}_i , since \mathbf{x}_i is very similar to itself, resulting in a high dot product. But other context words may also be similar to i , and the softmax will also assign some weight to those words. Then we use these weights as the α values in Eq. 9.6 to compute the weighted sum that is our \mathbf{a}_3 .

The simplified attention in equations 9.6 – 9.8 demonstrates the attention-based approach to computing \mathbf{a}_i : compare the \mathbf{x}_i to prior vectors, normalize those scores into a probability distribution used to weight the sum of the prior vector. But now we're ready to remove the simplifications.

attention head
head

A single attention head using query, key, and value matrices Now that we've seen a simple intuition of attention, let's introduce the actual **attention head**, the version of attention that's used in transformers. (The word **head** is often used in transformers to refer to specific structured layers). The attention head allows us to distinctly represent three different roles that each input embedding plays during the course of the attention process:

- query** • As *the current element* being compared to the preceding inputs. We'll refer to this role as a **query**.
- key** • In its role as a *preceding input* that is being compared to the current element to determine a similarity weight. We'll refer to this role as a **key**.
- value** • And finally, as a **value** of a preceding element that gets weighted and summed up to compute the output for the current element.

To capture these three different roles, transformers introduce weight matrices \mathbf{W}^Q , \mathbf{W}^K , and \mathbf{W}^V . These weights will project each input vector \mathbf{x}_i into a representation of its role as a key, query, or value:

$$\mathbf{q}_i = \mathbf{x}_i \mathbf{W}^Q; \quad \mathbf{k}_i = \mathbf{x}_i \mathbf{W}^K; \quad \mathbf{v}_i = \mathbf{x}_i \mathbf{W}^V \quad (9.9)$$

Given these projections, when we are computing the similarity of the current element \mathbf{x}_i with some prior element \mathbf{x}_j , we'll use the dot product between the current element's **query** vector \mathbf{q}_i and the preceding element's **key** vector \mathbf{k}_j . Furthermore, the result of a dot product can be an arbitrarily large (positive or negative) value, and exponentiating large values can lead to numerical issues and loss of gradients during training. To avoid this, we scale the dot product by a factor related to the size of the embeddings, via dividing by the square root of the dimensionality of the query and key vectors (d_k). We thus replace the simplified Eq. 9.7 with Eq. 9.11. The ensuing softmax calculation resulting in α_{ij} remains the same, but the output calculation for \mathbf{a}_i is now based on a weighted sum over the value vectors \mathbf{v} (Eq. 9.13).

Here's a final set of equations for computing self-attention for a single self-attention output vector \mathbf{a}_i from a single input vector \mathbf{x}_i . This version of attention computes \mathbf{a}_i by summing the *values* of the prior elements, each weighted by the similarity of its *key* to the *query* from the current element:

$$\mathbf{q}_i = \mathbf{x}_i \mathbf{W}^Q; \quad \mathbf{k}_j = \mathbf{x}_j \mathbf{W}^K; \quad \mathbf{v}_j = \mathbf{x}_j \mathbf{W}^V \quad (9.10)$$

$$\text{score}(\mathbf{x}_i, \mathbf{x}_j) = \frac{\mathbf{q}_i \cdot \mathbf{k}_j}{\sqrt{d_k}} \quad (9.11)$$

$$\alpha_{ij} = \text{softmax}(\text{score}(\mathbf{x}_i, \mathbf{x}_j)) \quad \forall j \leq i \quad (9.12)$$

$$\mathbf{a}_i = \sum_{j \leq i} \alpha_{ij} \mathbf{v}_j \quad (9.13)$$

We illustrate this in Fig. 9.4 for the case of calculating the value of the third output \mathbf{a}_3 in a sequence.

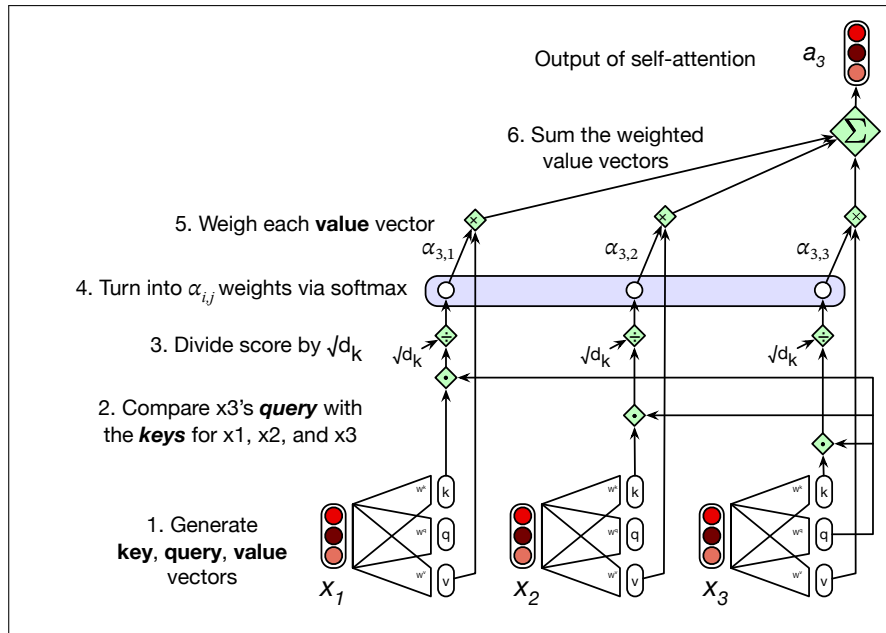


Figure 9.4 Calculating the value of a_3 , the third element of a sequence using causal (left-to-right) self-attention.

Let's talk shapes. The input to attention x_i and the output from attention a_i both have the same dimensionality $1 \times d$ (We often call d the **model dimensionality**, and indeed as we'll discuss in Section 9.2 the output h_i of each transformer block, as well as the intermediate vectors inside the transformer block also have the same dimensionality $1 \times d$).

We'll have a dimension d_k for the key and query vectors. The query vector and the key vector are both dimensionality $1 \times d_k$, so we can take their dot product $q_i \cdot k_j$. We'll have a separate dimension d_v for the value vectors. The transform matrix W^Q has shape $[d \times d_k]$, W^K is $[d \times d_k]$, and W^V is $[d \times d_v]$. In the original transformer work (Vaswani et al., 2017), d was 512, d_k and d_v were both 64.

Multi-head Attention Equations 9.11-9.13 describe a single **attention head**. But actually, transformers use multiple attention heads. The intuition is that each head might be attending to the context for different purposes: heads might be specialized to represent different linguistic relationships between context elements and the current token, or to look for particular kinds of patterns in the context.

multi-head attention

So in **multi-head attention** we have h separate attention heads that reside in parallel layers at the same depth in a model, each with its own set of parameters that allows the head to model different aspects of the relationships among inputs. Thus each head i in a self-attention layer has its own set of key, value, and query embeddings for each head: W^{K_i} , W^{Q_i} and W^{V_i} . These are used to project the inputs into separate key, value, and query embeddings for each head.

When using multiple heads the model dimension d is still used for the input and output, the key and query embeddings have dimensionality d_k , and the value embeddings are of dimensionality d_v (again, in the original transformer paper $d_k = d_v = 64$, $h = 8$, and $d = 512$). Thus for each head i , we have weight layers W^{Q_i} of shape $[d \times d_k]$, W^{K_i} of shape $[d \times d_k]$, and W^{V_i} of shape $[d \times d_v]$.

Below are the equations for attention augmented with multiple heads; Fig. 9.5

shows an intuition.

$$\mathbf{q}_i^c = \mathbf{x}_i \mathbf{W}^{Qc}; \quad \mathbf{k}_j^c = \mathbf{x}_j \mathbf{W}^{Kc}; \quad \mathbf{v}_j^c = \mathbf{x}_j \mathbf{W}^{Vc}; \quad \forall c \ 1 \leq c \leq h \quad (9.14)$$

$$\text{score}^c(\mathbf{x}_i, \mathbf{x}_j) = \frac{\mathbf{q}_i^c \cdot \mathbf{k}_j^c}{\sqrt{d_k}} \quad (9.15)$$

$$\alpha_{ij}^c = \text{softmax}(\text{score}^c(\mathbf{x}_i, \mathbf{x}_j)) \quad \forall j \leq i \quad (9.16)$$

$$\text{head}_i^c = \sum_{j \leq i} \alpha_{ij}^c \mathbf{v}_j^c \quad (9.17)$$

$$\mathbf{a}_i = (\text{head}^1 \oplus \text{head}^2 \dots \oplus \text{head}^h) \mathbf{W}^O \quad (9.18)$$

$$\text{MultiHeadAttention}(\mathbf{x}_i, [\mathbf{x}_1, \dots, \mathbf{x}_N]) = \mathbf{a}_i \quad (9.19)$$

The output of each of the h heads is of shape $1 \times d_v$, and so the output of the multi-head layer with h heads consists of h vectors of shape $1 \times d_v$. These are concatenated to produce a single output with dimensionality $1 \times hd_v$. Then we use yet another linear projection $\mathbf{W}^O \in \mathbb{R}^{hd_v \times d}$ to reshape it, resulting in the multi-head attention vector \mathbf{a}_i with the correct output shape $[1 \times d]$ at each input i .

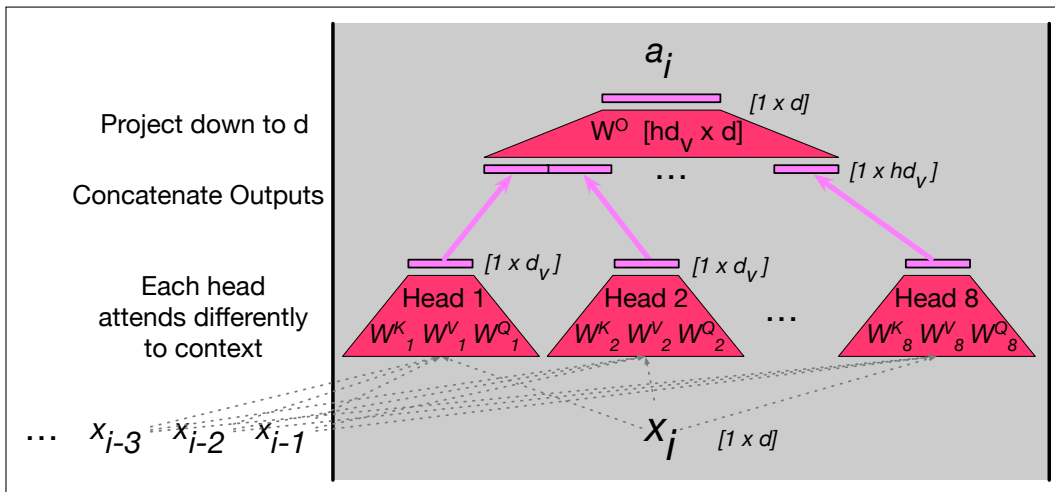


Figure 9.5 The multi-head attention computation for input \mathbf{x}_i , producing output \mathbf{a}_i . A multi-head attention layer has h heads, each with its own key, query and value weight matrices. The outputs from each of the heads are concatenated and then projected down to d , thus producing an output of the same size as the input.

9.2 Transformer Blocks

The self-attention calculation lies at the core of what’s called a transformer block, which, in addition to the self-attention layer, includes three other kinds of layers: (1) a feedforward layer, (2) residual connections, and (3) normalizing layers (colloquially called “layer norm”).

Fig. 9.6 illustrates a transformer block, sketching a common way of thinking about the block that is called the **residual stream** (Elhage et al., 2021). In the residual stream viewpoint, we consider the processing of an individual token i through the transformer block as a single stream of d -dimensional representations for token position i . This residual stream starts with the original input vector, and the various

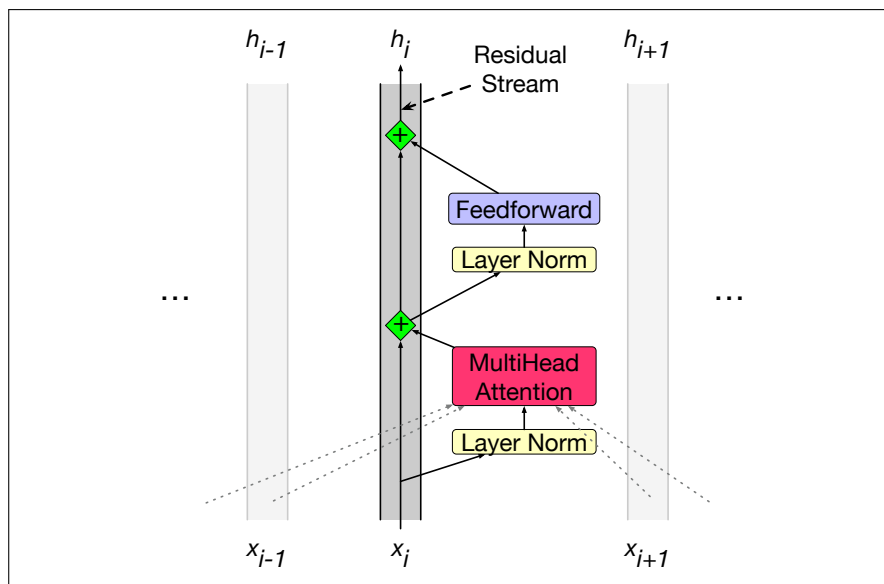


Figure 9.6 The architecture of a transformer block showing the **residual stream**. This figure shows the **prenorm** version of the architecture, in which the layer norms happen before the attention and feedforward layers rather than after.

components read their input from the residual stream and add their output back into the stream.

The input at the bottom of the stream is an embedding for a token, which has dimensionality d . This initial embedding gets passed up (by **residual connections**), and is progressively added to by the other components of the transformer: the **attention layer** that we have seen, and the **feedforward layer** that we will introduce. Before the attention and feedforward layer is a computation called the **layer norm**.

Thus the initial vector is passed through a layer norm and attention layer, and the result is added back into the stream, in this case to the original input vector \mathbf{x}_i . And then this summed vector is again passed through another layer norm and a feedforward layer, and the output of those is added back into the residual, and we'll use \mathbf{h}_i to refer to the resulting output of the transformer block for token i . (In earlier descriptions the residual stream was often described using a different metaphor as **residual connections** that add the input of a component to its output, but the residual stream is a more perspicuous way of visualizing the transformer.)

We've already seen the attention layer, so let's now introduce the feedforward and layer norm computations in the context of processing a single input \mathbf{x}_i at token position i .

Feedforward layer The feedforward layer is a fully-connected 2-layer network, i.e., one hidden layer, two weight matrices, as introduced in Chapter 7. The weights are the same for each token position i , but are different from layer to layer. It is common to make the dimensionality d_{ff} of the hidden layer of the feedforward network be larger than the model dimensionality d . (For example in the original transformer model, $d = 512$ and $d_{\text{ff}} = 2048$.)

$$\text{FFN}(\mathbf{x}_i) = \text{ReLU}(\mathbf{x}_i \mathbf{W}_1 + b_1) \mathbf{W}_2 + b_2 \quad (9.20)$$

Layer Norm At two stages in the transformer block we **normalize** the vector (Bayer et al., 2016). This process, called **layer norm** (short for layer normalization), is one

of many forms of normalization that can be used to improve training performance in deep neural networks by keeping the values of a hidden layer in a range that facilitates gradient-based training.

Layer norm is a variation of the **z-score** from statistics, applied to a single vector in a hidden layer. That is, the term layer norm is a bit confusing; layer norm is **not** applied to an entire transformer layer, but just to the embedding vector of a single token. Thus the input to layer norm is a single vector of dimensionality d and the output is that vector normalized, again of dimensionality d . The first step in layer normalization is to calculate the mean, μ , and standard deviation, σ , over the elements of the vector to be normalized. Given an embedding vector \mathbf{x} of dimensionality d , these values are calculated as follows.

$$\mu = \frac{1}{d} \sum_{i=1}^d x_i \quad (9.21)$$

$$\sigma = \sqrt{\frac{1}{d} \sum_{i=1}^d (x_i - \mu)^2} \quad (9.22)$$

Given these values, the vector components are normalized by subtracting the mean from each and dividing by the standard deviation. The result of this computation is a new vector with zero mean and a standard deviation of one.

$$\hat{\mathbf{x}} = \frac{(\mathbf{x} - \mu)}{\sigma} \quad (9.23)$$

Finally, in the standard implementation of layer normalization, two learnable parameters, γ and β , representing gain and offset values, are introduced.

$$\text{LayerNorm}(\mathbf{x}) = \gamma \frac{(\mathbf{x} - \mu)}{\sigma} + \beta \quad (9.24)$$

Putting it all together The function computed by a transformer block can be expressed by breaking it down with one equation for each component computation, using \mathbf{t} (of shape $[1 \times d]$) to stand for transformer and superscripts to demarcate each computation inside the block:

$$\mathbf{t}_i^1 = \text{LayerNorm}(\mathbf{x}_i) \quad (9.25)$$

$$\mathbf{t}_i^2 = \text{MultiHeadAttention}(\mathbf{t}_i^1, [\mathbf{x}_1^1, \dots, \mathbf{x}_N^1]) \quad (9.26)$$

$$\mathbf{t}_i^3 = \mathbf{t}_i^2 + \mathbf{x}_i \quad (9.27)$$

$$\mathbf{t}_i^4 = \text{LayerNorm}(\mathbf{t}_i^3) \quad (9.28)$$

$$\mathbf{t}_i^5 = \text{FFN}(\mathbf{t}_i^4) \quad (9.29)$$

$$\mathbf{h}_i = \mathbf{t}_i^5 + \mathbf{t}_i^3 \quad (9.30)$$

Notice that the only component that takes as input information from other tokens (other residual streams) is multi-head attention, which (as we see from (9.27)) looks at all the neighboring tokens in the context. The output from attention, however, is then added into this token's embedding stream. In fact, [Elhage et al. \(2021\)](#) show that we can view attention heads as literally moving information from the residual stream of a neighboring token into the current stream. The high-dimensional embedding space at each position thus contains information about the current token and about neighboring tokens, albeit in different subspaces of the vector space. Fig. 9.7 shows a visualization of this movement.




Figure 9.7 An attention head can move information from token A’s residual stream into token B’s residual stream.

Crucially, the input and output dimensions of transformer blocks are matched so they can be stacked. Each token vector \mathbf{x}_i at the input to the block has dimensionality d , and the output \mathbf{h}_i also has dimensionality d . Transformers for large language models stack many of these blocks, from 12 layers (used for the T5 or GPT-3-small language models) to 96 layers (used for GPT-3 large), to even more for more recent models. We’ll come back to this issue of stacking in a bit.

Equation (9.27) and following are just the equation for a single transformer block, but the residual stream metaphor goes through all the transformer layers, from the first transformer blocks to the 12th, in a 12-layer transformer. At the earlier transformer blocks, the residual stream is representing the current token. At the highest transformer blocks, the residual stream is usually representing the following token, since at the very end it’s being trained to predict the next token.

Once we stack many blocks, there is one more requirement: at the very end of the last (highest) transformer block, there is a single extra layer norm that is run on the last \mathbf{h}_i of each token stream (just below the language model head layer that we will define soon).³

9.3 Parallelizing computation using a single matrix \mathbf{X}

This description of multi-head attention and the rest of the transformer block has been from the perspective of computing a single output at a single time step i in a single residual stream. But as we pointed out earlier, the attention computation performed for each token to compute \mathbf{a}_i is independent of the computation for each other token, and that’s also true for all the computation in the transformer block computing \mathbf{h}_i from the input \mathbf{x}_i . That means we can easily parallelize the entire computation, taking advantage of efficient matrix multiplication routines.

We do this by packing the input embeddings for the N tokens of the input sequence into a single matrix \mathbf{X} of size $[N \times d]$. Each row of \mathbf{X} is the embedding of one token of the input. Transformers for large language models commonly have an input length $N = 1\text{K}, 2\text{K}$, or as many as 32K tokens (or more), so \mathbf{X} typically has between 1K and 32K rows, each of the dimensionality of the embedding d (the model

³ Note that we are using the most common current transformer architecture, which is called the **pre-norm** architecture. The original definition of the transformer in Vaswani et al. (2017) used an alternative architecture called the **post-norm** transformer in which the layer norm happens **after** the attention and FFN layers; it turns out moving the layer norm beforehand works better, but does require this one extra layer at the end.

dimension).

Parallelizing attention Let's first see this for a single attention head and then turn to multiple heads, and then add in the rest of the components in the transformer block. For one head we multiply \mathbf{X} by the key, query, and value matrices $\mathbf{W}^{\mathbf{Q}}$ of shape $[d \times d_k]$, $\mathbf{W}^{\mathbf{K}}$ of shape $[d \times d_k]$, and $\mathbf{W}^{\mathbf{V}}$ of shape $[d \times d_v]$, to produce matrices \mathbf{Q} of shape $[N \times d_k]$, $\mathbf{K} \in \mathbb{R}^{N \times d_k}$, and $\mathbf{V} \in \mathbb{R}^{N \times d_v}$, containing all the key, query, and value vectors:

$$\mathbf{Q} = \mathbf{X}\mathbf{W}^{\mathbf{Q}}; \mathbf{K} = \mathbf{X}\mathbf{W}^{\mathbf{K}}; \mathbf{V} = \mathbf{X}\mathbf{W}^{\mathbf{V}} \quad (9.31)$$

Given these matrices we can compute all the requisite query-key comparisons simultaneously by multiplying \mathbf{Q} and $\mathbf{K}^{\mathbf{T}}$ in a single matrix multiplication. The product is of shape $N \times N$, visualized in Fig. 9.8.

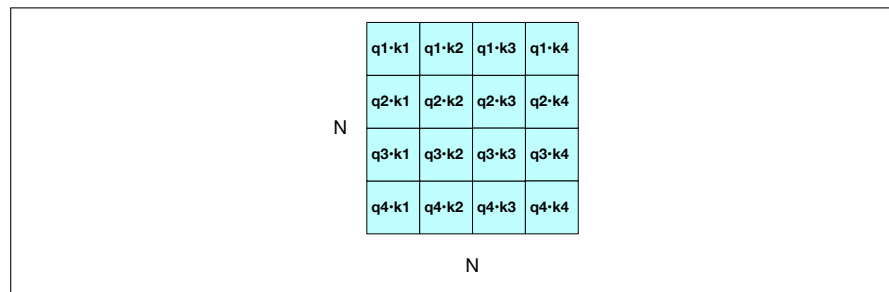


Figure 9.8 The $N \times N$ $\mathbf{QK}^{\mathbf{T}}$ matrix showing how it computes all $q_i \cdot k_j$ comparisons in a single matrix multiple.

Once we have this $\mathbf{QK}^{\mathbf{T}}$ matrix, we can very efficiently scale these scores, take the softmax, and then multiply the result by \mathbf{V} resulting in a matrix of shape $N \times d$: a vector embedding representation for each token in the input. We've reduced the entire self-attention step for an entire sequence of N tokens for one head to the following computation:

$$\mathbf{A} = \text{softmax} \left(\text{mask} \left(\frac{\mathbf{QK}^{\mathbf{T}}}{\sqrt{d_k}} \right) \right) \mathbf{V} \quad (9.32)$$

Masking out the future You may have noticed that we introduced a mask function in Eq. 9.32 above. This is because the self-attention computation as we've described it has a problem: the calculation in $\mathbf{QK}^{\mathbf{T}}$ results in a score for each query value to every key value, *including those that follow the query*. This is inappropriate in the setting of language modeling: guessing the next word is pretty simple if you already know it! To fix this, the elements in the upper-triangular portion of the matrix are zeroed out (set to $-\infty$), thus eliminating any knowledge of words that follow in the sequence. This is done in practice by adding a mask matrix M in which $M_{ij} = -\infty \forall j > i$ (i.e. for the upper-triangular portion) and $M_{ij} = 0$ otherwise. Fig. 9.9 shows the resulting masked $\mathbf{QK}^{\mathbf{T}}$ matrix. (we'll see in Chapter 11 how to make use of words in the future for tasks that need it).

Fig. 9.10 shows a schematic of all the computations for a single attention head parallelized in matrix form.

Fig. 9.8 and Fig. 9.9 also make it clear that attention is quadratic in the length of the input, since at each layer we need to compute dot products between each pair of tokens in the input. This makes it expensive to compute attention over very long documents (like entire novels). Nonetheless modern large language models manage to use quite long contexts of thousands or tens of thousands of tokens.

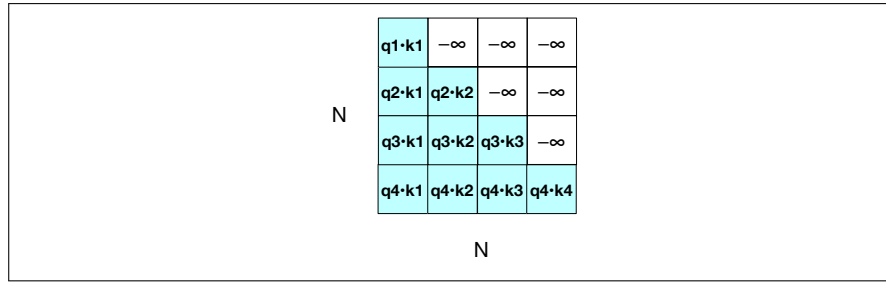


Figure 9.9 The $N \times N$ \mathbf{QK}^T matrix showing the $q_i \cdot k_j$ values, with the upper-triangular portion of the comparisons matrix zeroed out (set to $-\infty$, which the softmax will turn to zero).

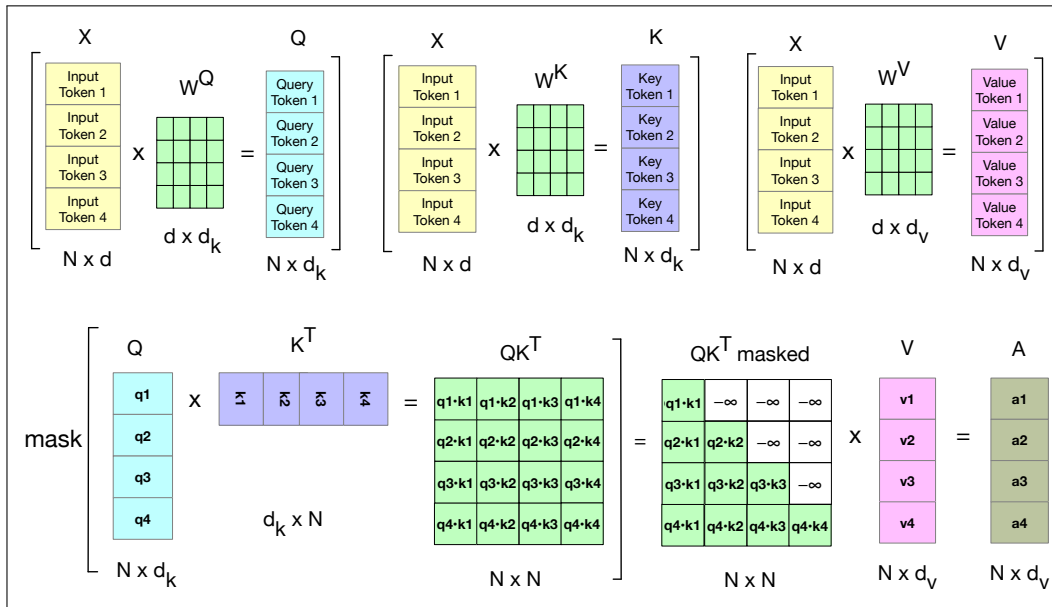


Figure 9.10 Schematic of the attention computation for a single attention head in parallel. The first row shows the computation of the \mathbf{Q} , \mathbf{K} , and \mathbf{V} matrices. The second row shows the computation of \mathbf{QK}^T , the masking (the softmax computation and the normalizing by dimensionality are not shown) and then the weighted sum of the value vectors to get the final attention vectors.

Parallelizing multi-head attention In multi-head attention, as with self-attention, the input and output have the model dimension d , the key and query embeddings have dimensionality d_k , and the value embeddings are of dimensionality d_v (again, in the original transformer paper $d_k = d_v = 64$, $h = 8$, and $d = 512$). Thus for each head i , we have weight layers $\mathbf{W}_i^Q \in \mathbb{R}^{d \times d_k}$, $\mathbf{W}_i^K \in \mathbb{R}^{d \times d_k}$, and $\mathbf{W}_i^V \in \mathbb{R}^{d \times d_v}$, and these get multiplied by the inputs packed into \mathbf{X} to produce $\mathbf{Q} \in \mathbb{R}^{N \times d_k}$, $\mathbf{K} \in \mathbb{R}^{N \times d_k}$, and $\mathbf{V} \in \mathbb{R}^{N \times d_v}$. The output of each of the h heads is of shape $N \times d_v$, and so the output of the multi-head layer with h heads consists of h matrices of shape $N \times d_v$. To make use of these matrices in further processing, they are concatenated to produce a single output with dimensionality $N \times hd_v$. Finally, we use yet another linear projection $\mathbf{W}^O \in \mathbb{R}^{hd_v \times d}$, that reshape it to the original output dimension for each token. Multiplying the concatenated $N \times hd_v$ matrix output by $\mathbf{W}^O \in \mathbb{R}^{hd_v \times d}$ yields

the self-attention output \mathbf{A} of shape $[N \times d]$.

$$\mathbf{Q}^i = \mathbf{XW}^{\mathbf{Q}^i}; \quad \mathbf{K}^i = \mathbf{XW}^{\mathbf{K}^i}; \quad \mathbf{V}^i = \mathbf{XW}^{\mathbf{V}^i} \quad (9.33)$$

$$\text{head}_i = \text{SelfAttention}(\mathbf{Q}^i, \mathbf{K}^i, \mathbf{V}^i) = \text{softmax}\left(\frac{\mathbf{Q}^i \mathbf{K}^{i\top}}{\sqrt{d_k}}\right) \mathbf{V}^i \quad (9.34)$$

$$\text{MultiHeadAttention}(\mathbf{X}) = (\text{head}_1 \oplus \text{head}_2 \dots \oplus \text{head}_h) \mathbf{W}^{\mathbf{O}} \quad (9.35)$$

Putting it all together with the parallel input matrix \mathbf{X} The function computed in parallel by an entire layer of N transformer block over the entire N input tokens can be expressed as:

$$\mathbf{O} = \text{LayerNorm}(\mathbf{X} + \text{MultiHeadAttention}(\mathbf{X})) \quad (9.36)$$

$$\mathbf{H} = \text{LayerNorm}(\mathbf{O} + \text{FFN}(\mathbf{O})) \quad (9.37)$$

Or we can break it down with one equation for each component computation, using \mathbf{T} (of shape $[N \times d]$) to stand for transformer and superscripts to demarcate each computation inside the block:

$$\mathbf{T}^1 = \text{MultiHeadAttention}(\mathbf{X}) \quad (9.38)$$

$$\mathbf{T}^2 = \mathbf{X} + \mathbf{T}^1 \quad (9.39)$$

$$\mathbf{T}^3 = \text{LayerNorm}(\mathbf{T}^2) \quad (9.40)$$

$$\mathbf{T}^4 = \text{FFN}(\mathbf{T}^3) \quad (9.41)$$

$$\mathbf{T}^5 = \mathbf{T}^4 + \mathbf{T}^3 \quad (9.42)$$

$$\mathbf{H} = \text{LayerNorm}(\mathbf{T}^5) \quad (9.43)$$

Here when we use a notation like $\text{FFN}(\mathbf{T}^3)$ we mean that the same FFN is applied in parallel to each of the N embedding vectors in the window. Similarly, each of the N tokens is normed in parallel in the LayerNorm. Crucially, the input and output dimensions of transformer blocks are matched so they can be stacked. Since each token x_i at the input to the block has dimensionality d , that means the input \mathbf{X} and output \mathbf{H} are both of shape $[N \times d]$.

9.4 The input: embeddings for token and position

embedding

Let's talk about where the input \mathbf{X} comes from. Given a sequence of N tokens (N is the context length in tokens), the matrix \mathbf{X} of shape $[N \times d]$ has an **embedding** for each word in the context. The transformer does this by separately computing two embeddings: an input token embedding, and an input positional embedding.

A token embedding, introduced in Chapter 7 and Chapter 8, is a vector of dimension d that will be our initial representation for the input token. (As we pass vectors up through the transformer layers in the residual stream, this embedding representation will change and grow, incorporating context and playing a different role depending on the kind of language model we are building.) The set of initial embeddings are stored in the embedding matrix \mathbf{E} , which has a row for each of the $|V|$ tokens in the vocabulary. Thus each word is a row vector of d dimensions, and \mathbf{E} has shape $[|V| \times d]$.

Given an input token string like *Thanks for all the* we first convert the tokens into vocabulary indices (these were created when we first tokenized the input using

BPE or SentencePiece). So the representation of *thanks for all the* might be $\mathbf{w} = [5, 4000, 10532, 2224]$. Next we use indexing to select the corresponding rows from \mathbf{E} , (row 5, row 4000, row 10532, row 2224).

one-hot vector

Another way to think about selecting token embeddings from the embedding matrix is to represent tokens as one-hot vectors of shape $[1 \times |V|]$, i.e., with one dimension for each word in the vocabulary. Recall that in a **one-hot vector** all the elements are 0 except one, the element whose dimension is the word’s index in the vocabulary, which has value 1. So if the word “thanks” has index 5 in the vocabulary, $x_5 = 1$, and $x_i = 0 \forall i \neq 5$, as shown here:

$$\begin{bmatrix} 0 & 0 & 0 & 0 & 1 & 0 & 0 & \dots & 0 & 0 & 0 & 0 \end{bmatrix}$$

$$\begin{matrix} 1 & 2 & 3 & 4 & 5 & 6 & 7 & \dots & \dots & |V| \end{matrix}$$

Multiplying by a one-hot vector that has only one non-zero element $x_i = 1$ simply selects out the relevant row vector for word i , resulting in the embedding for word i , as depicted in Fig. 9.11.




Figure 9.11 Selecting the embedding vector for word V_5 by multiplying the embedding matrix \mathbf{E} with a one-hot vector with a 1 in index 5.

We can extend this idea to represent the entire token sequence as a matrix of one-hot vectors, one for each of the N positions in the transformer’s context window, as shown in Fig. 9.12.




Figure 9.12 Selecting the embedding matrix for the input sequence of token ids W by multiplying a one-hot matrix corresponding to W by the embedding matrix \mathbf{E} .

These token embeddings are not position-dependent. To represent the position of each token in the sequence, we combine these token embeddings with **positional embeddings** specific to each position in an input sequence.

positional embeddings

Where do we get these positional embeddings? The simplest method, called **absolute position**, is to start with randomly initialized embeddings corresponding to each possible input position up to some maximum length. For example, just as we have an embedding for the word *fish*, we’ll have an embedding for the position 3. As with word embeddings, these positional embeddings are learned along with other parameters during training. We can store them in a matrix E_{pos} of shape $[1 \times N]$.

absolute position

To produce an input embedding that captures positional information, we just add the word embedding for each input to its corresponding positional embedding. The individual token and position embeddings are both of size $[1 \times d]$, so their sum is

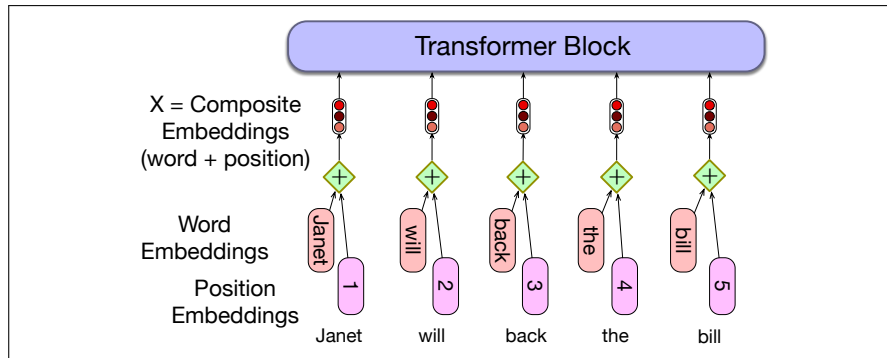


Figure 9.13 A simple way to model position: add an embedding of the absolute position to the token embedding to produce a new embedding of the same dimensionality.

also $[1 \times d]$. This new embedding serves as the input for further processing. Fig. 9.13 shows the idea.

The final representation of the input, the matrix \mathbf{X} , is an $[N \times d]$ matrix in which each row i is the representation of the i th token in the input, computed by adding $\mathbf{E}[id(i)]$ —the embedding of the id of the token that occurred at position i —, to $\mathbf{P}[i]$, the positional embedding of position i .

A potential problem with the simple absolute position embedding approach is that there will be plenty of training examples for the initial positions in our inputs and correspondingly fewer at the outer length limits. These latter embeddings may be poorly trained and may not generalize well during testing. An alternative approach to absolute positional embeddings is to choose a static function that maps integer inputs to real-valued vectors in a way that captures the inherent relationships among the positions. That is, it captures the fact that position 4 in an input is more closely related to position 5 than it is to position 17. A combination of sine and cosine functions with differing frequencies was used in the original transformer work. Even more complex positional embedding methods exist, such as ones that represent **relative position** instead of absolute position, often implemented in the attention mechanism at each layer rather than being added once at the initial input.

9.5 The Language Modeling Head

language
modeling head
head

The last component of the transformer we must introduce is the **language modeling head**. Here we are using the word **head** to mean the additional neural circuitry we add on top of the basic transformer architecture when we apply pretrained transformer models to various tasks. The language modeling head is the circuitry we need to do language modeling.

Recall that language models, from the simple n-gram models of Chapter 3 through the feedforward and RNN language models of Chapter 7 and Chapter 8, are word predictors. Given a context of words, they assign a probability to each possible next word. For example, if the preceding context is “Thanks for all the” and we want to know how likely the next word is “fish” we would compute:

$$P(\text{fish}|\text{Thanks for all the})$$

Language models give us the ability to assign such a conditional probability to every possible next word, giving us a distribution over the entire vocabulary. The n-gram

language models of Chapter 3 compute the probability of a word given counts of its occurrence with the $n - 1$ prior words. The context is thus of size $n - 1$. For transformer language models, the context is the size of the transformer’s context window, which can be quite large: 2K, 4K, even 32K tokens for very large models.

The job of the language modeling head is to take the output of the final transformer layer from the last token N and use it to predict the upcoming word at position $N + 1$. Fig. 9.14 shows how to accomplish this task, taking the output of the last token at the last layer (the d -dimensional output embedding of shape $[1 \times d]$) and producing a probability distribution over words (from which we will choose one to generate).

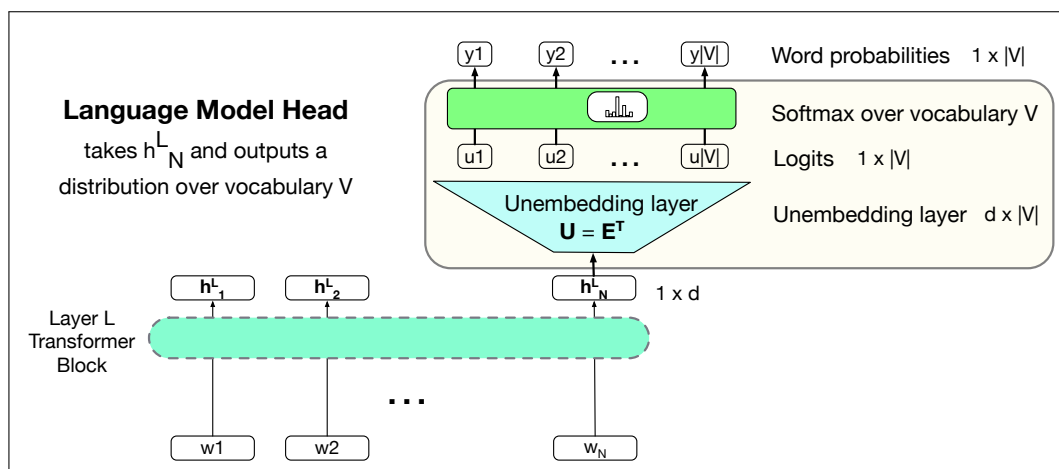


Figure 9.14 The language modeling head: the circuit at the top of a transformer that maps from the output embedding for token N from the last transformer layer (h_N^L) to a probability distribution over words in the vocabulary V .

The first module in Fig. 9.14 is a linear layer, whose job is to project from the output h_N^L , which represents the output token embedding at position N from the final block L , (hence of shape $[1 \times d]$) to the **logit** vector, or score vector, that will have a single score for each of the $|V|$ possible words in the vocabulary V . The logit vector \mathbf{u} is thus of dimensionality $1 \times |V|$.

This linear layer can be learned, but more commonly we tie this matrix to (the transpose of) the embedding matrix \mathbf{E} . Recall that in **weight tying**, we use the same weights for two different matrices in the model. Thus at the input stage of the transformer the embedding matrix (of shape $[|V| \times d]$) is used to map from a one-hot vector over the vocabulary (of shape $[1 \times |V|]$) to an embedding (of shape $[1 \times d]$). And then in the language model head, \mathbf{E}^T , the transpose of the embedding matrix (of shape $[d \times |V|]$) is used to map back from an embedding (shape $[1 \times d]$) to a vector over the vocabulary (shape $[1 \times |V|]$). In the learning process, \mathbf{E} will be optimized to be good at doing both of these mappings. We therefore sometimes call the transpose \mathbf{E}^T the **unembedding** layer because it is performing this reverse mapping.

A softmax layer turns the logits \mathbf{u} into the probabilities \mathbf{y} over the vocabulary.

$$\mathbf{u} = \mathbf{h}_N^L \mathbf{E}^T \quad (9.44)$$

$$\mathbf{y} = \text{softmax}(\mathbf{u}) \quad (9.45)$$

We can use these probabilities to do things like help assign a probability to a given text. But the most important usage to generate text, which we do by **sampling**

a word from these probabilities y . We might sample the highest probability word (‘greedy’ decoding), or use another of the sampling methods we’ll introduce in Section ?? . In either case, whatever entry y_k we choose from the probability vector y , we generate the word that has that index k .




Figure 9.15 A transformer language model (decoder-only), stacking transformer blocks and mapping from an input token w_i to a predicted next token w_{i+1} .

Fig. 9.15 shows the total stacked architecture for one token i . Note that the input to each transformer layer x_i is the same as the output from the preceding layer h_i^- .

logit lens

Now that we see all these transformer layers spread out on the page, we can point out another useful feature of the unembedding layer: as a tool for interpretability of the internals of the transformer that we call the **logit lens** (Nostalgebraist, 2020). We can take a vector from any layer of the transformer and, pretending that it is the prefinal embedding, simply multiply it by the unembedding layer to get logits, and compute a softmax to see the distribution over words that that vector might be representing. This can be a useful window into the internal representations of the model. Since the network wasn’t trained to make the internal representations function in this way, the logit lens doesn’t always work perfectly, but this can still be a useful trick.

decoder-only model

A terminological note before we conclude: You will sometimes see a transformer used for this kind of unidirectional causal language model called a **decoder-only model**. This is because this model constitutes roughly half of the **encoder-**

decoder model for transformers that we'll see how to apply to machine translation in Chapter 13. (Confusingly, the original introduction of the transformer had an encoder-decoder architecture, and it was only later that the standard paradigm for causal language model was defined by using only the decoder part of this original architecture).

9.6 Summary

This chapter has introduced the transformer and its components for the task of language modeling. We'll continue the task of language modeling including issues like training and sampling in the next chapter.

Here's a summary of the main points that we covered:

- Transformers are non-recurrent networks based on **multi-head attention**, a kind of **self-attention**. A multi-head attention computation takes an input vector \mathbf{x}_i and maps it to an output \mathbf{a}_i by adding in vectors from prior tokens, weighted by how relevant they are for the processing of the current word.
- A **transformer block** consists of a **residual stream** in which the input from the prior layer is passed up to the next layer, with the output of different components added to it. These components include a **multi-head attention layer** followed by a **feedforward layer**, each preceded by **layer normalizations**. Transformer blocks are stacked to make deeper and more powerful networks.
- The input to a transformer is a computing by adding an embedding (computed with an **embedding matrix**) to a **positional encoding** that represents the sequential position of the token in the window.
- Language models can be built out of stacks of transformer blocks, with a **language model head** at the top, which applies an **unembedding** matrix to the output \mathbf{H} of the top layer to generate the **logits**, which are then passed through a softmax to generate word probabilities.
- Transformer-based language models have a wide context window (as wide as 32768 tokens for very large models) allowing them to draw on enormous amounts of context to predict upcoming words.

Bibliographical and Historical Notes

The transformer (Vaswani et al., 2017) was developed drawing on two lines of prior research: **self-attention** and **memory networks**.

Encoder-decoder attention, the idea of using a soft weighting over the encodings of input words to inform a generative decoder (see Chapter 13) was developed by Graves (2013) in the context of handwriting generation, and Bahdanau et al. (2015) for MT. This idea was extended to self-attention by dropping the need for separate encoding and decoding sequences and instead seeing attention as a way of weighting the tokens in collecting information passed from lower layers to higher layers (Ling et al., 2015; Cheng et al., 2016; Liu et al., 2016).

Other aspects of the transformer, including the terminology of key, query, and value, came from **memory networks**, a mechanism for adding an external read-

write memory to networks, by using an embedding of a query to match keys representing content in an associative memory ([Sukhbaatar et al., 2015](#); [Weston et al., 2015](#); [Graves et al., 2014](#)).

MORE HISTORY TBD IN NEXT DRAFT.

- Ba, J. L., J. R. Kiros, and G. E. Hinton. 2016. [Layer normalization](#). *NeurIPS workshop*.
- Bahdanau, D., K. H. Cho, and Y. Bengio. 2015. Neural machine translation by jointly learning to align and translate. *ICLR 2015*.
- Cheng, J., L. Dong, and M. Lapata. 2016. [Long short-term memory-networks for machine reading](#). *EMNLP*.
- Elhage, N., N. Nanda, C. Olsson, T. Henighan, N. Joseph, B. Mann, A. Askell, Y. Bai, A. Chen, T. Conerly, N. Das-Sarma, D. Drain, D. Ganguli, Z. Hatfield-Dodds, D. Hernandez, A. Jones, J. Kernion, L. Lovitt, K. Ndousse, D. Amodei, T. Brown, J. Clark, J. Kaplan, S. McCandlish, and C. Olah. 2021. [A mathematical framework for transformer circuits](#). White paper.
- Graves, A. 2013. [Generating sequences with recurrent neural networks](#). ArXiv.
- Graves, A., G. Wayne, and I. Danihelka. 2014. [Neural Turing machines](#). ArXiv.
- Ling, W., C. Dyer, A. W. Black, I. Trancoso, R. Fernandez, S. Amir, L. Marujo, and T. Luis. 2015. [Finding function in form: Compositional character models for open vocabulary word representation](#). *EMNLP*.
- Liu, Y., C. Sun, L. Lin, and X. Wang. 2016. [Learning natural language inference using bidirectional LSTM model and inner-attention](#). ArXiv.
- Nostalgebraist. 2020. [Interpreting gpt: the logit lens](#). White paper.
- Sukhbaatar, S., A. Szlam, J. Weston, and R. Fergus. 2015. End-to-end memory networks. *NeurIPS*.
- Uszkoreit, J. 2017. [Transformer: A novel neural network architecture for language understanding](#). Google Research blog post, Thursday August 31, 2017.
- Vaswani, A., N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, Ł. Kaiser, and I. Polosukhin. 2017. [Attention is all you need](#). *NeurIPS*.
- Weston, J., S. Chopra, and A. Bordes. 2015. [Memory networks](#). *ICLR 2015*.