Deep Learning & Generative AI in Healthcare

Session 08

Neural Network Model of Language

- A neural probabilistic language model (Y. Bengio, et al.)
- Fixed window is small
- No window is large enough

We need a neural architecture that can process *any length input*

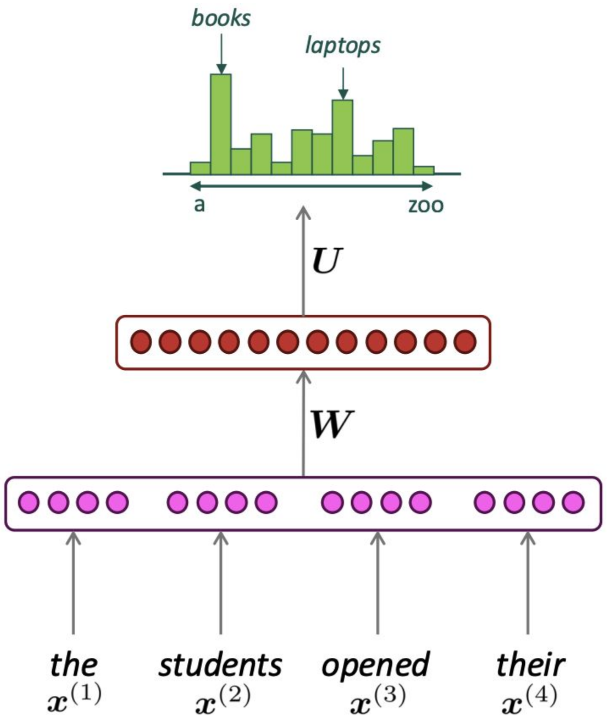

Recurrent Neural Networks

output distribution

$$\hat{\boldsymbol{y}}^{(t)} = \mathrm{softmax}\left(\boldsymbol{U}\boldsymbol{h}^{(t)} + \boldsymbol{b}_2\right) \in \mathbb{R}^{|V|}$$

hidden states

$$\boldsymbol{h}^{(t)} = \sigma\left(\boldsymbol{W}_h\boldsymbol{h}^{(t-1)} + \boldsymbol{W}_e\boldsymbol{e}^{(t)} + \boldsymbol{b}_1\right)$$

$\boldsymbol{h}^{(0)}$ is the initial hidden state

word embeddings

$$\boldsymbol{e}^{(t)} = \boldsymbol{E}\boldsymbol{x}^{(t)}$$

words / one-hot vectors

$$\boldsymbol{x}^{(t)} \in \mathbb{R}^{|V|}$$

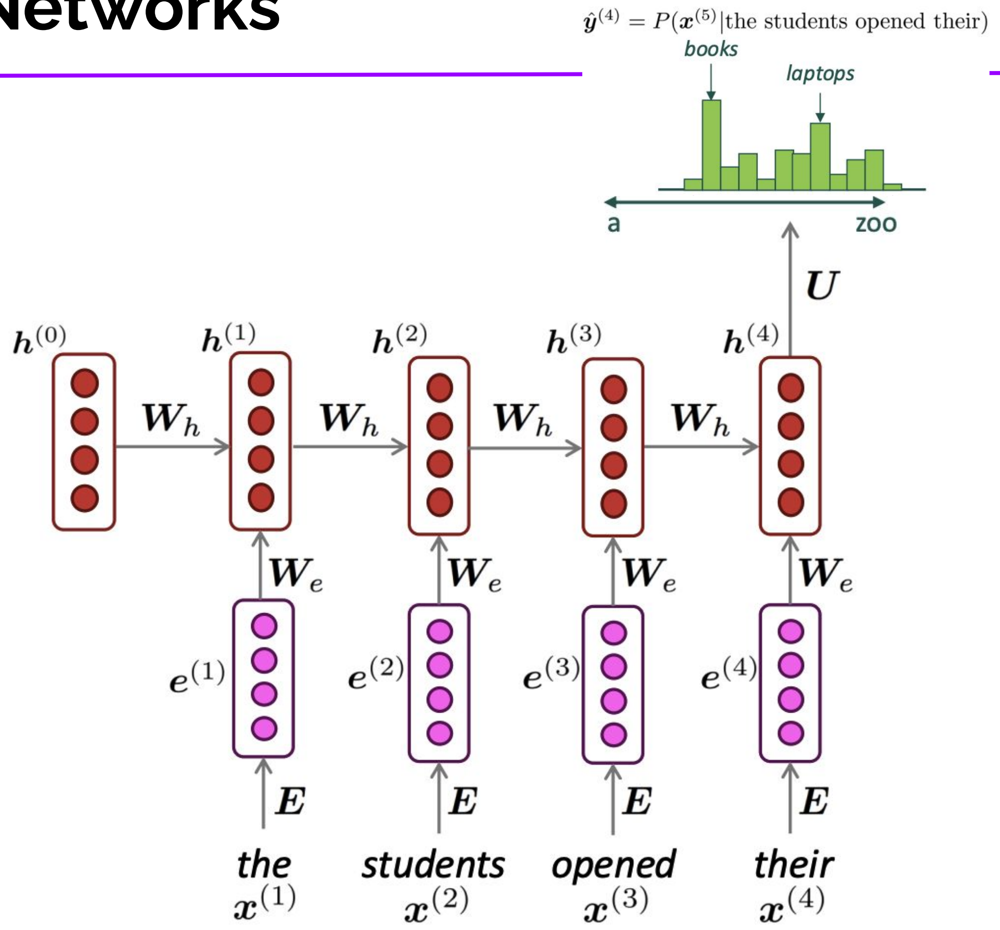

Recurrent Neural Networks

- Advantages
  - The process any length!
  - Can use information from previous steps
  - Model size does not increase for longer input context
  - Same weights applied on every timestep
- Disadvantages
  - Slow
  - Difficult to access information from many steps back

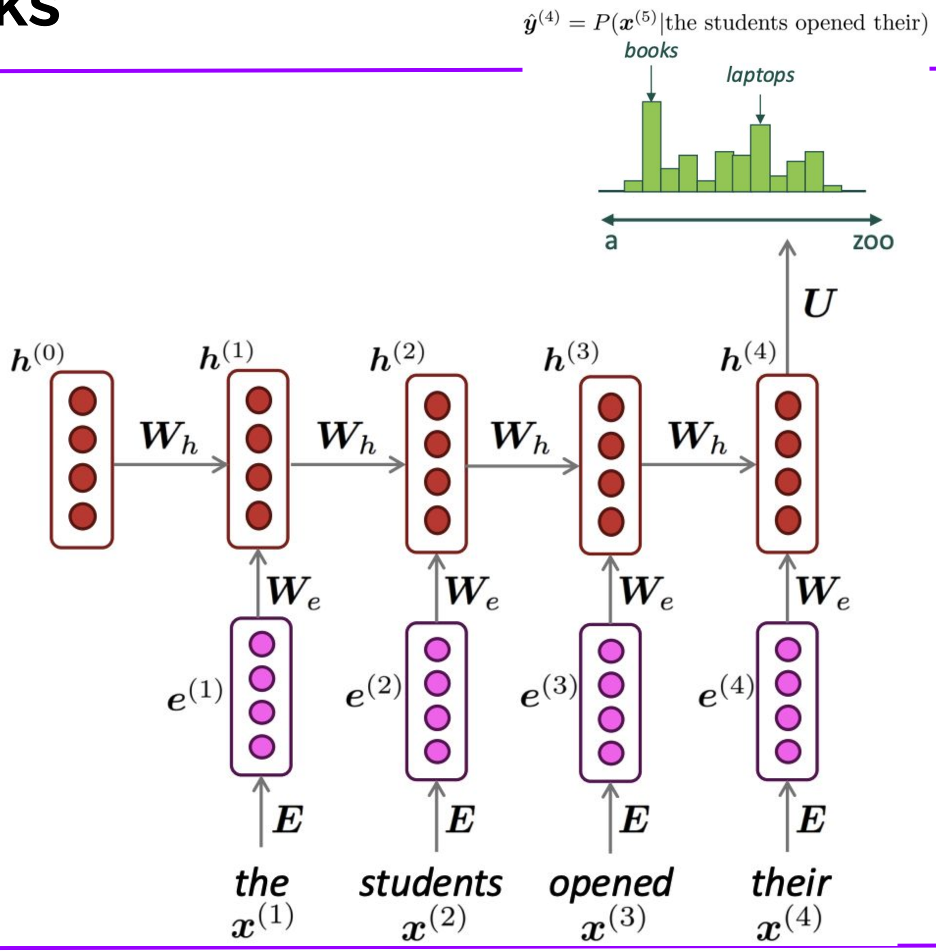

Attention is a solution!

- Attention provides a solution to the bottleneck problem!
- Core idea: on each step of the decoder, use direct connection to the encoder to focus on a particular part of the source sequence!
- In attention, the query matches all keys softly, to a weight between 0 and 1. The key's values are multiplied by the weights and summed!

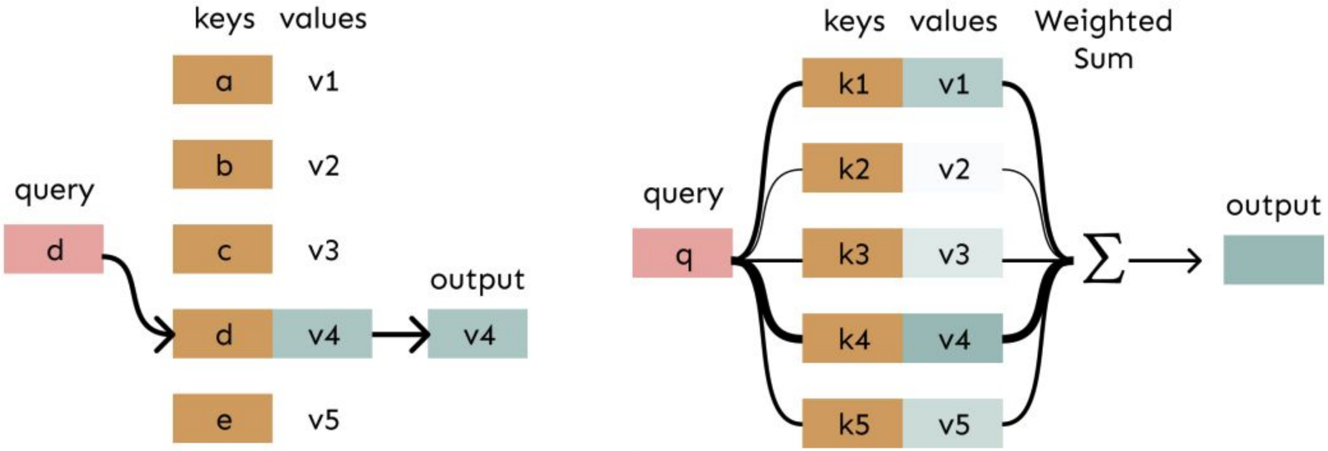

Attention in RNNs

- On each step of the decoder, use direct connection to the encoder to focus on a particular part of the source sequence.

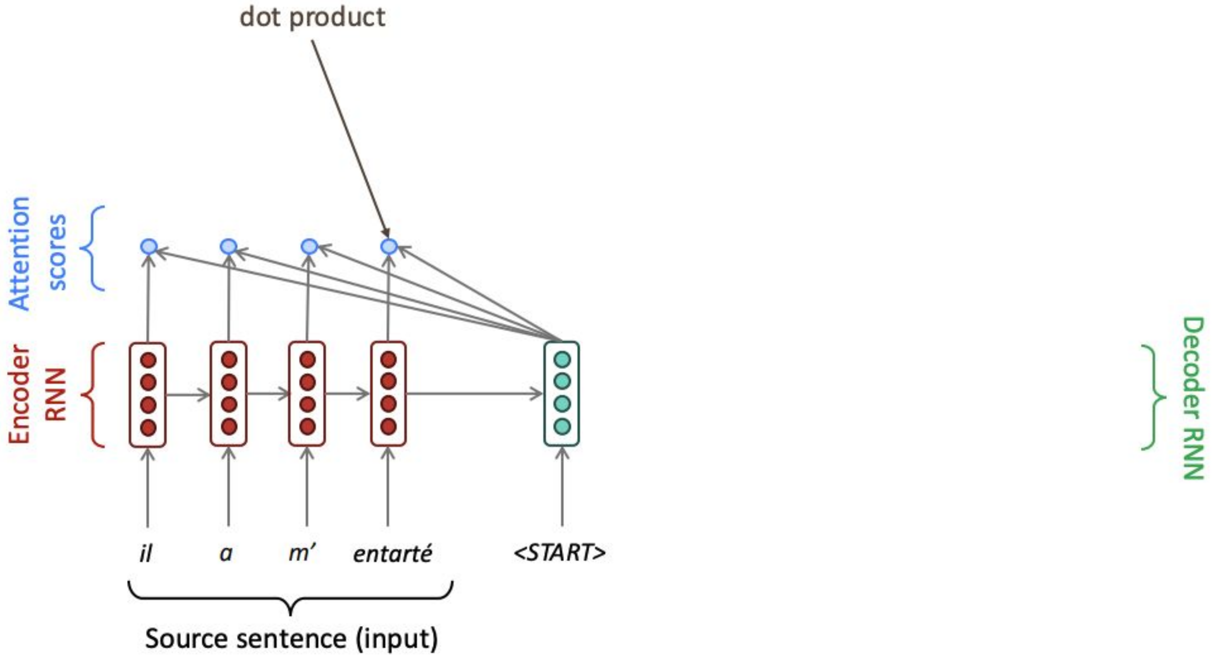

Attention in RNNs

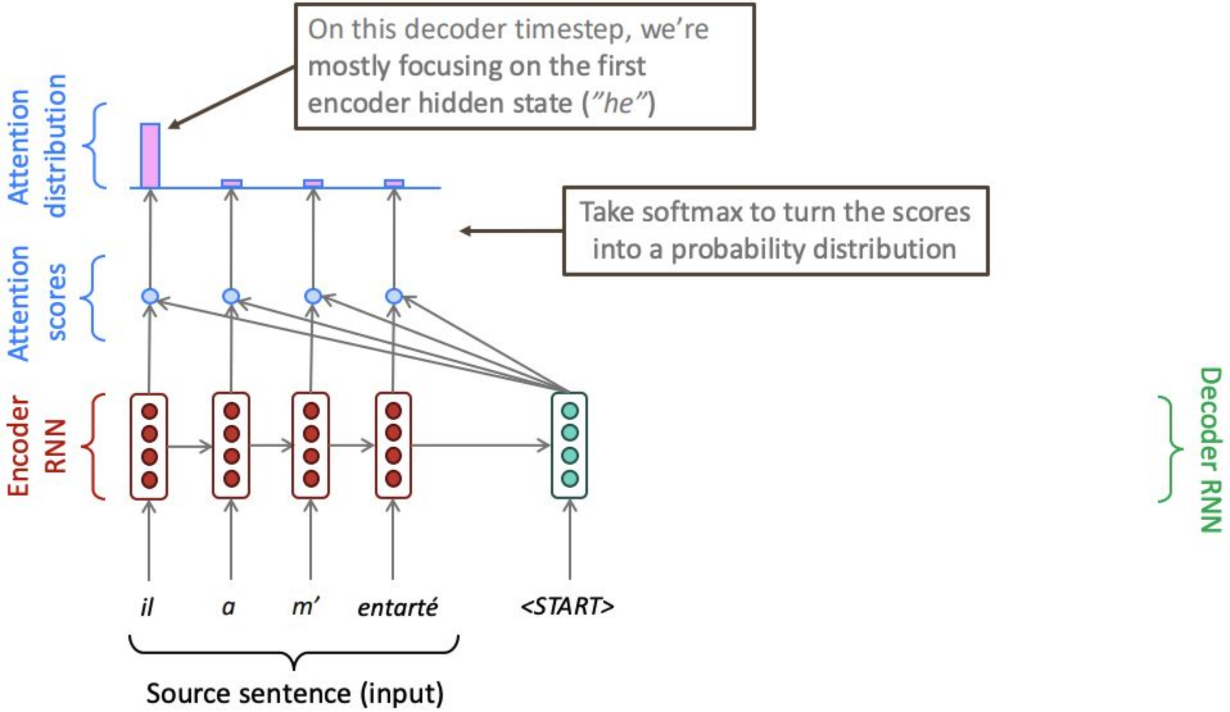

Attention in RNNs

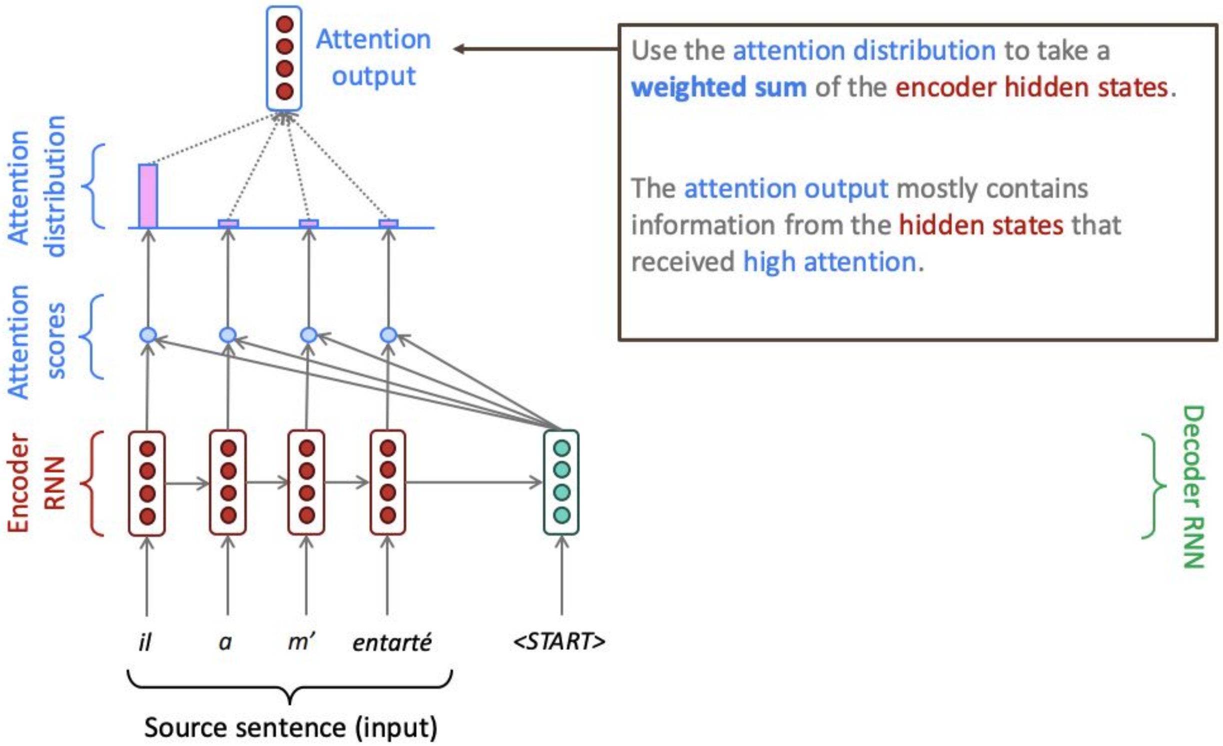

Attention in RNNs

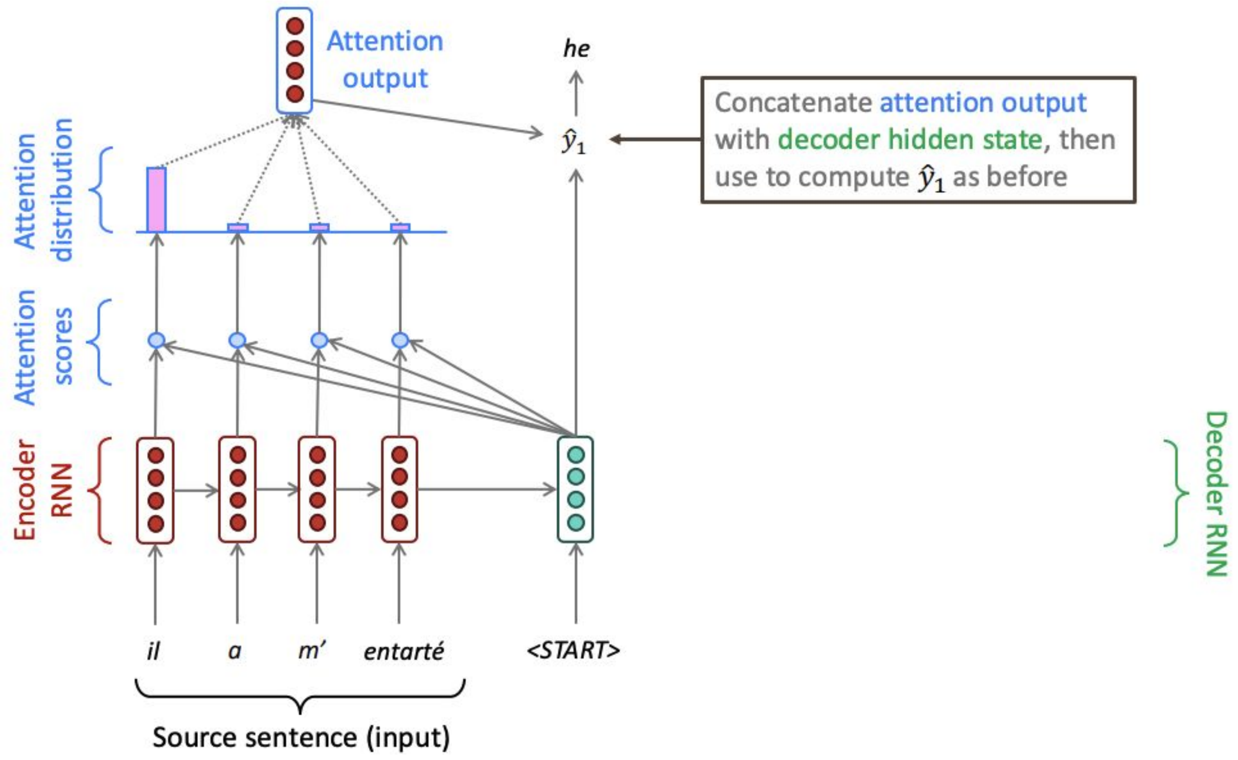

Attention in RNNs

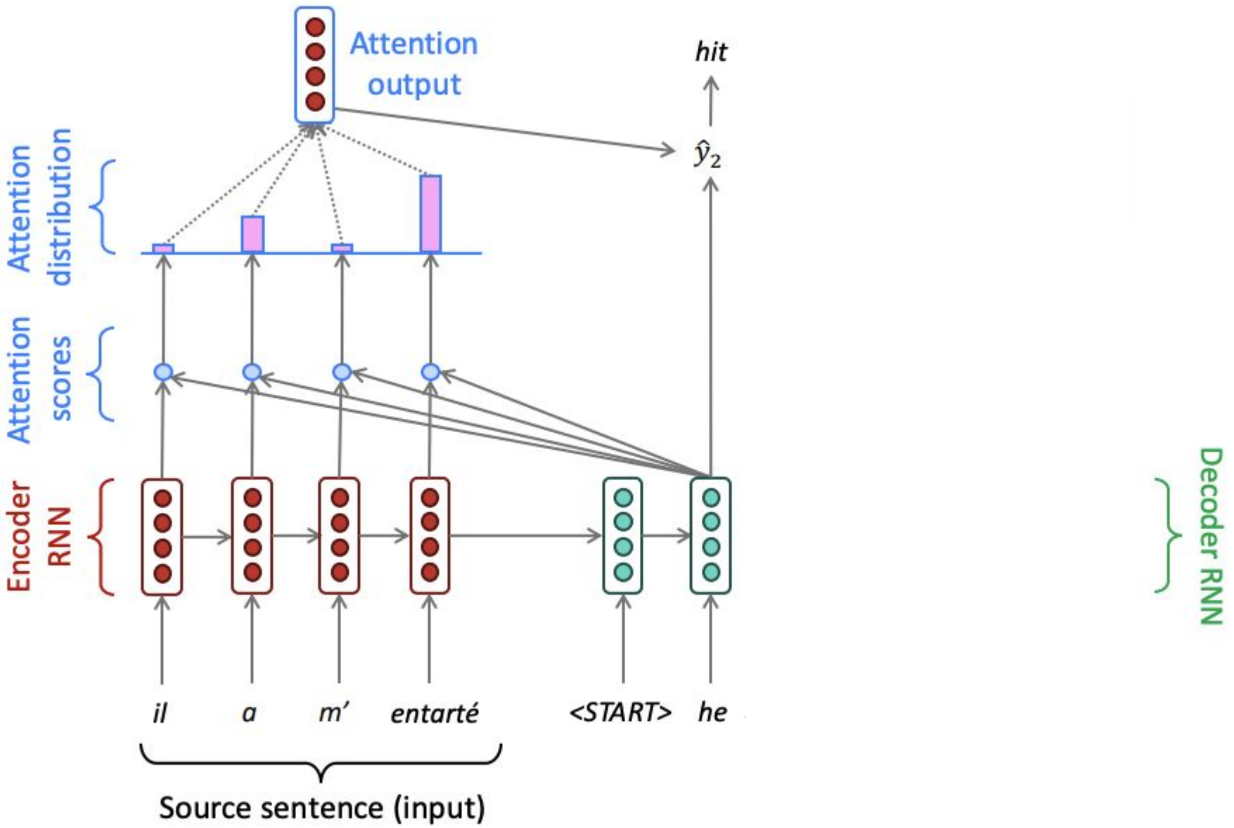

Attention in RNNs

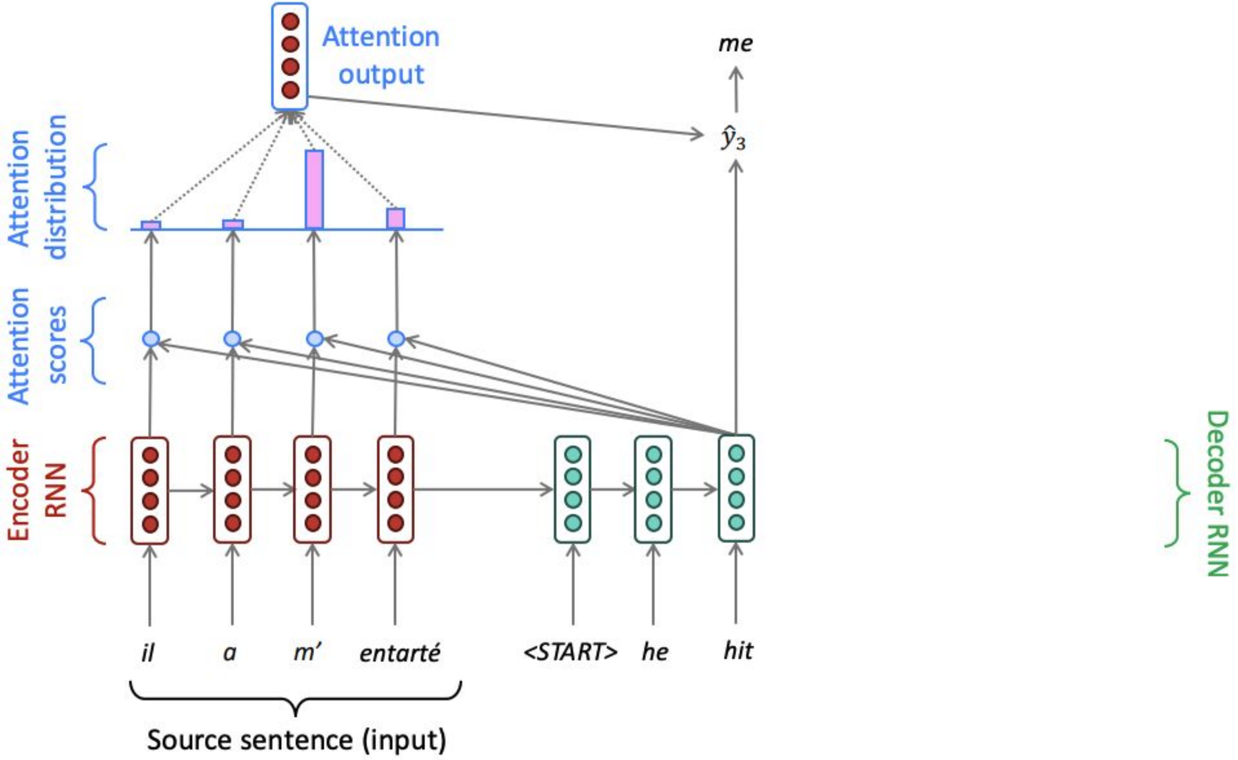

Is Recurrent Necessary at All?

- Abstractly: Attention is a way to pass information from a sequence (x) to a neural network input. (h_t)
- This is also exactly what RNNs are used for – to pass information!
- Can we just get rid of the RNN entirely? Maybe attention is just a better way to pass information!
- The building block we need is self Attention!
- So far we saw cross-attention!

Attention in RNNs

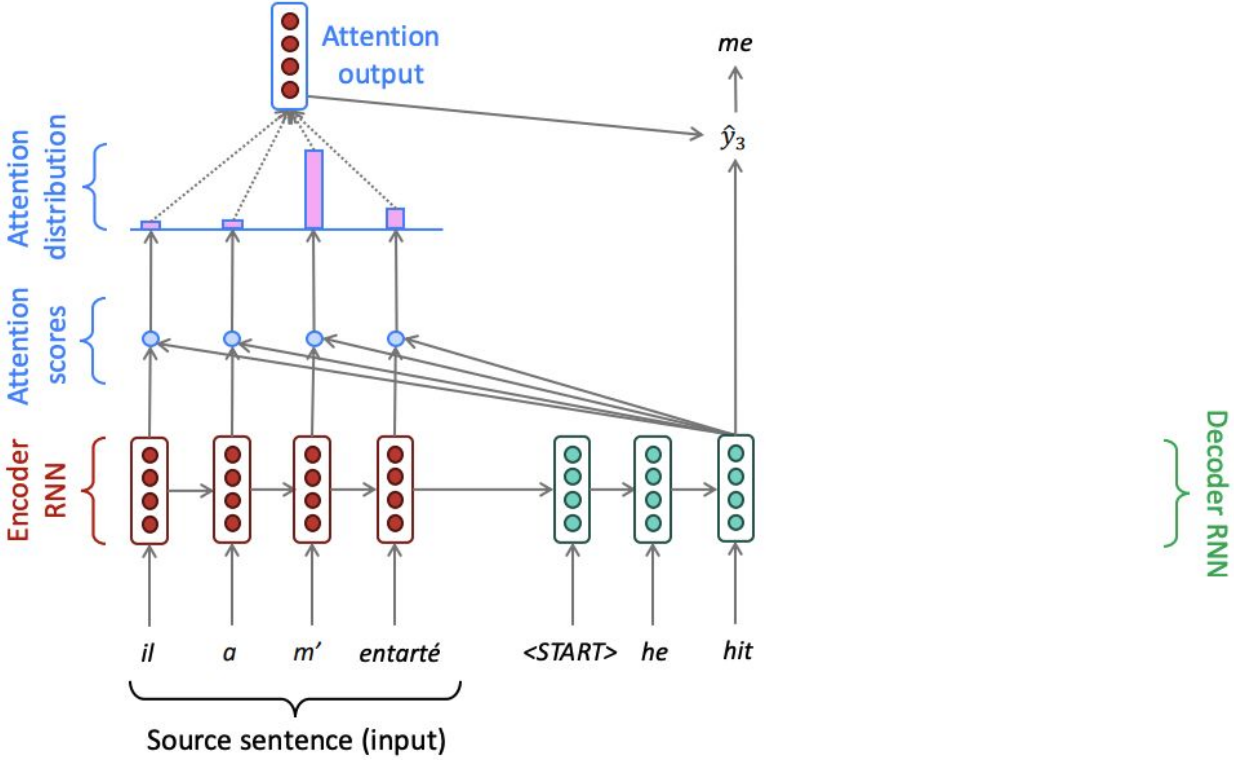

Self-attention: Keys, Queries, and Values

Let $\boldsymbol{w}_{1:n}$ be a sequence of words in vocabulary V, like *Zuko made his uncle tea*.

For each $\boldsymbol{w}_i$, let $\boldsymbol{x}_i = E\boldsymbol{w}_{\boldsymbol{i}}$, where $E \in \mathbb{R}^{d\times|V|}$ is an embedding matrix.

1. Transform each word embedding with weight matrices $\mathrm{Q}, \mathrm{K}, \mathrm{V}$, each in $\mathbb{R}^{d\times d}$

$\boldsymbol{q}_i = Q\boldsymbol{x}_{\boldsymbol{i}}$ (queries) $\boldsymbol{k}_i = K\boldsymbol{x}_{\boldsymbol{i}}$ (keys) $\boldsymbol{v}_i = V\boldsymbol{x}_{\boldsymbol{i}}$ (values)

2. Compute pairwise similarities between keys and queries; normalize with softmax

$$\boldsymbol{e}_{ij} = \boldsymbol{q}_{\boldsymbol{i}}^{\top}\boldsymbol{k}_{\boldsymbol{j}} \qquad \boldsymbol{\alpha}_{ij} = \frac{\exp(\boldsymbol{e}_{ij})}{\sum_{j'}\exp(\boldsymbol{e}_{ij'})}$$

3. Compute output for each word as weighted sum of values

$$\boldsymbol{o}_i = \sum_{\boldsymbol{j}} \boldsymbol{\alpha}_{ij}\,\boldsymbol{v}_i$$

Positional Embedding in Self-Attention

- Since self-attention doesn't build in order information, we need to encode the order of the sentence in our keys, queries, and values
- Consider representing each sequence index as a vector

$\boldsymbol{p}_i \in \mathbb{R}^d$, for $i \in \{1,2,\ldots,n\}$ are position vectors

$$\widetilde{\boldsymbol{x}}_i = \boldsymbol{x}_i + \boldsymbol{p}_i$$

In deep self-attention networks, we do this at the first layer! You could concatenate them as well, but people mostly just add...

Sinusoida Positional Embedding

- Sinusoidal position representations: concatenate sinusoidal functions of varying periods
- Periodicity indicates that maybe “absolute position” isn’t as important
- It can extrapolate to longer sequences as periods restart!

$$\boldsymbol{p}_i = \begin{pmatrix} \sin(i/10000^{2*1/d}) \\ \cos(i/10000^{2*1/d}) \\ \vdots \\ \sin(i/10000^{2*\frac{d}{2}/d}) \\ \cos(i/10000^{2*\frac{d}{2}/d}) \end{pmatrix}$$

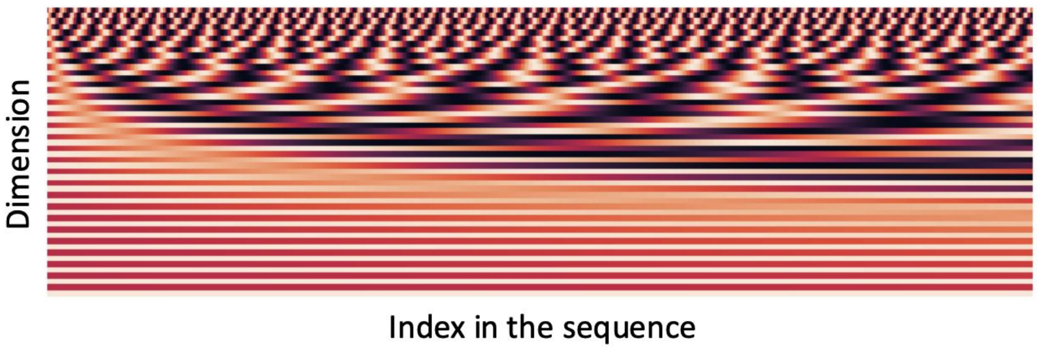

Non-Linearity in Self-Attention

- Easy fix: add a feed-forward network to post-process each output vector.

$$
\begin{aligned}
m_i &= MLP(\text{output}_i) \\
&= W_2 * \text{ReLU}(W_1\ \text{output}_i + b_1) + b_2
\end{aligned}
$$

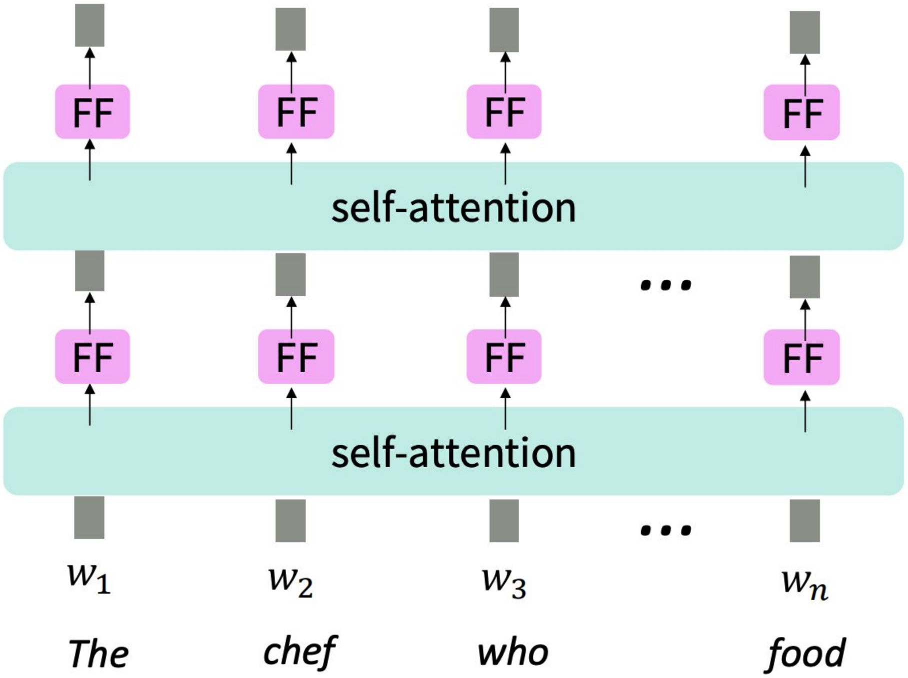

Causal Masking in Self-Attention

- For causality, we need to ensure not to peek at the future.
- At each timestep, we could change the set of keys and queries to only include past words!

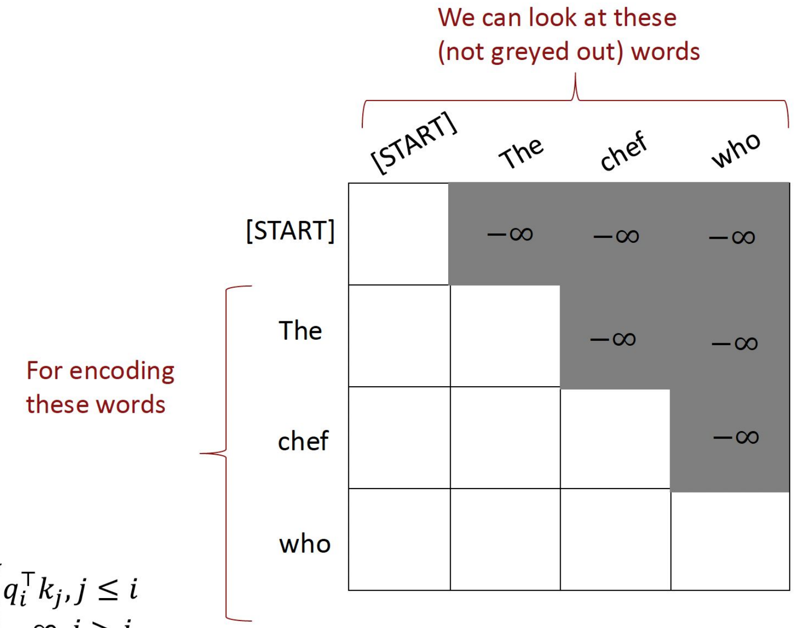

$$\begin{cases} q_i^\top k_j, j \leq i \\ -\infty, j > i \end{cases}$$

Multi-Head Self-Attention Layer

- The Attention module splits its Query, Key, and Value parameters N-ways and passes each split independently through a separate head.
- Calculations are combined together to produce a final attention score.
- Greater power to encode multiple relationships and nuances for each word.

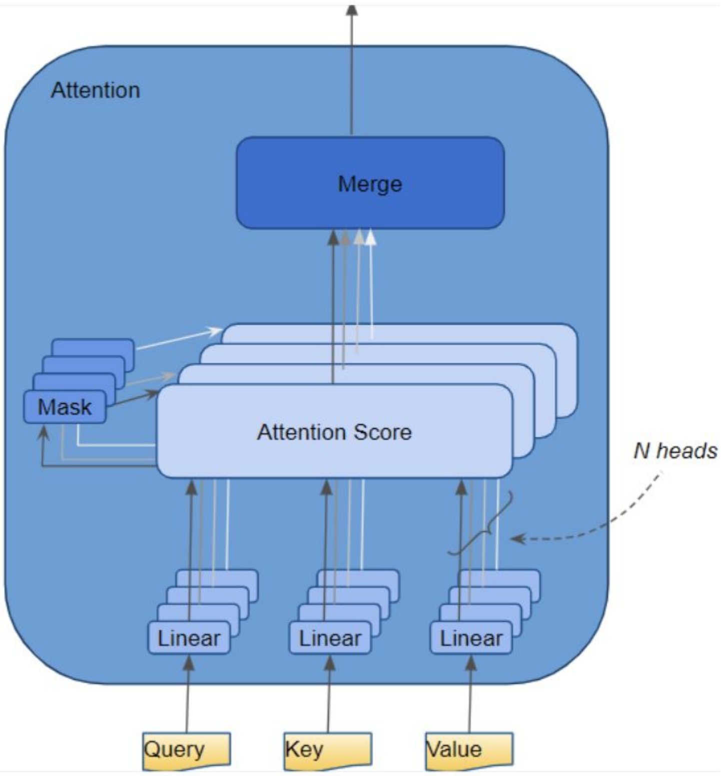

Residual Connections

- A trick to help models learn better!
- Gradient is 1 through residual connection
- Bias toward identity function.

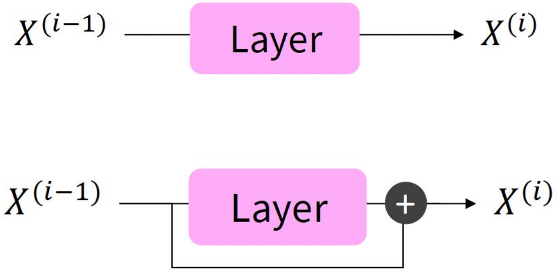

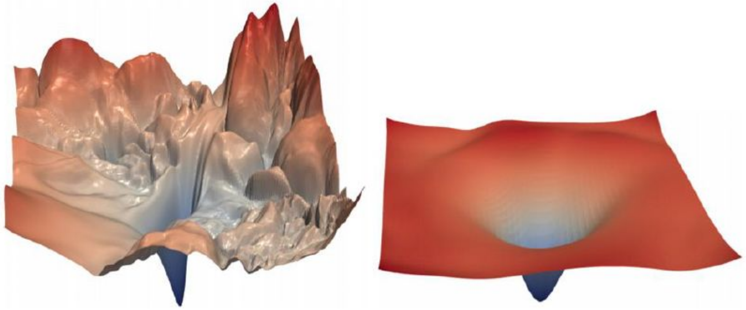

[no residuals] [residuals]

[Loss landscape visualization, Li et al., 2018, on a ResNet]

Layer Normalization

- A trick to help models train faster.
- Cut down on uninformative variation in hidden vectors by normalizing to unit mean and standard deviation within each layer.

$$\text{output} = \frac{x - \mu}{\sqrt{\sigma} + \epsilon} * \gamma + \beta$$

Normalize by scalar mean and variance

Modulate by learned elementwise gain and bias

Transformer Encoder

- Position representation
  - Specify the sequence order, since self-attention is an unordered function of its inputs.
- Nonlinearities
  - Frequently implemented as a simple feedforward network.
- Masking
  - Keep information about the future from "leaking" to the past.

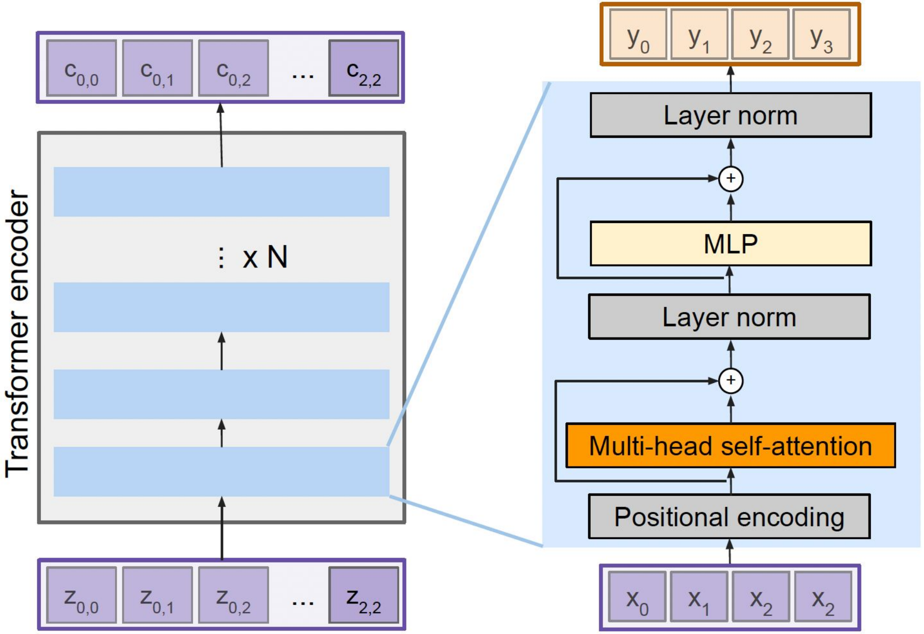

Transformer Decoder

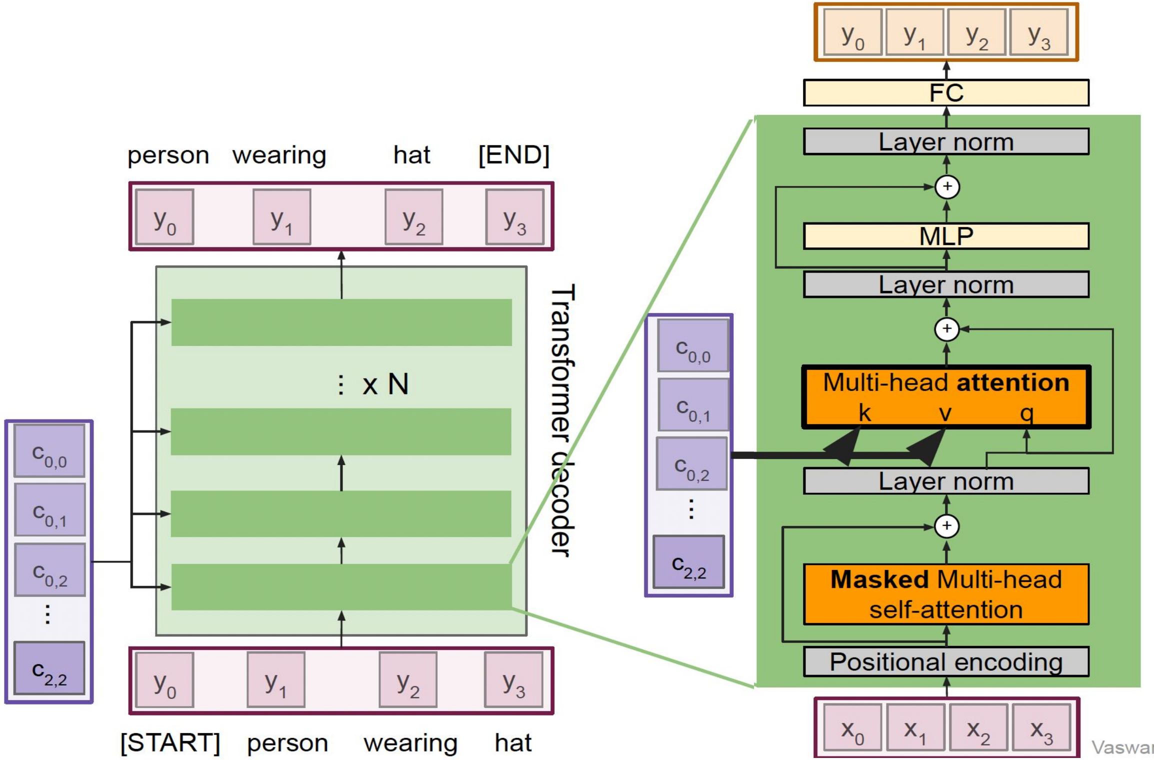

Multi-head attention block attends over the transformer encoder outputs.

For image captions, this is how we inject image features into the decoder.

Vaswani et al, "Attention is all you need", NeurIPS 2017

Cross Attention

- Self-attention:
  - Keys, queries, and values from same source
- Cross-attention
  - The keys and values are from encoder (like a memory)
  - The queries are from the decoder

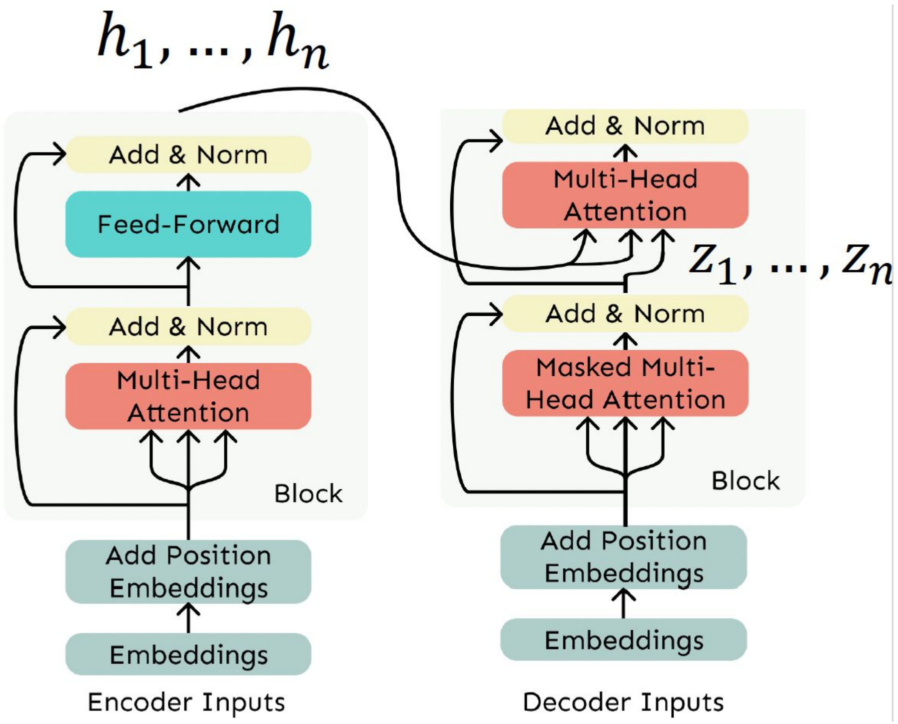


Inside an
LLM

Vision-Language Models

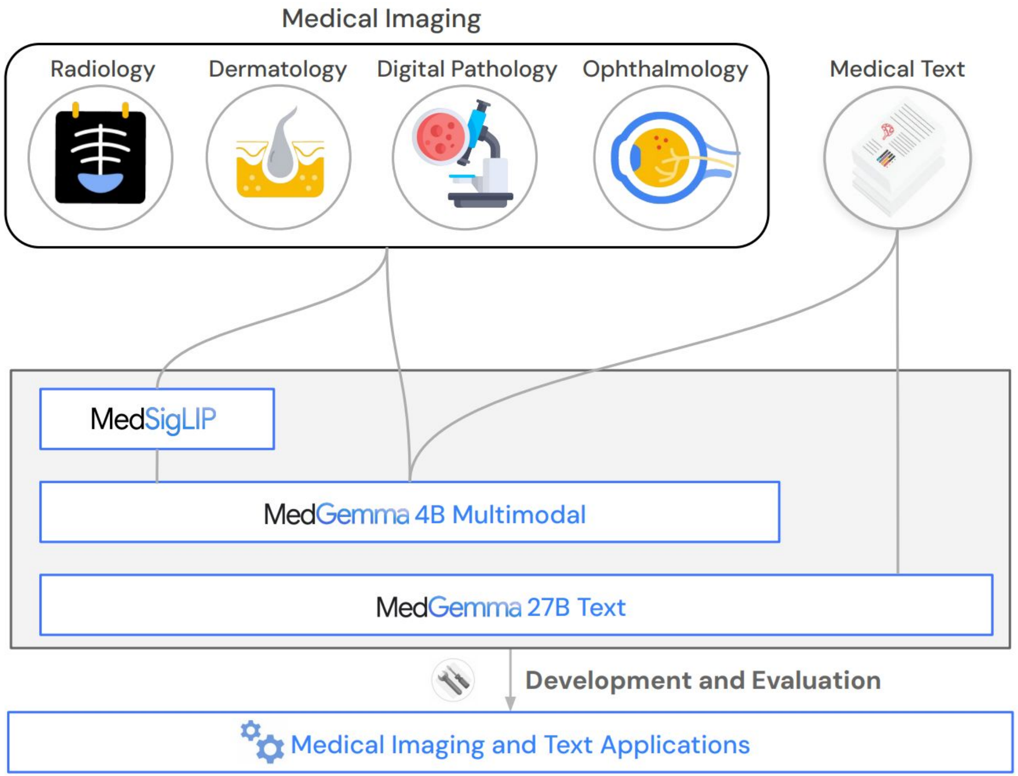

Prompt

You are an expert radiologist. Please succinctly describe the findings for the above chest x-ray.

MedGemma

The chest X-ray shows a large left pneumothorax with a chest tube in place. There is also a possible nodule in the right upper lobe. The mediastinum is midline. The heart size is normal. There is no evidence of pulmonary edema.

Reference Information: Radiologist Impression
Moderately large left hydropneumothorax (predominantly air-containing) with incompletely expanded left lung. If persistent, consider ex vacuo pneumothorax/trapped lung condition. Multiple right lung nodules, highly suspicious for metastases.

Prompt

Describe this image and provide the most likely condition. Keep your answer brief

MedGemma

The chest X-ray shows a significant opacity in the left lower lung field, obscuring the costophrenic angle. This suggests a possible pleural effusion or consolidation. The right lung appears relatively clear. The heart size is within normal limits. The mediastinum is unremarkable.

Reference Information: Radiologist Impression
Lower left lung consolidation. Small-medium left pleural effusion. Similar though lesser findings right side.

Prompt

Describe the findings and most likely diagnosis in this image. Be concise.

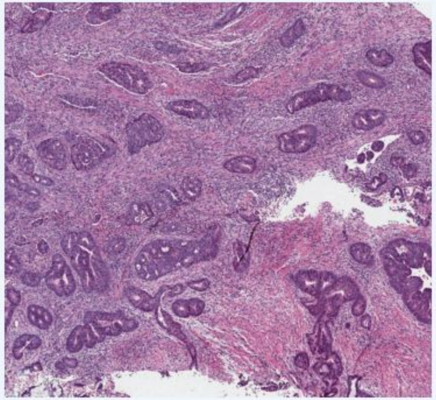

MedGemma

The image shows a histological section of colon tissue. The findings include a mix of glandular structures and stromal tissue. There are areas of inflammation and some degree of cellular atypia.

Most likely diagnosis: **Colorectal adenocarcinoma**

Reference Information:
Colon Adenocarcinoma

Prompt

Describe this image and provide the most likely condition. Keep your answer brief.

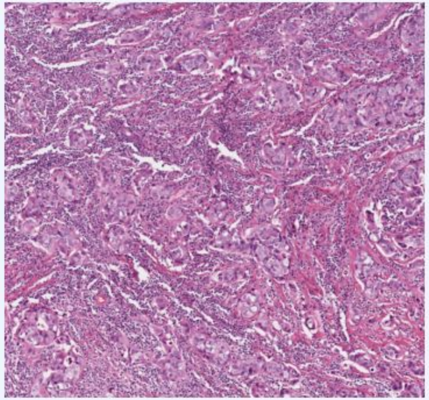

MedGemma

The image shows a high-grade invasive carcinoma with significant cellular atypia, prominent nucleoli, and increased mitotic activity. This indicates a malignant tumor with aggressive growth.

Most likely diagnosis: **High-grade invasive carcinoma** (the specific type would require further information about the tissue origin).

Reference Information:
Invasive Ductal Carcinoma (Breast Cancer)