CS371N Lecture 11
Transformers, Transformer Language Modeling

Announcements

- A3 out

Recap Attention: places a probability distribution over a sequence of n tokens with embeddings $e_1 \dots e_n$

Simplified version:

(1) Form keys $k_i = W^K e_i$

$$K = E(W^K)^T$$

query q

(2) Compute scores $s_i = k_i^T q$ $\quad s = [0\ 0\ 1\ 0]$ $\quad \alpha = [\frac{1}{6}\ \frac{1}{6}\ \frac{1}{2}\ \frac{1}{6}]$

(3) Compute attn weights $\alpha = \text{softmax}(s)$

(4) Result (output) $= \sum \alpha_i e_i$

Self-attention

E: seq len x d

E now gives rise to q_i and k_j for each word

W^K: d x d matrix $\quad K = E(W^K)^T$

W^Q: d x d matrix $\quad Q = E(W^Q)^T$

K, Q: seq len x d

$S = QK^T \qquad S_{ij} = q_i \cdot k_j$

Suppose $E = \begin{bmatrix} 1 & 0 \\ 1 & 0 \\ 0 & 1 \\ 1 & 0 \end{bmatrix}$ $\quad W^Q = W^K = I$

$$S = \begin{bmatrix} 1 & 1 & 0 & 1 \\ 1 & 1 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 1 & 1 & 0 & 1 \end{bmatrix}$$

matrix of similarities

$$A = \text{softmax}(S) = \begin{bmatrix} 3/10 & 3/10 & 1/10 & 3/10 \\ \cdots & \cdots & \cdots & \cdots \\ 1/6 & 1/6 & 1/2 & 1/6 \\ \cdots & \cdots & \cdots & \cdots \end{bmatrix}$$

Last step:

W^V produces "values"

$$\text{Output} = A\left(E\left(W^V\right)^T\right)$$

$$(\text{seq len} \times \text{seq len}) \cdot (\text{seq len} \times d) \cdot (d \times d)$$

$$\text{Out} = \text{seq len} \times d$$

A takes a weighted sum of values according to attention weights at each position

$$\text{First row of output} = \frac{3}{10} \cdot V_1 + \frac{3}{10} \cdot V_2 + \frac{1}{10} \cdot V_3 + \frac{3}{10} \cdot V_4$$

$$\text{Third row} = 1/6 \cdot V_1 + 1/6\, V_2 + \frac{1}{2} V_3 + \frac{1}{6} V_4$$