Plan

Recap

Reminders

Embedding

Motivation

Self-attention

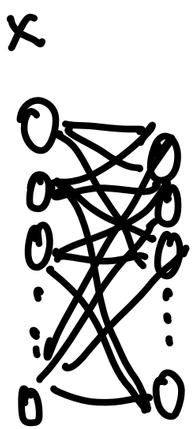

Wx

Recap

1) How to encode text?

2) How to include long range dependency?

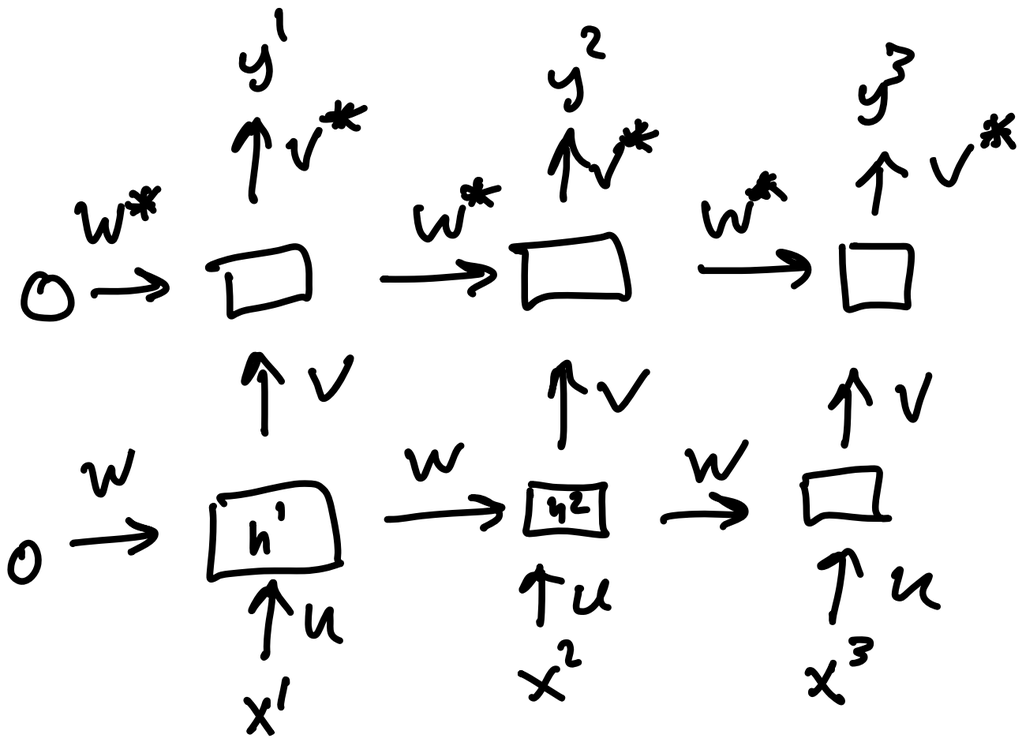

$$h^2 = \sigma(Wh^1 + Ux^2)$$

RNN

- variable length
- memory (but overwritten)

LSTM

cell state

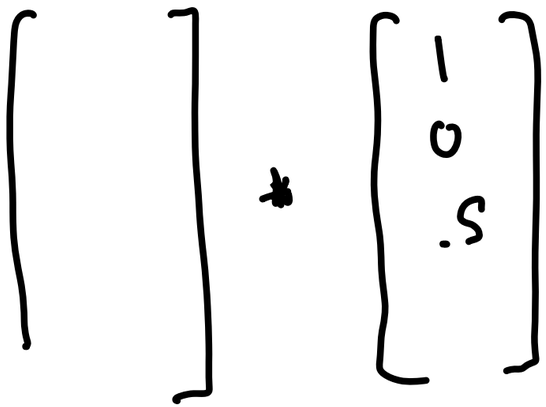

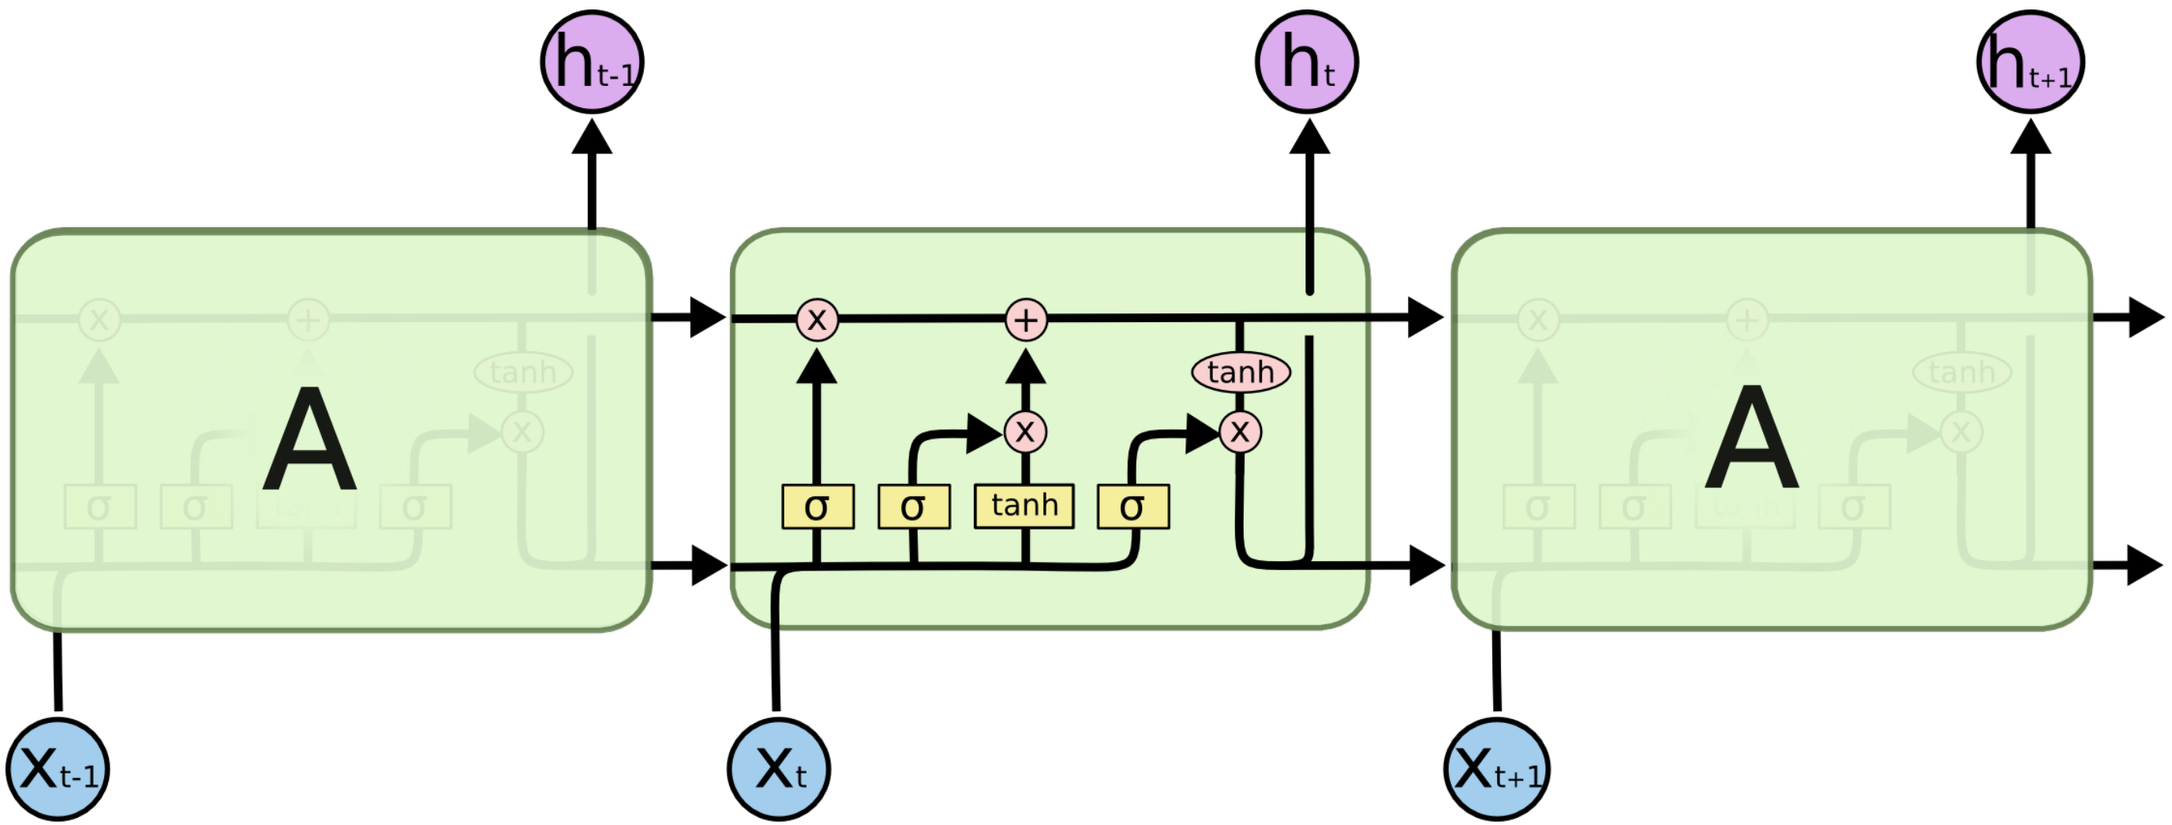

Reminders

- Form 24/26
- 24 hours total lateness
- work in 202 during office hours
- Project proposal due Monday

Embedding

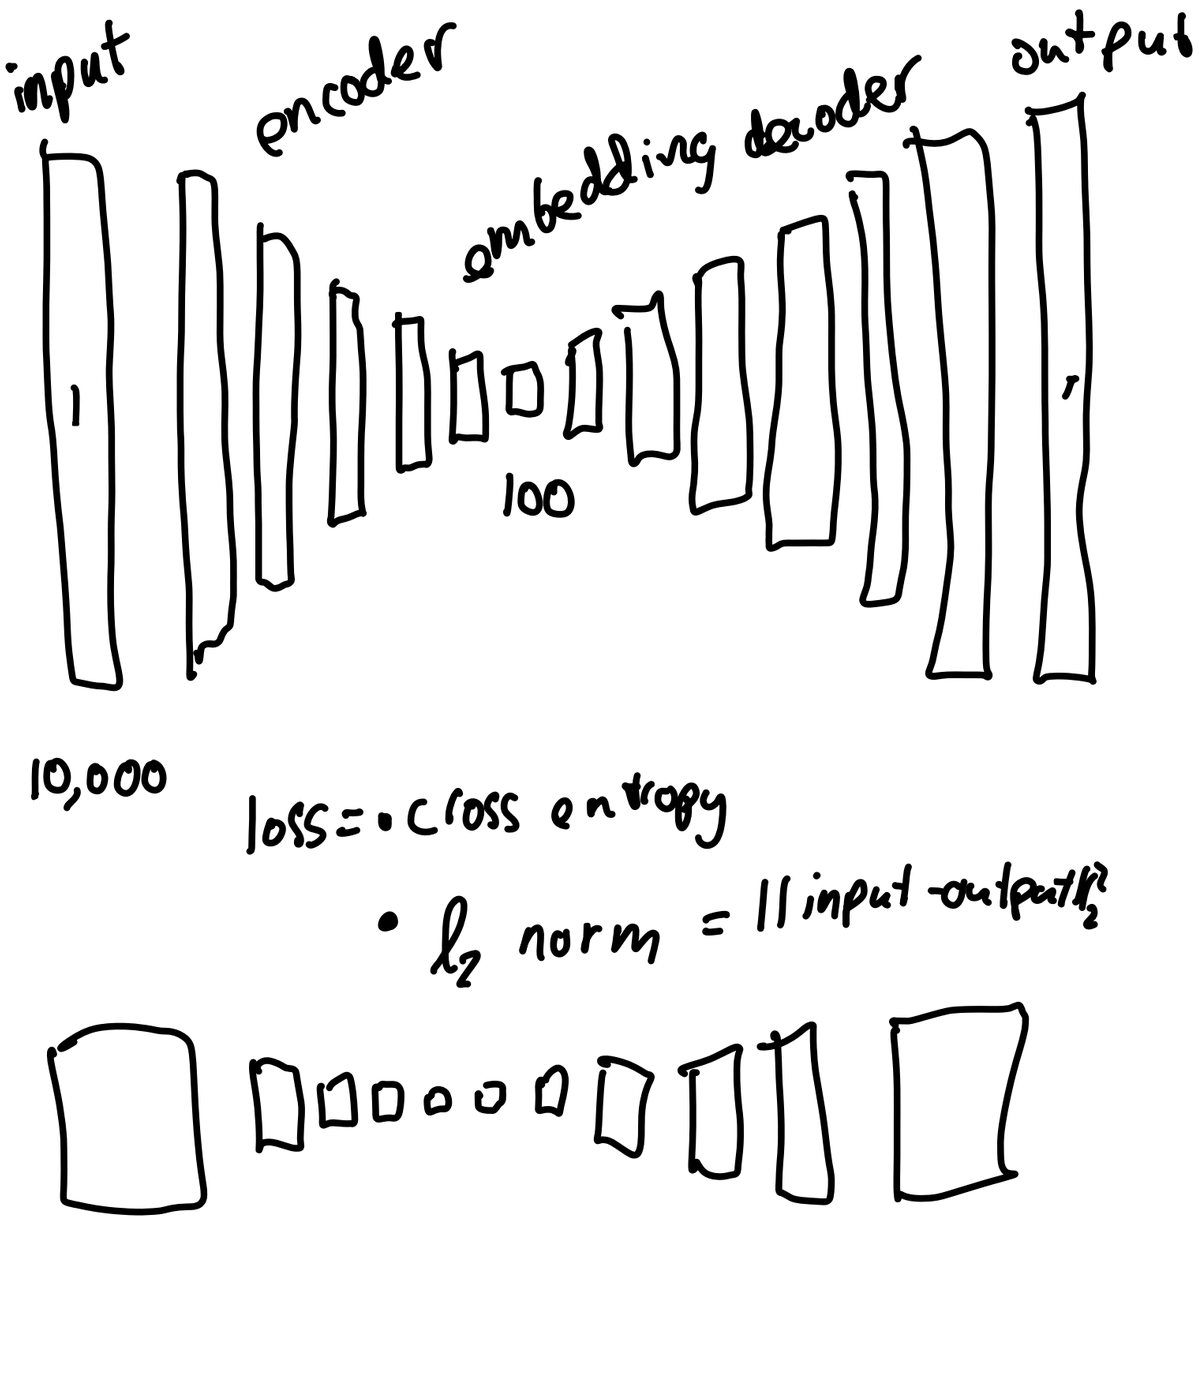

RNN work well on

- next word prediction
- sequence classification

but less well on

- translation
- sentence generation

The dog ran really fast

RNN ☺

El perro corrio muy rapido

I love you a lot

RNN ☹

Te amo mucho

1. different # words 5 vs 3
2. mixed order
3. need whole context

Approach 1: sentences instead of words

sentences of length 5

$5^{\#\text{words}}$

no long range dependency

Wish list:

- long range dependency
- variable length
- association between output and input

Approach 2: bidirectional RNN

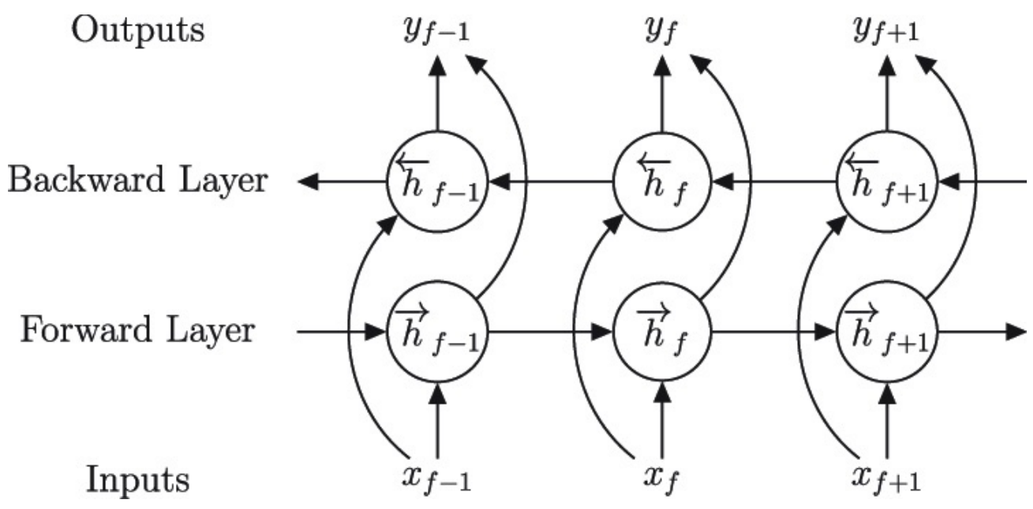

input: $x_1, \dots, x_n \in \mathbb{R}^d$

output: $y_1, \dots, y_n \in \mathbb{R}^d$

$$y_i = \sum_{j=1}^{n} w_{ij} x_j$$

x_j ← value

w_{ij} ↑ row normalized

$$\sum_{j=1}^{n} w_{ij} = 1$$

W

$$\begin{bmatrix} \\ \\ \end{bmatrix}$$

$n \times n$

$$w_{ij} = \frac{\exp(w_{ij})}{\sum_{j'} \exp(w_{ij'})}$$

softmax ↓

w_{ij} = similarity between y_i, x_j

$$= x_i^T x_j$$

x_i ↑ query, x_j ↑ key

I love you a lot

Te amo mucho

amo = $\frac{1}{4}$ I + $\frac{3}{4}$ love + 0 you + 0 a + 0 lot

- set to set ☺
- see all input
  - ↳ no distance
- parameters

x_i: appeared as

- part of output
- weight for own output
- weight for other output

$q_i = W_q x_i$ # query

$k_i = W_k x_i$ # key

$v_i = W_v x_i$ # value

$w_{ij} = q_i^T k_j \quad w_{ij} = \text{softmax}(w_{ij})$

$y_i = \sum_{j=1}^{n} w_{ij} v_j$

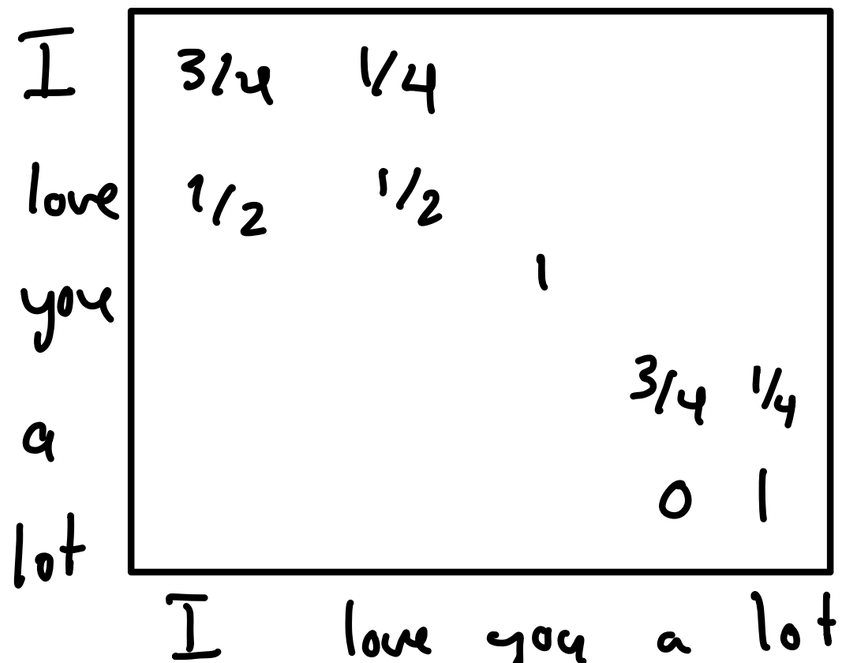

w_{ij} where $i \rightarrow$ love, $j \rightarrow$ you

$= q_i^T k_j$

w_{ji}

$= q_j^T k_i$

Transformer

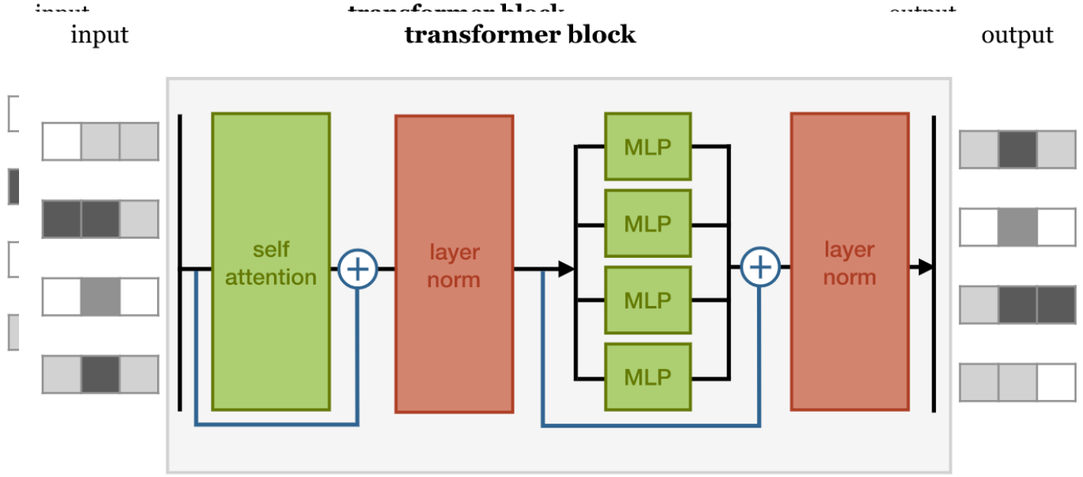

Positional Encoding

①	②	③ ← t	④	⑤
Jack	gave	water	to	Jill
Jill	gave	water	to	Jack

x_i ← embedding "Jack"

encodes position
↳ put in index
↳ one hot encoding

Jack gave water

$$\begin{bmatrix}1\\0\\0\end{bmatrix} \quad \begin{bmatrix}0\\1\\0\end{bmatrix} \quad \begin{bmatrix}0\\0\\1\end{bmatrix}$$

omega

$$\begin{bmatrix}\sin(\omega_1 t)\\ \sin(\omega_2 t)\\ \sin(\omega_3 t)\\ \vdots\end{bmatrix}$$