Applications of sequence models

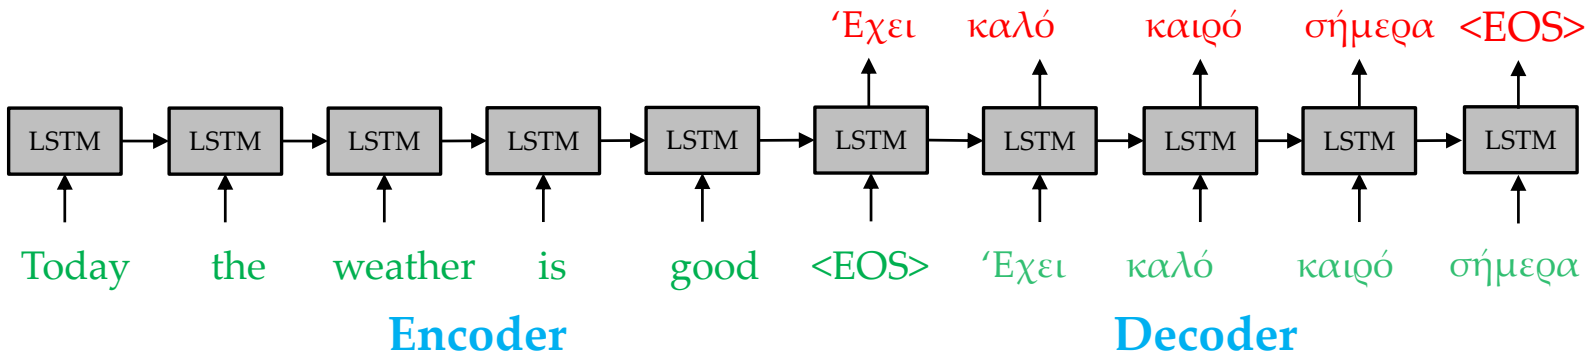

One-hot vectors

- A vector with all zeros except for the active dimension
- Example: 12 words in a sequence → 12 One-hot vectors
- After the one-hot vectors apply an embedding
  - Word2Vec, GloVE

Vocabulary	**One-hot vectors**			
I	I 1	I 0	I 0	I 0
am	am 0	am 1	am 0	am 0
Bond	Bond 0	Bond 0	Bond 1	Bond 0
James	James 0	James 0	James 0	James 1
tired	tired 0	tired 0	tired 0	tired 0
,	, 0	, 0	, 0	, 0
McGuire	McGuire 0	McGuire 0	McGuire 0	McGuire 0
!	! 0	! 0	! 0	! 0

Why not indices instead of one-hot vectors?

- Indices as representations introduce an artificial bias
- Some words suddenly become ‘closer’ because of artificial ordering

OR?

One-hot representation

I	am	James	McGuire
1	0	0	0
0	1	0	0
0	0	0	0
0	0	1	0
0	0	0	0
0	0	0	0
0	0	0	1
0	0	0	0
x_t	x_{t+1}	x_{t+2}	x_{t+3}

$$\ell_2(x_{am}, x_{McQuire}) = \sqrt{2}$$

$$=$$

$$\ell_2(x_I, x_{am}) = \sqrt{2}$$

Index representation

I am James McGuire

$$x_{"I"} = 1$$

$$x_{"am"} = 2$$

$$x_{"James"} = 4$$

$$x_{"McGuire"} = 7$$

$$\ell_2(x_{am}, x_{McQuire}) = (7-2)^2 = 5$$

$$\neq$$

$$\ell_2(x_I, x_{am}) = (2-1)^2 = 1$$

Text generation

- Decompose the text sequence
- Given starting word, sample a new word at a time

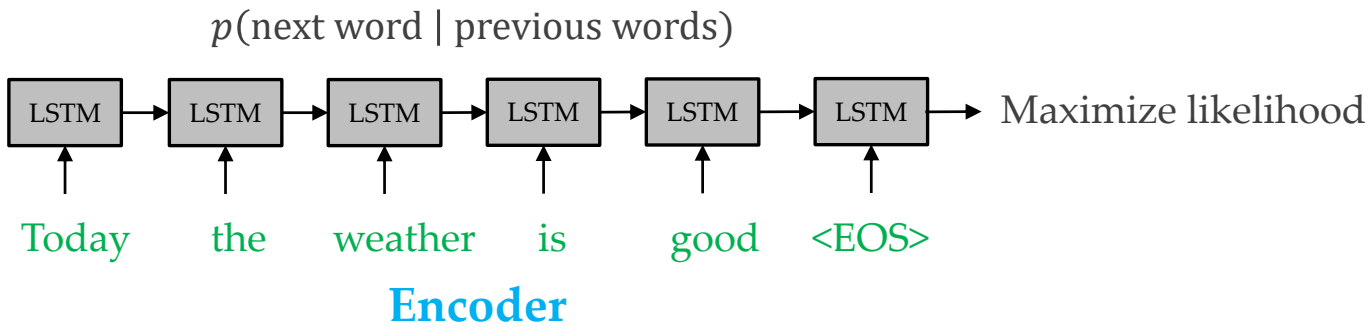

Encoder-Decoder architecture: Machine translation

- The source phrase a sequence encoded by an LSTM
  - "Today the weather is good"
- The phrase in the target language is a sequence decoded by another LSTM
  - "'Εχει καλό καιρό σήμερα"

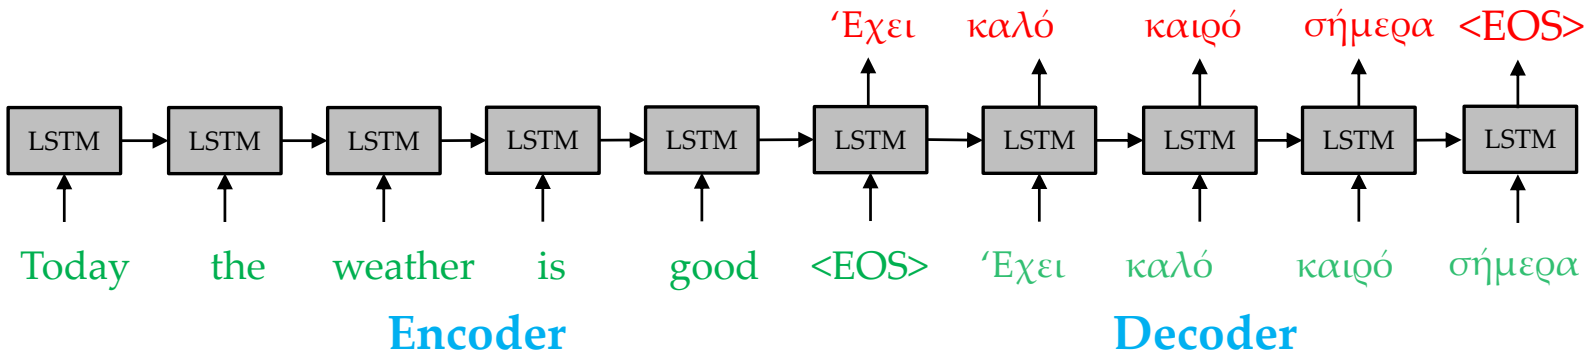

Encoder-Decoder architecture: Image captioning

- Similar to image translation
- The only difference is that the Encoder LSTM is an image ConvNet
  - VGG, ResNet, …
- Keep decoder the same

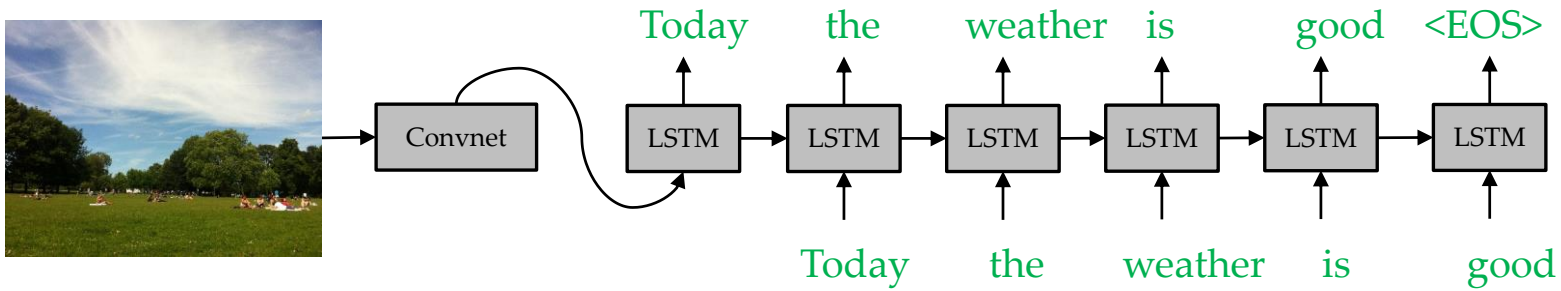

Image captioning demo

Click

NeuralTalk and Walk, recognition, text description of the image while walking

Summary

- Inductive bias: what makes sequences special?
- Backpropagation through time
- LSTMs and variants
- Attention, Transformers
- Applications of sequence models