



# Recurrent Neural Networks


Roberto Basili, Danilo Croce  
Machine Learning, Web Mining & Retrieval 2020/2021

# Syllabus

- (13/5) From Statistical Learning to Deep learning
- (17/5) Convolutional Networks for Image Processing
- (19/5) Laboratory: Programming neural learners
- (20/5) Recurrent Neural Networks
- (24/5) Neural Language Modeling
- (26/5) Laboratory: Designing Convolutional Networks in Python
- (27/5) Encoder-Decoder Networks: Transformers
- (31/5) Designing Recurrent Networks for streaming analytics


# Outline

- Recurrent and Recursive Networks
- Training Recurrent Networks
- LSTMS
- Applications to Language Processing
- Perspectives



# Recurrent Neural Networks


- Used mainly to model sequences
  - naturally applied to textual and speech problems
- A representation at time step  $i$  is made dependent on the representations of the preceding steps
  - connections between units form a directed cycle



# Recurrent Neural Networks

- Common tasks are
  - **language models**: predict the next word in a sentence given the already seen word
  - **speech recognition**: predict a word given the current wave form and the preceding words
  - **machine translation**: produce a sequence in a target language given an input sequence in a source language
- The most famous and effective model of RNNs are the Long-Short Term Memory (LSTM) Networks (Sepp Hochreiter and Jürgen Schmidhuber, 1997)
  - they are meant to better deal with long-range dependencies

# Neural Networks for Natural Language Processing



- Linguistic features have been highly enriched since NN language models have been introduced
  - Words,  $n$ -grams as well as sentences, paragraphs have been modeled through efficient and highly robust neural learners
- Representation are usually dense Embeddings
- Making explicit Use of the contexts: Recurrent Networks
- Beyond Classification: Transducing, Ranking, Encoding, Decoding

# Recurrent Neural Networks

- For example, consider the classical form of a dynamical system

$$\mathbf{s}^{(t)} = f(\mathbf{s}^{(t-1)}; \boldsymbol{\theta}),$$

- Its corresponding unfolded computational graph is as follows


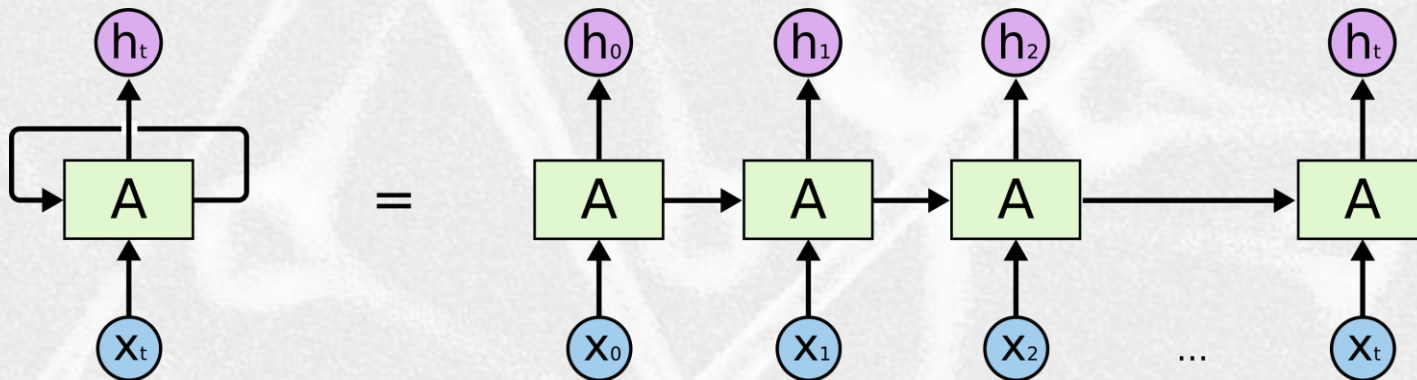


Figure 10.1: The classical dynamical system described by equation 10.1, illustrated as an unfolded computational graph. Each node represents the state at some time  $t$ , and the function  $f$  maps the state at  $t$  to the state at  $t + 1$ . The same parameters (the same value of  $\boldsymbol{\theta}$  used to parametrize  $f$ ) are used for all time steps.






Figure 10.2: A recurrent network with no outputs. This recurrent network just processes information from the input  $\mathbf{x}$  by incorporating it into the state  $\mathbf{h}$  that is passed forward through time. (Left) Circuit diagram. The black square indicates a delay of a single time step. (Right) The same network seen as an unfolded computational graph, where each node is now associated with one particular time instance.

Many recurrent neural networks use equation 10.5 or a similar equation to define the values of their hidden units. To indicate that the state is the hidden units of the network, we now rewrite equation 10.4 using the variable  $\mathbf{h}$  to represent the state,

$$\mathbf{h}^{(t)} = f(\mathbf{h}^{(t-1)}, \mathbf{x}^{(t)}; \boldsymbol{\theta}), \quad (10.5)$$

illustrated in figure 10.2; typical RNNs will add extra architectural features such as output layers that read information out of the state  $\mathbf{h}$  to make predictions.





We can represent the unfolded recurrence after  $t$  steps with a function  $g^{(t)}$ :

$$\mathbf{h}^{(t)} = g^{(t)}(\mathbf{x}^{(t)}, \mathbf{x}^{(t-1)}, \mathbf{x}^{(t-2)}, \dots, \mathbf{x}^{(2)}, \mathbf{x}^{(1)}) \quad (10.6)$$

$$= f(\mathbf{h}^{(t-1)}, \mathbf{x}^{(t)}; \boldsymbol{\theta}). \quad (10.7)$$

The function  $g^{(t)}$  takes the whole past sequence  $(\mathbf{x}^{(t)}, \mathbf{x}^{(t-1)}, \mathbf{x}^{(t-2)}, \dots, \mathbf{x}^{(2)}, \mathbf{x}^{(1)})$  as input and produces the current state, but the unfolded recurrent structure allows us to factorize  $g^{(t)}$  into repeated application of a function  $f$ .

# Using a RNN




Figure 10.3: The computational graph to compute the training loss of a recurrent network that maps an input sequence of  $\mathbf{x}$  values to a corresponding sequence of output  $\mathbf{o}$  values. A loss  $L$  measures how far each  $\mathbf{o}$  is from the corresponding training target  $\mathbf{y}$ . When using softmax outputs, we assume  $\mathbf{o}$  is the unnormalized log probabilities. The loss  $L$  internally computes  $\hat{\mathbf{y}} = \text{softmax}(\mathbf{o})$  and compares this to the target  $\mathbf{y}$ . The RNN has input to hidden connections parametrized by a weight matrix  $U$ , hidden-to-hidden recurrent connections parametrized by a weight matrix  $W$ , and hidden-to-output connections parametrized by a weight matrix  $V$ . Equation 10.8 defines forward propagation in this model. (Left) The RNN and its loss drawn with recurrent connections. (Right) The same seen as a time-unfolded computational graph, where each node is now associated with one particular time instance.

# Simple RNN

## 11.1 Simple RNN

The simplest RNN formulation, known as an Elman Network or Simple-RNN (S-RNN), was proposed by Elman (1990) and explored for use in language modeling by Mikolov (2012). The S-RNN takes the following form:


$$\begin{aligned} \mathbf{s}_i &= R_{\text{SRNN}}(\mathbf{s}_{i-1}, \mathbf{x}_i) = g(\mathbf{x}_i \mathbf{W}^x + \mathbf{s}_{i-1} \mathbf{W}^s + \mathbf{b}) \\ \mathbf{y}_i &= O_{\text{SRNN}}(\mathbf{s}_i) = \mathbf{s}_i \end{aligned} \tag{38}$$

$$\mathbf{s}_i, \mathbf{y}_i \in \mathbb{R}^{d_s}, \quad \mathbf{x}_i \in \mathbb{R}^{d_x}, \quad \mathbf{W}^x \in \mathbb{R}^{d_x \times d_s}, \quad \mathbf{W}^s \in \mathbb{R}^{d_s \times d_s}, \quad \mathbf{b} \in \mathbb{R}^{d_s}$$

That is, the state at position  $i$  is a linear combination of the input at position  $i$  and the previous state, passed through a non-linear activation (commonly tanh or ReLU). The output at position  $i$  is the same as the hidden state in that position.<sup>71</sup>

# Recurrent neural networks (RNNs)

- An RNN can be unwrapped and implemented using the same weights and biases at each step to link units over time as shown below
- The resulting unwrapped RNN is similar to a hidden Markov model, but keep in mind that the hidden units in RNNs are not stochastic



# Recurrent neural networks (RNNs)



- Recurrent neural networks apply linear matrix operations to the current observation and the hidden units from the previous time step, and the resulting linear terms serve as arguments of activation functions  $\text{act}()$ :

$$\mathbf{h}_t = g(\mathbf{W}_h \mathbf{x}_t + \mathbf{U}_h \mathbf{h}_{t-1} + \mathbf{b}_h)$$

$$\mathbf{o}_t = f(\mathbf{W}_o \mathbf{h}_t + \mathbf{b}_o)$$

- The same matrix  $\mathbf{U}_h$  is used at each time step
- The hidden units in the previous step  $\mathbf{h}_{t-1}$  influence the computation of  $\mathbf{h}_t$  where the current observation  $\mathbf{x}_t$  contributes to a  $\mathbf{W}_h \mathbf{x}_t$  term that is combined with  $\mathbf{U}_h \mathbf{h}_{t-1}$  and bias  $\mathbf{b}_h$  terms
- Both  $\mathbf{W}_h$  and  $\mathbf{b}_h$  are typically replicated over time
- The output layer is modeled by a classical neural network activation function applied to a linear transformation of the hidden units, the operation is replicated at each step.

# BPPTT

- For training a recurrent network, a solution is to unfold the recurrent structure and expand it as a feedforward neural network with a certain number of time steps: then apply traditional backpropagation onto this unfolded neural network.
- This solution is known as Backpropagation through Time (BPTT), independently invented by several researchers including (Robinson and Fallside, 1987; Werbos, 1988; Mozer, 1989)

# The loss, exploding and vanishing gradients



- The loss for a particular sequence in the training data can be computed either at each time step or just once, at the end of the sequence.
- In either case, predictions will be made after many processing steps and this brings us to an important problem.
- The gradient for feedforward networks decomposes the gradient of parameters at layer  $l$  into a term that involves the product of matrix multiplications of the form  $\delta^{(l)}\mathbf{W}^{\text{T}(l+1)}$  (remind lessons on backpropagation in feedforward network)
- A recurrent network uses the same matrix at each time step, and over many steps the gradient can very easily either diminish to zero or explode to infinity—just as the magnitude of any number other than one taken to a large power either approaches zero or increases indefinitely

# BPTT: the algorithm

1. Present a sequence of  $k_1$  timesteps of input and output pairs to the network.
2. Unroll the network then calculate and accumulate errors across  $k_2$  timesteps.
3. Roll-up the network and update weights.
4. Repeat

- The TBPTT algorithm requires the consideration of two parameters:
  - **$k_1$** : The *number of forward-pass timesteps between updates*.
    - this influences how slow or fast training will be, given how often weight updates are performed.
  - **$k_2$** : The *number of timesteps to which to apply BPTT*.
    - it should be large enough to capture the temporal structure in the problem for the network to learn.
    - Too large a value results in vanishing gradients






# Dealing with exploding gradients

- The use of L1 or L2 regularization can mitigate the problem of exploding gradients by encouraging weights to be small.
- Another strategy is to simply detect if the norm of the gradient exceeds some threshold, and if so, scale it down.
- This is sometimes called **gradient (norm) clipping** where for a gradient vector  $\mathbf{g}$  and threshold  $T$ ,


$$\text{if } \|\mathbf{g}\| \geq T \text{ then } \mathbf{g} \leftarrow \frac{T}{\|\mathbf{g}\|} \mathbf{g}$$

- where  $T$  is a hyperparameter, which can be set to the average norm over several previous updates where clipping was not used.

# Long-term Dependencies with one single layer




# LSTMS (Hochreiter & Schmidhuber, 1997)





- The Long Short-Term Memory (LSTM) architecture (Hochreiter & Schmidhuber, 1997) was designed to solve the vanishing gradients problem.
- Main idea: to introduce as part of the state representation also specialized **memory cells** (a vector  $C$ ) that can preserve gradients across time.
- Access to the memory cells is controlled by gating components, i.e. smooth mathematical functions that simulate logical gates.

# LSTMS





- At each input state, a gate is used to decide:
  - **how much of the new input** should be written to the memory cell,
  - **how much of the current content** of the memory cell should be **forgotten**.
- Concretely, a gate  $g$  in  $[0;1]^n$  is a vector of values in the range  $[0; 1]$  that is multiplied component-wise with another vector  $C$  in  $R^n$ , and the result is then added to another vector.
- Indices in  $C$  corresponding to near-one values in  $g$  are allowed to pass, while those corresponding to near-zero values are blocked.


# 4 layer RNNs



# ... The memory component and the gates

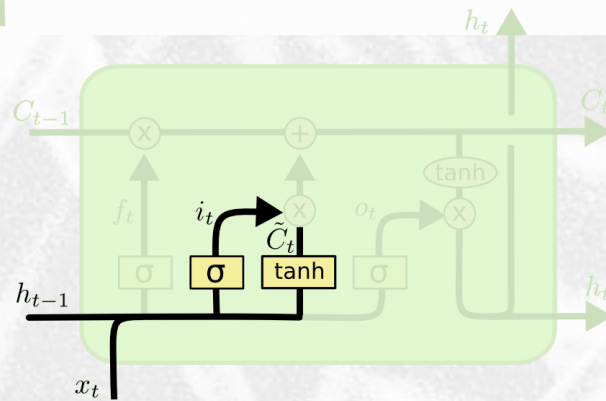




- The FORGET gate



$$f_t = \sigma(W_f \cdot [h_{t-1}, x_t] + b_f)$$

# ... The memory component and the gates





$$i_t = \sigma(W_i \cdot [h_{t-1}, x_t] + b_i)$$

$$\tilde{C}_t = \tanh(W_C \cdot [h_{t-1}, x_t] + b_C)$$

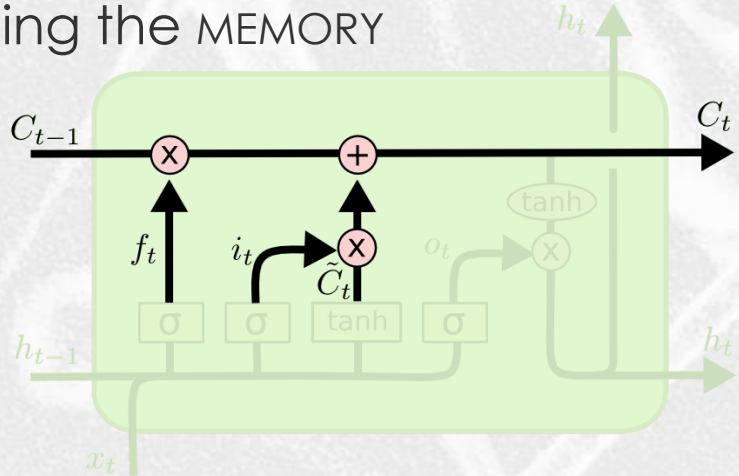
- The INPUT gate



# ... The memory component and the gates





## ■ Updating the MEMORY




$$C_t = f_t * C_{t-1} + i_t * \tilde{C}_t$$



# ... The memory component and the gates



## ■ Computing the OUTPUT



$$o_t = \sigma(W_o [h_{t-1}, x_t] + b_o)$$

$$h_t = o_t * \tanh(C_t)$$

# LSTMS

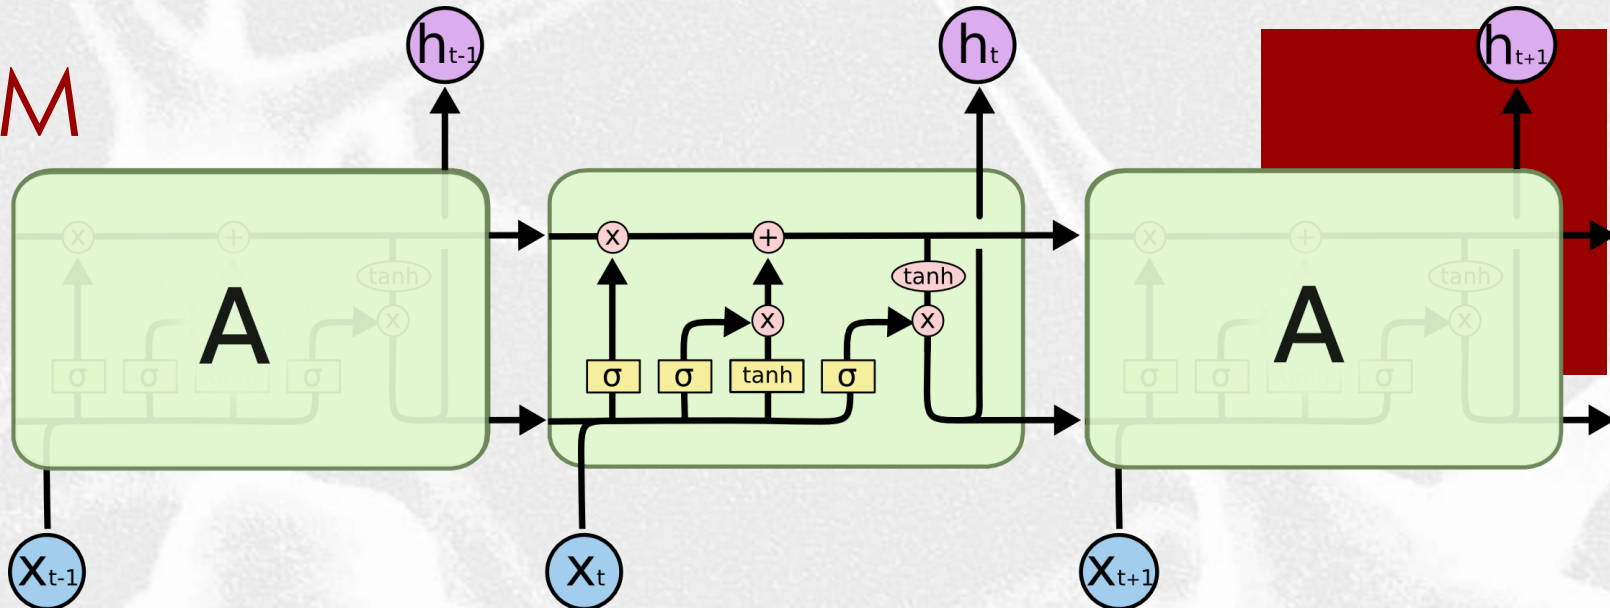
Mathematically, the LSTM architecture is defined as:<sup>72</sup>

$$\begin{aligned} \mathbf{s}_j &= R_{\text{LSTM}}(\mathbf{s}_{j-1}, \mathbf{x}_j) = [\mathbf{c}_j; \mathbf{h}_j] \\ \mathbf{c}_j &= \mathbf{c}_{j-1} \odot \mathbf{f} + \mathbf{g} \odot \mathbf{i} \\ \mathbf{h}_j &= \tanh(\mathbf{c}_j) \odot \mathbf{o} \\ \mathbf{i} &= \sigma(\mathbf{x}_j \mathbf{W}^{\text{xi}} + \mathbf{h}_{j-1} \mathbf{W}^{\text{hi}}) \\ \mathbf{f} &= \sigma(\mathbf{x}_j \mathbf{W}^{\text{xf}} + \mathbf{h}_{j-1} \mathbf{W}^{\text{hf}}) \\ \mathbf{o} &= \sigma(\mathbf{x}_j \mathbf{W}^{\text{xo}} + \mathbf{h}_{j-1} \mathbf{W}^{\text{ho}}) \\ \mathbf{g} &= \tanh(\mathbf{x}_j \mathbf{W}^{\text{xg}} + \mathbf{h}_{j-1} \mathbf{W}^{\text{hg}}) \end{aligned} \tag{39}$$

$$\mathbf{y}_j = O_{\text{LSTM}}(\mathbf{s}_j) = \mathbf{h}_j$$

$$\mathbf{s}_j \in \mathbb{R}^{2 \cdot d_h}, \quad \mathbf{x}_i \in \mathbb{R}^{d_x}, \quad \mathbf{c}_j, \mathbf{h}_j, \mathbf{i}, \mathbf{f}, \mathbf{o}, \mathbf{g} \in \mathbb{R}^{d_h}, \quad \mathbf{W}^{\text{xo}} \in \mathbb{R}^{d_x \times d_h}, \quad \mathbf{W}^{\text{ho}} \in \mathbb{R}^{d_h \times d_h},$$

# LSTM



Mathematically, the LSTM architecture is defined as:<sup>72</sup>

Neural Network Layer

Pointwise Operation

Vector Transfer

Concatenate

Copy


$$\begin{aligned}
 s_j &= R_{\text{LSTM}}(s_{j-1}, x_j) = [c_j; h_j] \\
 c_j &= c_{j-1} \odot f + g \odot i \\
 h_j &= \tanh(c_j) \odot o \\
 i &= \sigma(x_j W^{xi} + h_{j-1} W^{hi}) \\
 f &= \sigma(x_j W^{xf} + h_{j-1} W^{hf}) \\
 o &= \sigma(x_j W^{xo} + h_{j-1} W^{ho}) \\
 g &= \tanh(x_j W^{xg} + h_{j-1} W^{hg})
 \end{aligned}
 \tag{39}$$

$$y_j = O_{\text{LSTM}}(s_j) = h_j$$

$$s_j \in \mathbb{R}^{2 \cdot d_h}, x_j \in \mathbb{R}^{d_x}, c_j, h_j, i, f, o, g \in \mathbb{R}^{d_h}, W^{x_o} \in \mathbb{R}^{d_x \times d_h}, W^{h_o} \in \mathbb{R}^{d_h \times d_h},$$

# Gated Recurrent Unit networks

- Variants



$$z_t = \sigma(W_z \cdot [h_{t-1}, x_t])$$

$$r_t = \sigma(W_r \cdot [h_{t-1}, x_t])$$

$$\tilde{h}_t = \tanh(W \cdot [r_t * h_{t-1}, x_t])$$

$$h_t = (1 - z_t) * h_{t-1} + z_t * \tilde{h}_t$$


# LSTMs and vanishing gradients



- The so-called “long short term memory” (LSTM) RNN architecture was specifically created to address the vanishing gradient problem.
- Uses a combination of hidden units, elementwise products and sums between units to implement gates that control “memory cells”.
- Memory cells are designed to retain information without modification for long periods of time.
- They have their own input and output gates, which are controlled by learnable weights that are a function of the current observation and the hidden units at the previous time step.
- As a result, backpropagated error terms from gradient computations can be stored and propagated backwards without degradation.


# Other RNN architectures

- a) Recurrent networks can be made bidirectional, propagating information in both directions
  - They have been used for a wide variety of applications, including protein secondary structure prediction and handwriting recognition
- b) An “encoder-decoder” network creates a fixed-length vector representation for variable-length inputs, the encoding can be used to generate a variable-length sequence as the output
  - Particularly useful for machine translation



a)

b)



# Training different Types of RNNs




Figure 7: Acceptor RNN Training Graph.

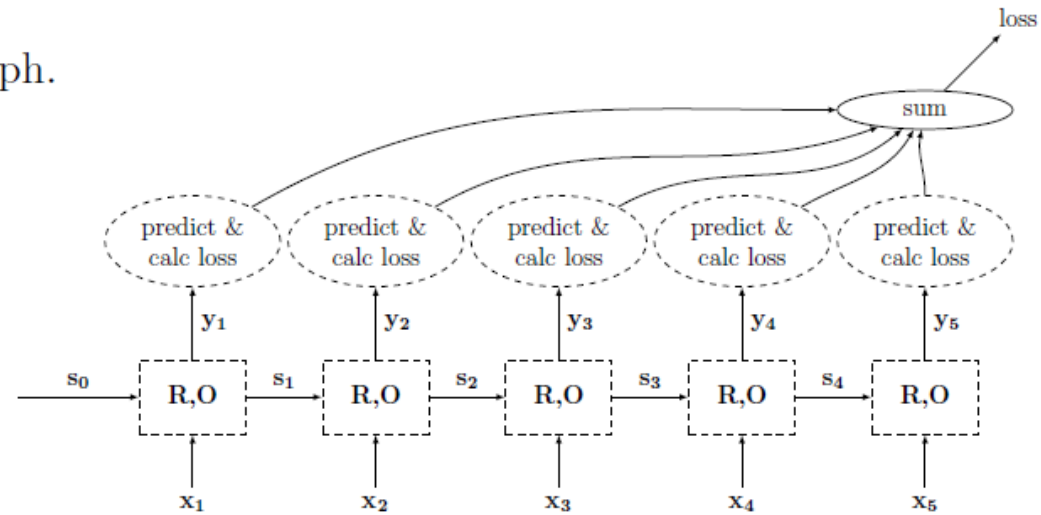


Figure 8: Transducer RNN Training Graph.

# Training different Types of RNNs

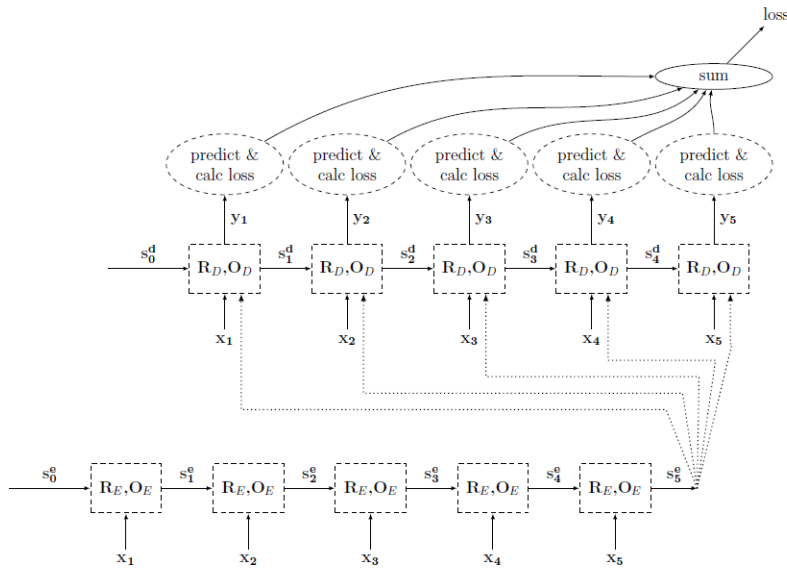


Figure 9: Encoder-Decoder RNN Training Graph.

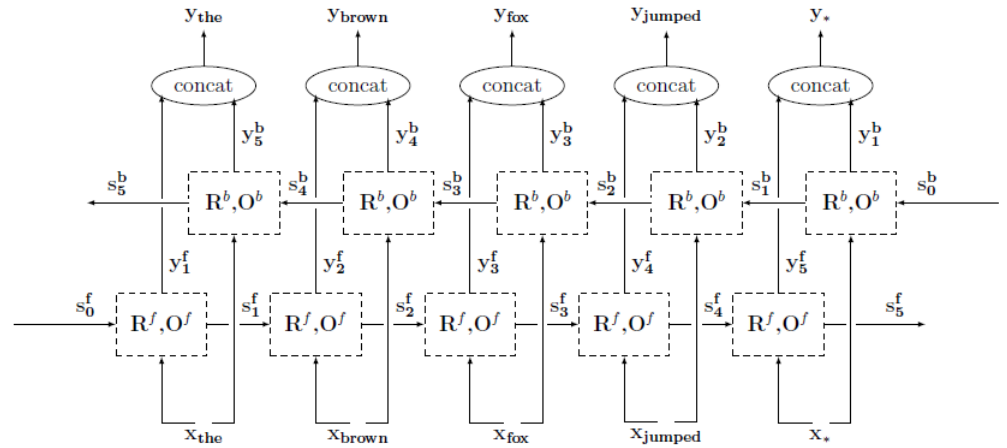
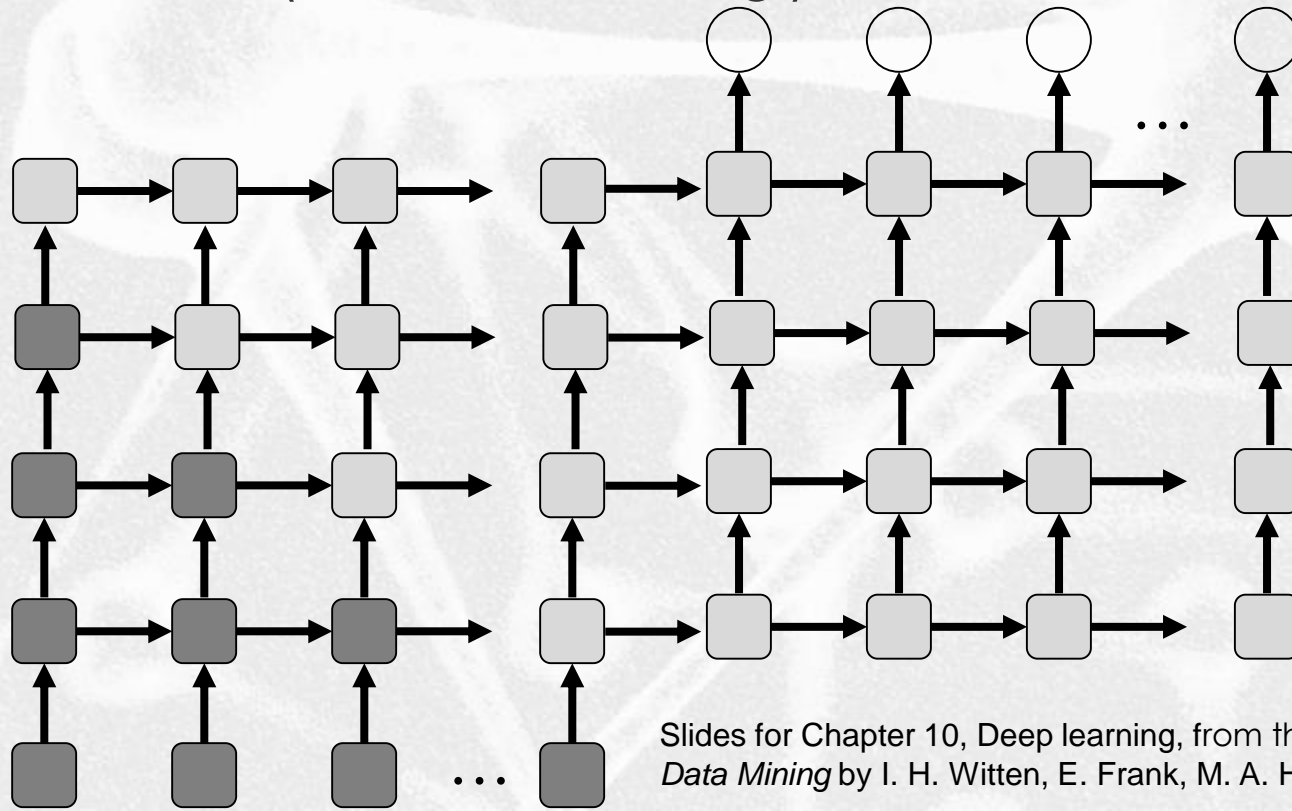


Figure 11: biRNN over the sentence "the brown fox jumped .".



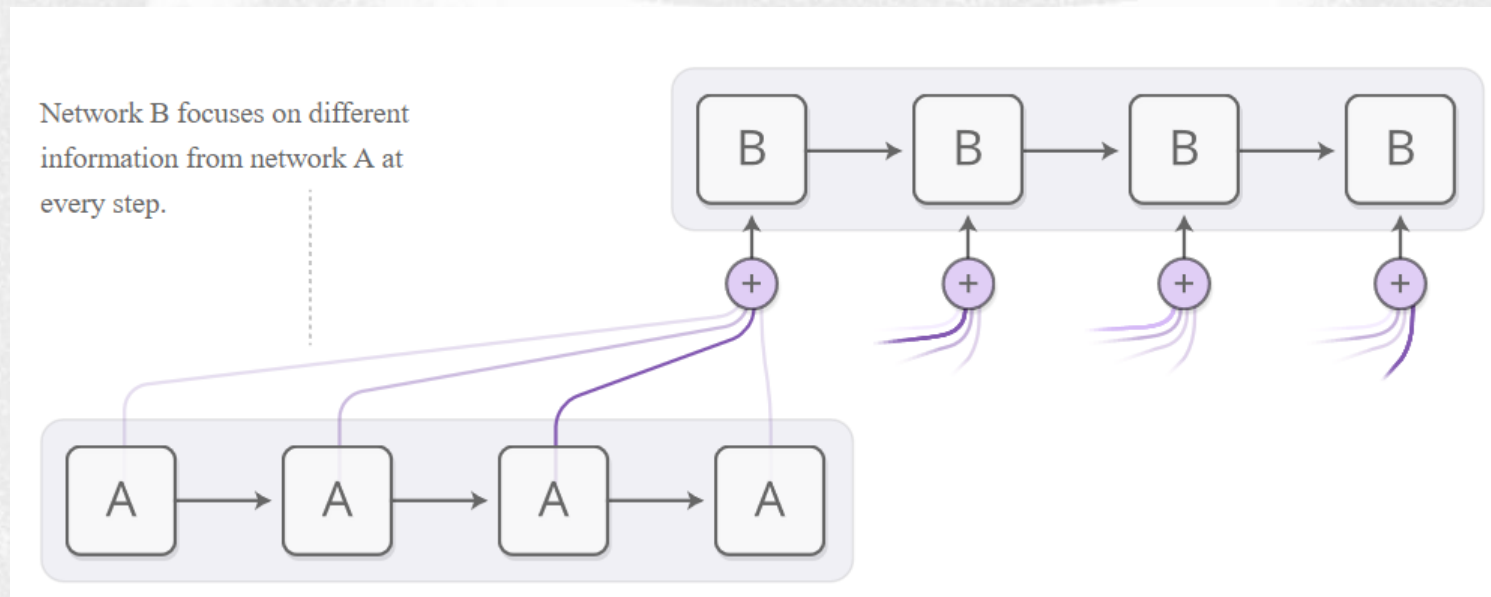
# Encoder-decoder deep architectures

- Given enough data, a deep encoder-decoder architecture (see below) can yield results that compete with hand-engineered translation systems.
- The connectivity structure means that partial computations in the model can flow through the graph in a wave (darker nodes in fig.)



# Attention-based RNNs

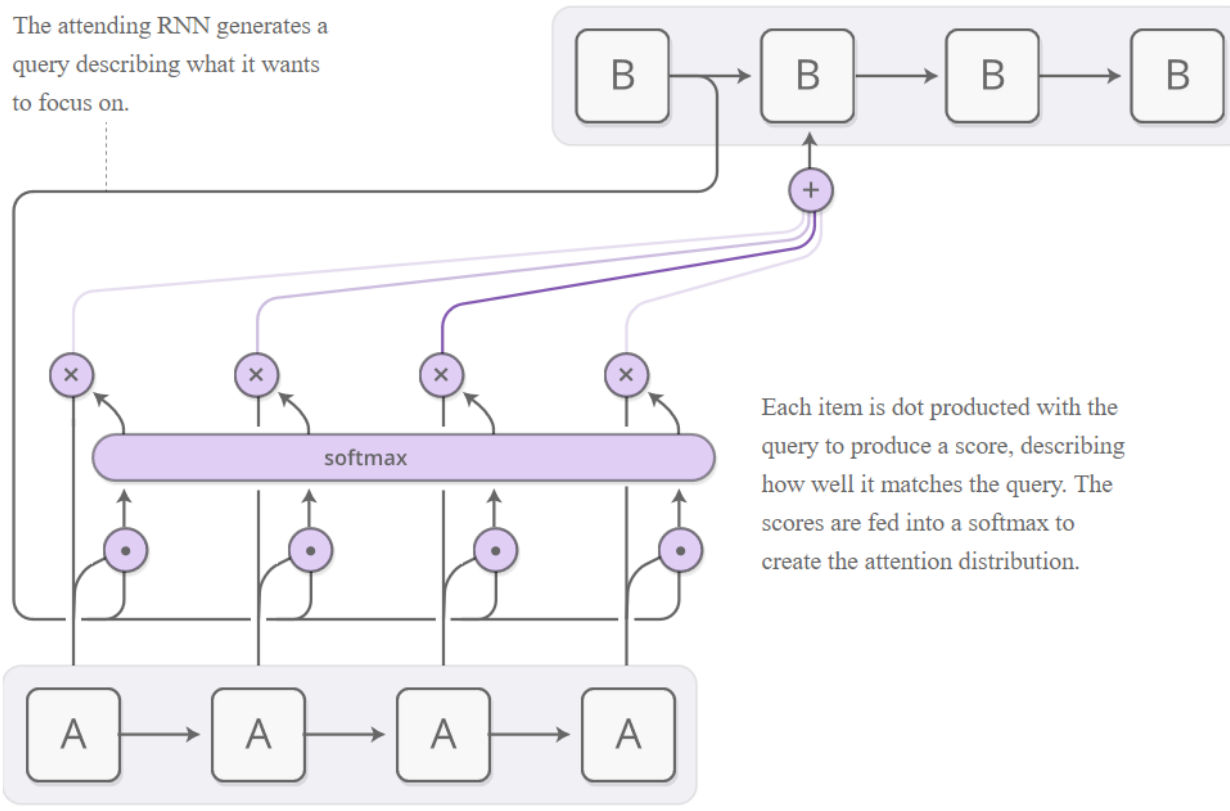
- A NN (e.g. B) is used to attend the outcome of a second network A



# Attention-based RNNs



The attending RNN generates a query describing what it wants to focus on.



Each item is dot producted with the query to produce a score, describing how well it matches the query. The scores are fed into a softmax to create the attention distribution.

# Attention mechanisms in Machine Translation



l' accord sur la zone économique européenne a été signé en août 1992 . <end>

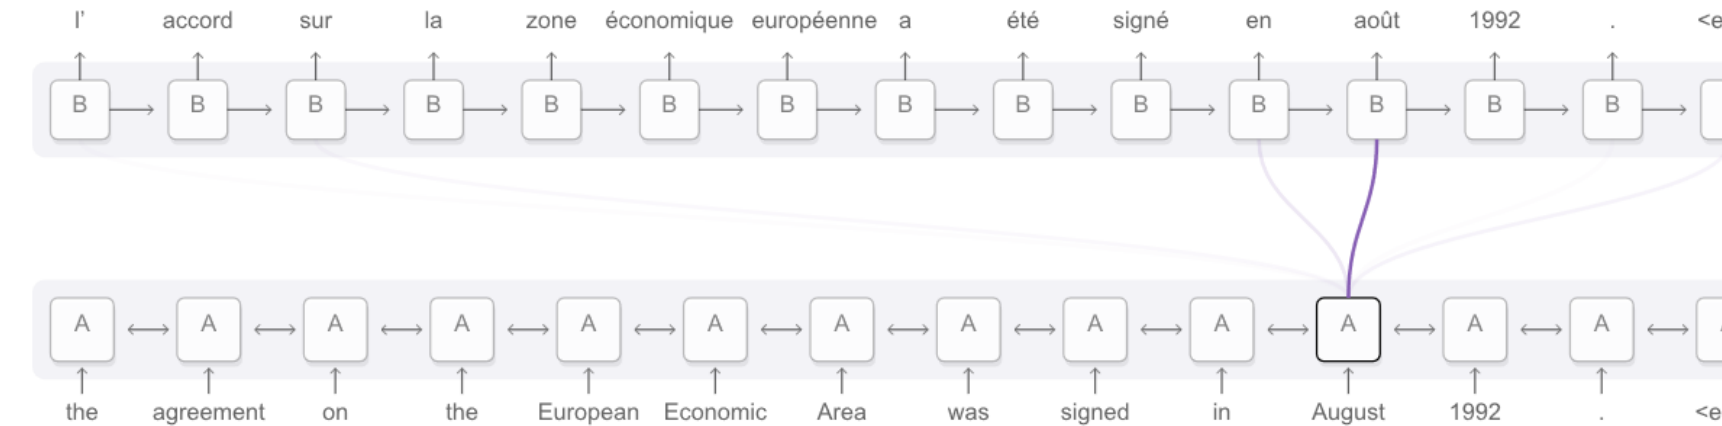
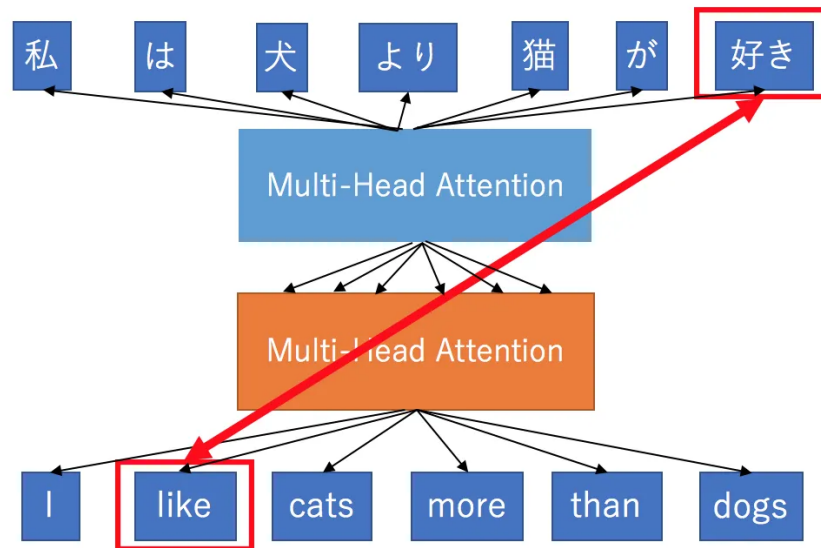
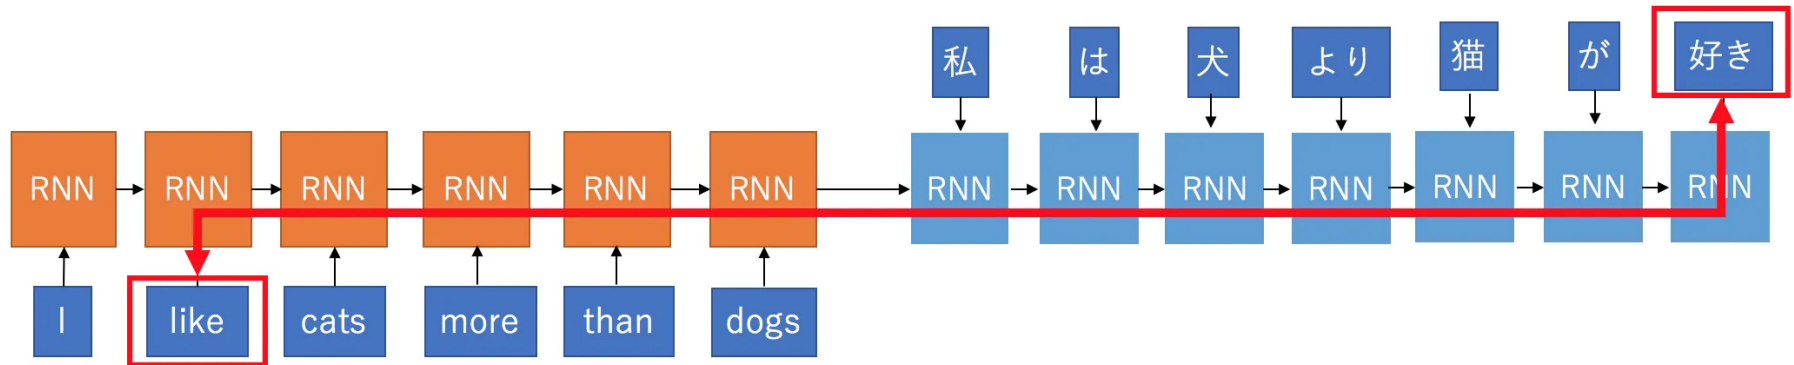



Diagram derived from Fig. 3 of Bahdanau, *et al.* 2014

# Attention & encoding

- IN a decoding process (e.g. machine translation) there are **three** kinds of dependencies for neural architectures
- Dependencies can establish between
- (1) the ***input and output*** tokens
- (2) the ***input tokens themselves***
- (3) the ***output tokens themselves***
- Examples:
  - MT
  - QA where the query the answer paragraph is the input and the matched answer is the output

# Attention in MT: long distance dependencies



# From RNNs to Transformers

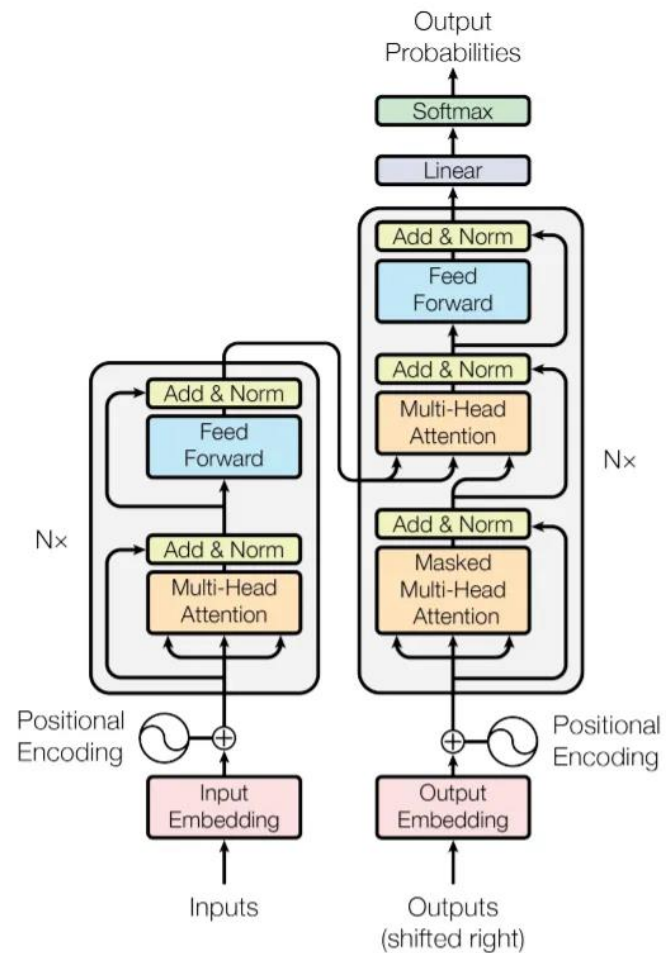


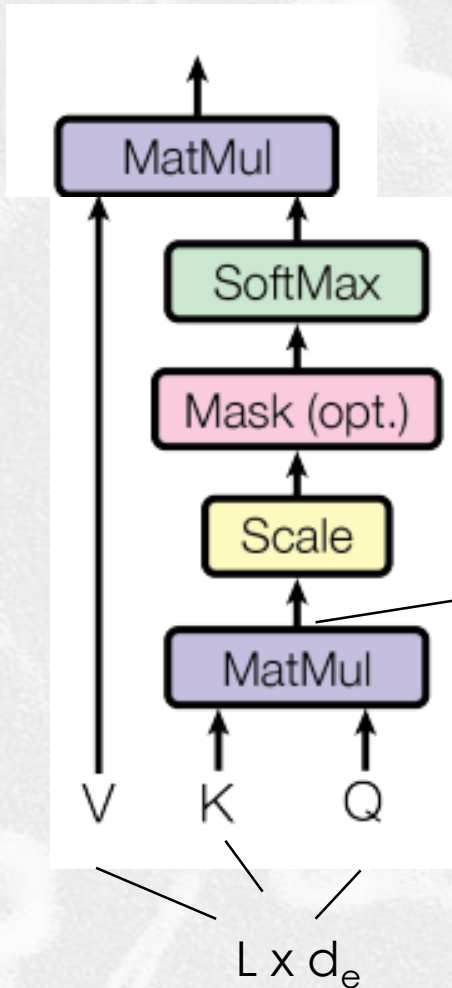


Figure 1: The Transformer - model architecture.

# Bidirectional Encoder Representations from **BERT** - Transformers (Devlin et al. '18)

## Scaled Dot-Product Attention

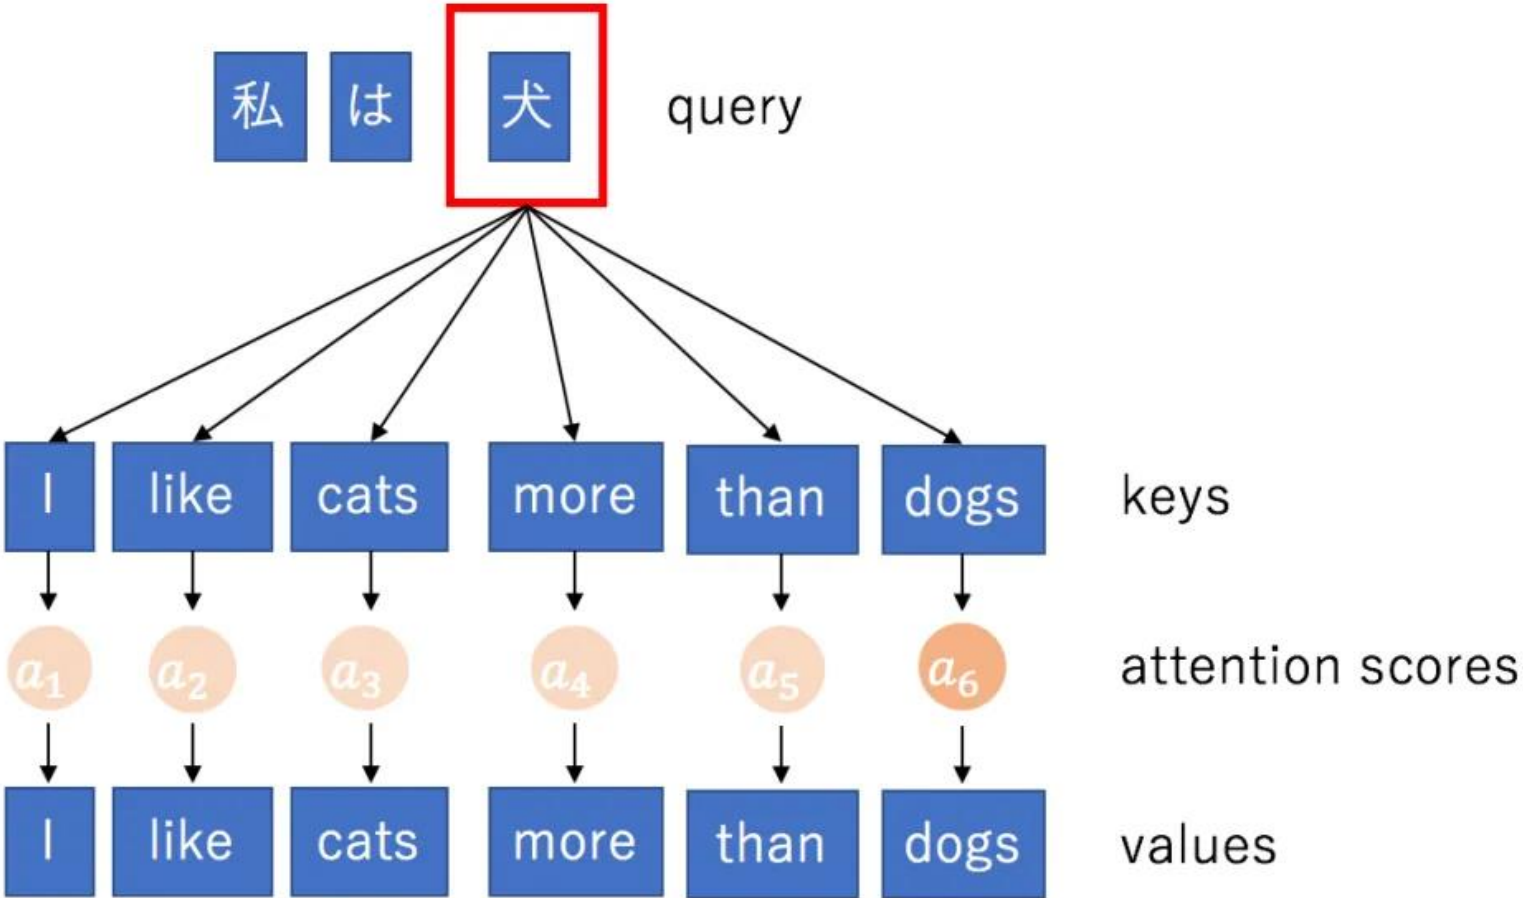
Attention is a function that maps a query  $Q$  and a set of key-value pairs  $\langle K, V \rangle$  to an output





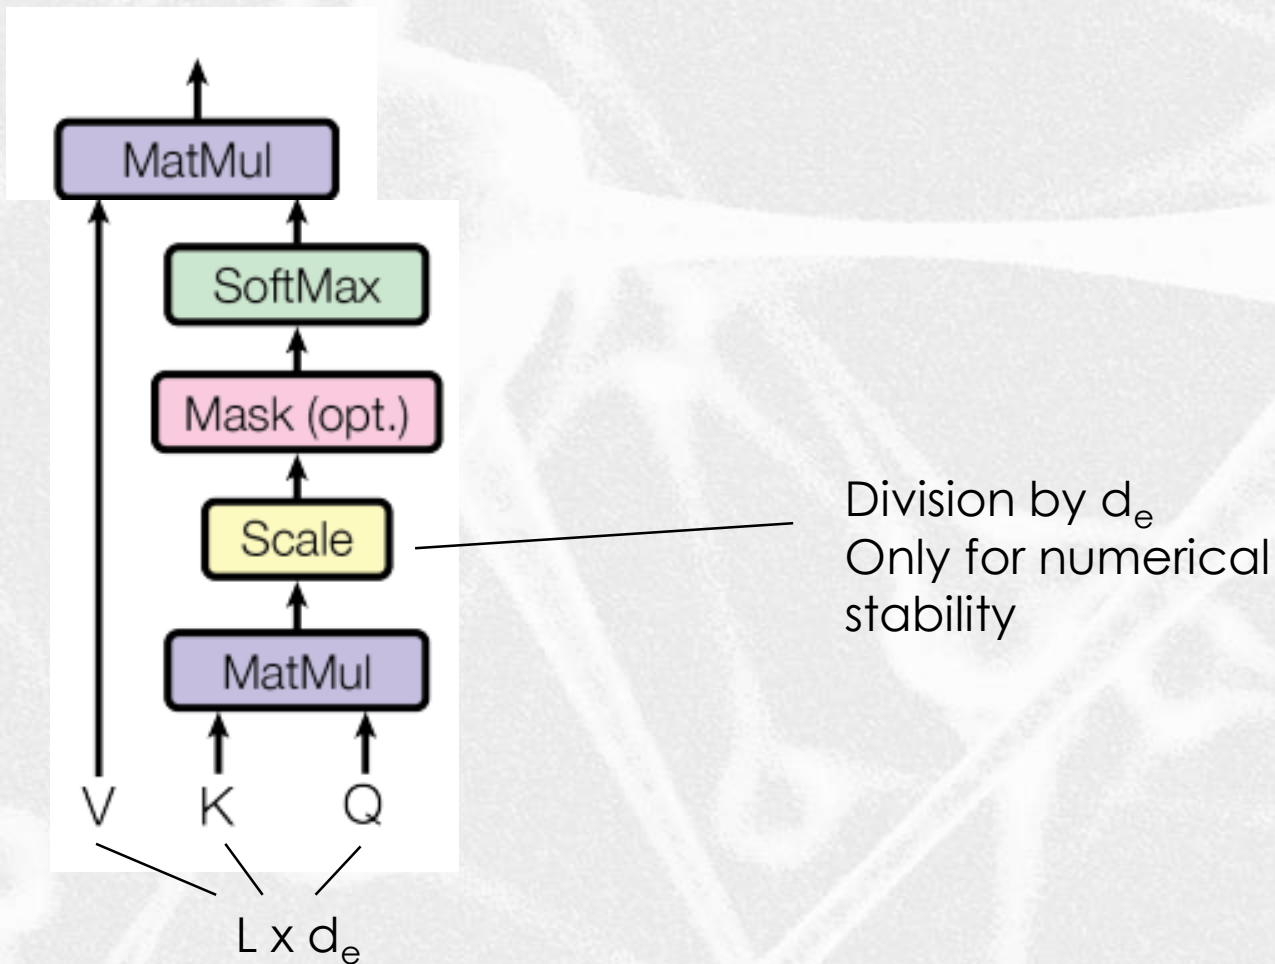


# Input-Output Attention



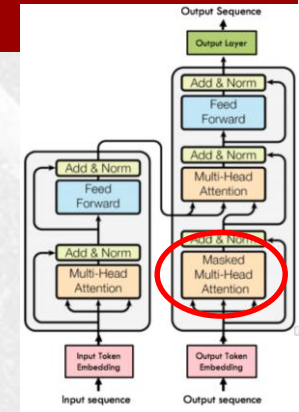


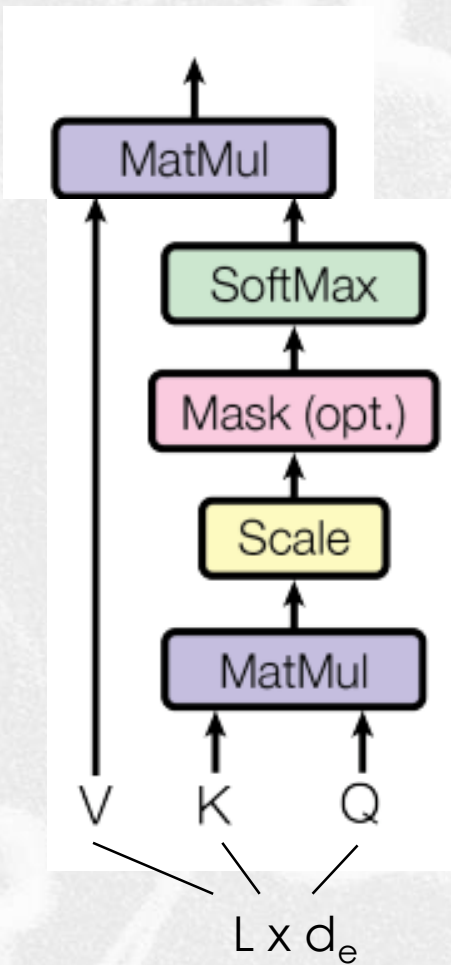
# BERT (Devlin et al. '18)

Scaled Dot-Product Attention





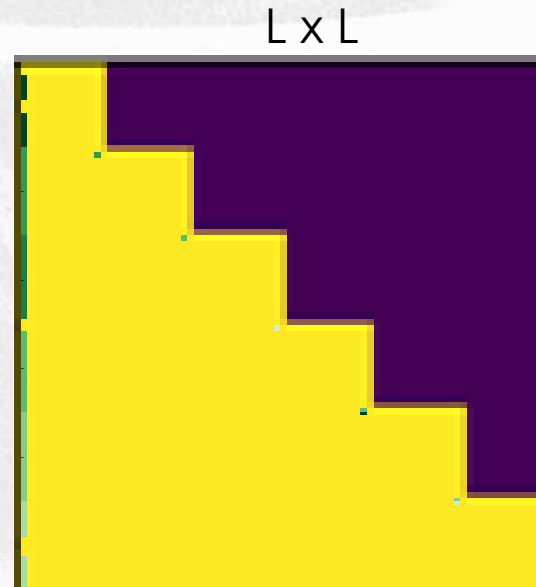
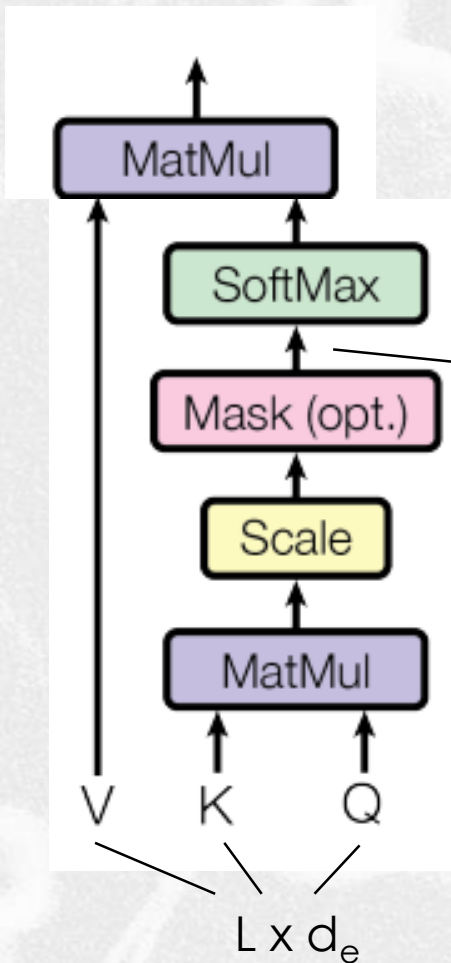
# BERT (Devlin et al. '18)

Scaled Dot-Product Attention





# BERT (Devlin et al. '18)

Scaled Dot-Product Attention



# BERT & NLP


## Encoder



# BERT & NLP (2)



- How to optimize the encoding?
- General and complex tasks defined in (Devlin et al., 2018) are
  - Masked Language Modeling (15%)
    - Inspired by Distributional Hypothesis
    - Can be Simulated and does not require any labeling
  - Next Sentence Prediction
    - Inspired by Textual Inference tasks (e.g. Textual Entailment)
    - Can be Simulated and does not require any labeling
- Source Representations
  - Words? And why not subword (in the BERT jargon: word pieces)?
    - Useful to deal with out-of-vocabulary phenomena

# BERT (Devlin et al. '18)



BERT for single sentence classification (Sentiment analysis, Intent Classification, etc.)


# BERT (Devlin et al. '18)



BERT for Sequence Tagging Tasks (e.g., POS tagging, Named Entity Recognition, etc.)




# BERT (Devlin et al. '18)



BERT for sentence pairs classification (Paraphrase Identification, answer selection in QA, Recognizing Textual Entailment)

# BERT (Devlin et al. '18)



BERT for Answer Span Selection in Question Answering

# A QA example on SquAD

- Cross-lingual Question Answering

COVID-19 ITQA

Insert your question here:

Q How is Covid-19 transmitted? SEARCH

ex. Sintomi covid-19 sui bambini?

Order by Solr  Order by BERT  Order by Solr+BERT

**In-flight Transmission Cluster of COVID-19: A Retrospective Case Series Running title: In-flight Transmission Cluster of COVID-19 figure**

Naibin Yang, Yuefei Shen, Chunwei Shi, Ada Hoi, Yan Ma, Xie Zhang, Xiaomin Jian, Liping Wang, Jiejun Shi, Chunyang Wu, Guoxiang Li, Yuan Fu, Keyin Wang, Mingqin Lu, Guoqing Qian, \* N Yang, Y Shen, C Shi, A Ma

easily transmitted than SARS-CoV [25]. Different from SARS, COVID-19 can be transmitted during the incubation period [26], or by an asymptomatic patient [27]. Features of transmission between SARS and COVID-19 were largely different. For example, health workers account for majority of persons infected with SARS-CoV, while infection with SARS-CoV-2 usually develops in social clusters or family clusters [3]. Wider-Smith reported the first case in-flight transmission of SARS from Singapore [28]. They suggested that it is unlikely to have mass infection of SARS on airplanes. However, we believe it is very likely that mass infection of COVID-19 can occur during a flight, especially when respiratory and contact precautions were not in place. How the SARS-CoV-2 in our study transmitted among the ten passengers was largely unknown. Transmission via aerosol is a possible way for SARS-CoV-2, especially when persons are placed for a long-time under high concentration of aerosol

# BERT (Devlin et al. '18)



**Pretraining** on two unsupervised prediction tasks:

- **Masked Language Model:** given a sentence  $s$  with missing words, reconstruct  $s$ 
  - Example: Amazon <MASK> amazing  $\rightarrow$  Amazon is amazing
  - In BERT the language modeling is deeply Bidirectional, while in ELMo the forward and backward LMs were two independent branches of the NN
- **Next Sentence Prediction:** given two sentences  $s_1$  and  $s_2$ , the task is to understand whether  $s_1$  is the actual sentence that follows  $s_2$ 
  - 50% of the training data are positive examples:  $s_1$  and  $s_2$  are actually consecutive sentences
  - 50% of the training data are negative examples:  $s_1$  and  $s_2$  are randomly chosen from the corpus

# BERT pretraining: Input representations



INPUT

[CLS] my dog is cute [SEP] he MASK play ##ing [SEP]

WordPieces  
Embeddings

$E_{[CLS]}$   $E_{my}$   $E_{dog}$   $E_{is}$   $E_{cute}$   $E_{[SEP]}$   $E_{he}$   $E_{MASK}$   $E_{play}$   $E_{##ing}$   $E_{[SEP]}$

Sentence  
Embeddings

+ + + + + + + + + + +  
 $E_A$   $E_A$   $E_A$   $E_A$   $E_A$   $E_A$   $E_B$   $E_B$   $E_B$   $E_B$   $E_B$

Position  
Embeddings

+ + + + + + + + + + +  
 $E_0$   $E_1$   $E_2$   $E_3$   $E_4$   $E_5$   $E_6$   $E_7$   $E_8$   $E_9$   $E_{10}$

All these embeddings  
are learned during the  
(pre)training process

In pre-training 15% of the input tokens  
are masked for the masked LM task

# Attention mechanisms in Speech Recognition






Figure derived from [Chan, et al. 2015](#)

<https://arxiv.org/pdf/1508.01211.pdf>


# A complex application of LSTM (and recently Transformers): Image captioning



A woman is throwing a frisbee in a park.




A dog is standing on a hardwood floor.




A stop sign is on a road with a mountain in the background.

# Image Captioning

- Image to captions
  - Convolutional Neural Network to learn a representation of the image
  - (Bi-directional) Recurrent Neural Network to generate a caption describing the image
    - its input is the representation computed from the CNN
    - its output is a sequence of words, i.e. the caption



"baseball player is throwing ball in game."






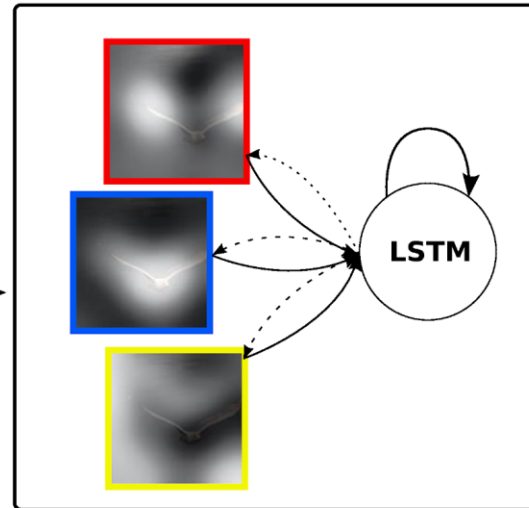


1. Input Image

14x14 Feature Map



2. Convolutional Feature Extraction




3. RNN with attention over the image

**A**  
**bird**  
flying  
over  
a  
body  
of  
**water**

4. Word by word generation


# Attention: a dynamic rendering



A(0.99)




A(1.00)



# Perspectives

- Injecting bias (e.g. linguistic structures) within the learning architectures
- Making use of hybrid architectures integrating visual and linguistic knowledge
- Extending the epistemological transparency of current architectures: Explainable AI
- Making natural language data to work as a representation layer for different cognitive functions (e.g HRI in robotics but also vision)



# RNNs - Bibliographic Notes & Further Readings

- Graves et al. (2009) demonstrate how recurrent neural networks are particularly effective at handwriting recognition,
- Graves et al. (2013) apply recurrent neural networks to speech.
- The form of gradient clipping presented above was proposed by Pascanu et al. (2013).
- Hochreiter and Schmidhuber (1997) is the seminal work on the “Long Short-term Memory” architecture for recurrent neural networks;
  - our explanation follows Graves and Schmidhuber (2005)'s formulation.
- Yoav Goldberg, A Primer on Neural Network Models for Natural Language Processing, Journal of Artificial Intelligence Research volume 57 pp 345-420, 2016
- Greff et al. (2015)'s paper “LSTM: A search space odyssey” explored a wide variety of variants and finds that:
  - none of them significantly outperformed the standard LSTM architecture; and
  - forget gates and the output activation function were the most critical components. Forget gates were added by Gers et al. (2000).

# RNNs - Bibliographic Notes & Further Readings

- IRNNs were proposed by Le et al. (2015)
- Chung et al. (2014) proposed gated recurrent units
- Schuster and Paliwal (1997) proposed bidirectional recurrent neural networks
- Chen and Chaudhari (2004) used bi-directional networks for protein structure prediction; Graves et al. (2009) used them for handwriting recognition
- Cho et al. (2014) used encoder-decoder networks for machine translation, while Sutskever et al. (2014) proposed deep encoder-decoder networks and used them with massive quantities of data
- For further accounts of advances in deep learning and a more extensive history of the field, consult the reviews of LeCun et al. (2015), Bengio (2009), and Schmidhuber (2015)