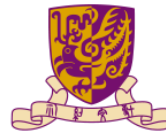# CSC6203/CIE6021:
Large Language Model

香港中文大學（深圳）
The Chinese University of Hong Kong, Shenzhen

Lecture 3: Architecture engineering and scaling law: Transformer and beyond

Winter 2023
Benyou Wang
School of Data Science

To recap…

Language models: Narrow Sense

A probabilistic model that assigns a probability to every finite sequence (grammatical or not)

Sentence: "the cat sat on the mat"

$$P(\text{the cat sat on the mat}) = P(\text{the}) * P(\text{cat}|\text{the}) * P(\text{sat}|\text{the cat})$$
$$* P(\text{on}|\text{the cat sat}) * P(\text{the}|\text{the cat sat on})$$
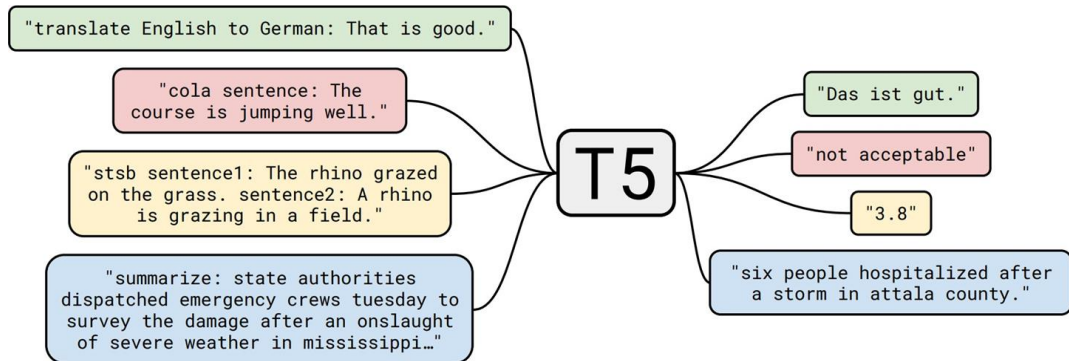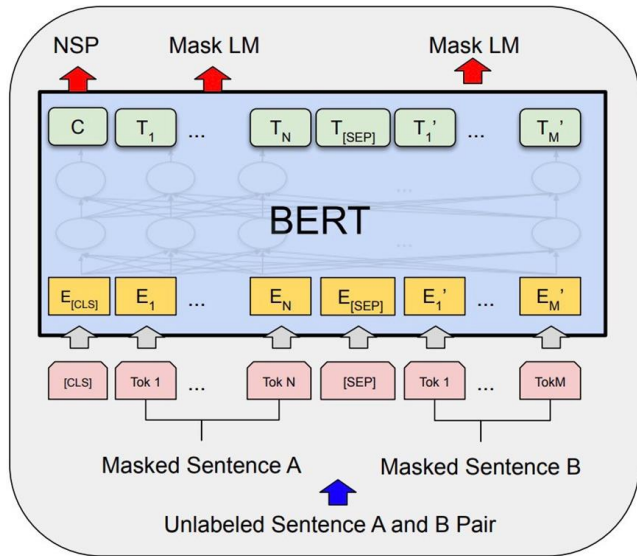$$* P(\text{mat}|\text{the cat sat on the})$$

Implicit order ←

GPT-3 still acts in this way but the model is implemented as a very large neural network of 175-billion parameters!
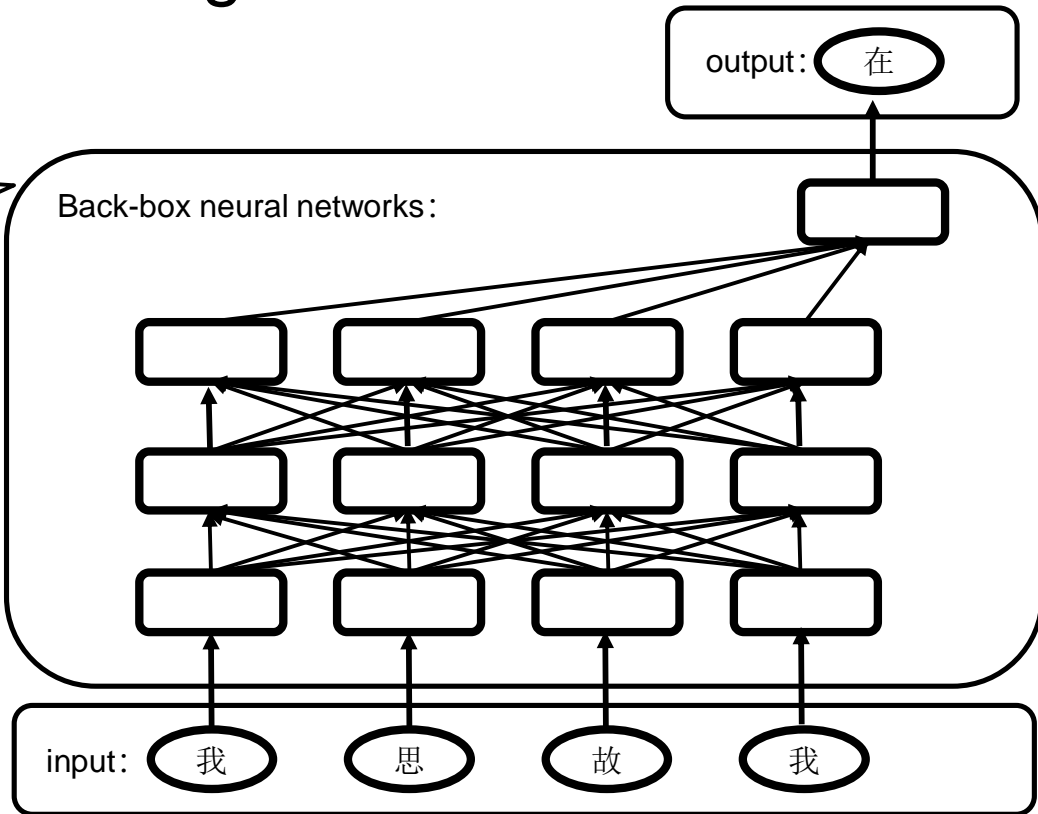
Language models:Broad Sense

❖ Decoder-only models (GPT-x models)
❖ Encoder-only models (BERT, RoBERTa, ELECTRA)
❖ Encoder-decoder models (T5, BART)

The latter two usually involve a different **pre-training** objective.

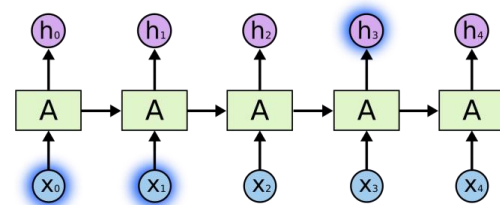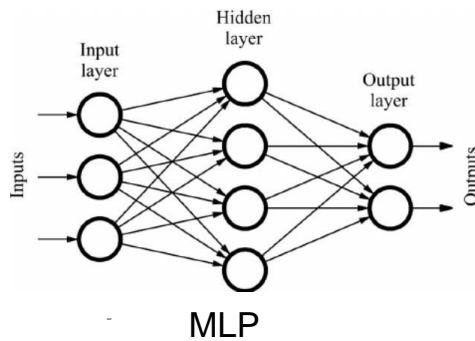# Language model using neural networks

output: 在

GPT-3/ChatGPT/GPT4 have 175B+ parameters
Humans have 100B+ neurons

Back-box neural networks:
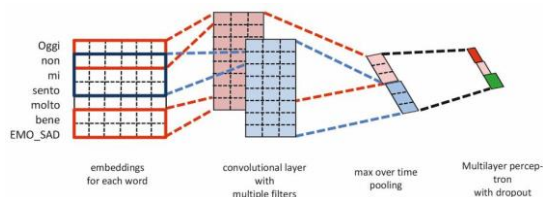
input: 我 思 故 我

Today's Lecture– Big Picture

Which neural networks should be used for LLM?

- ✓ Multilayer Perceptron (MLP)
- ✓ Convolutional neural network
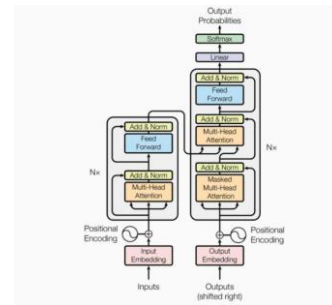- ✓ Recurrent neural network
- ✓ **Transformer**

MLP

Recurrent NNs

Convolutional NNs

Transformer

Which Transformer is so powerful?

Today's lecture

- **MLP**
  - +: Strongest inductive bias: if all words are concatenated
  - +: Weakest inductive bias: if all words are averaged
  - - : The interaction at the token-level is too weak
- CNN & RNN
  - +: The interaction at the token-level is slightly better.
    - CNN: Bringing the global token-level interaction to the window-level
      - - : Make simplifications, its global dependencies are limited
    - RNN: An ideal method for processing token sequences
      - - : Its recursive nature has the problem of disaster forgetting.
- Transformer
  - +: Achieve **global dependence** at the **token-level** by **decoupling** token-level interaction and feature-level abstraction into two components, in **SAN** and **FNN**.

- Scaling law and emergent ability

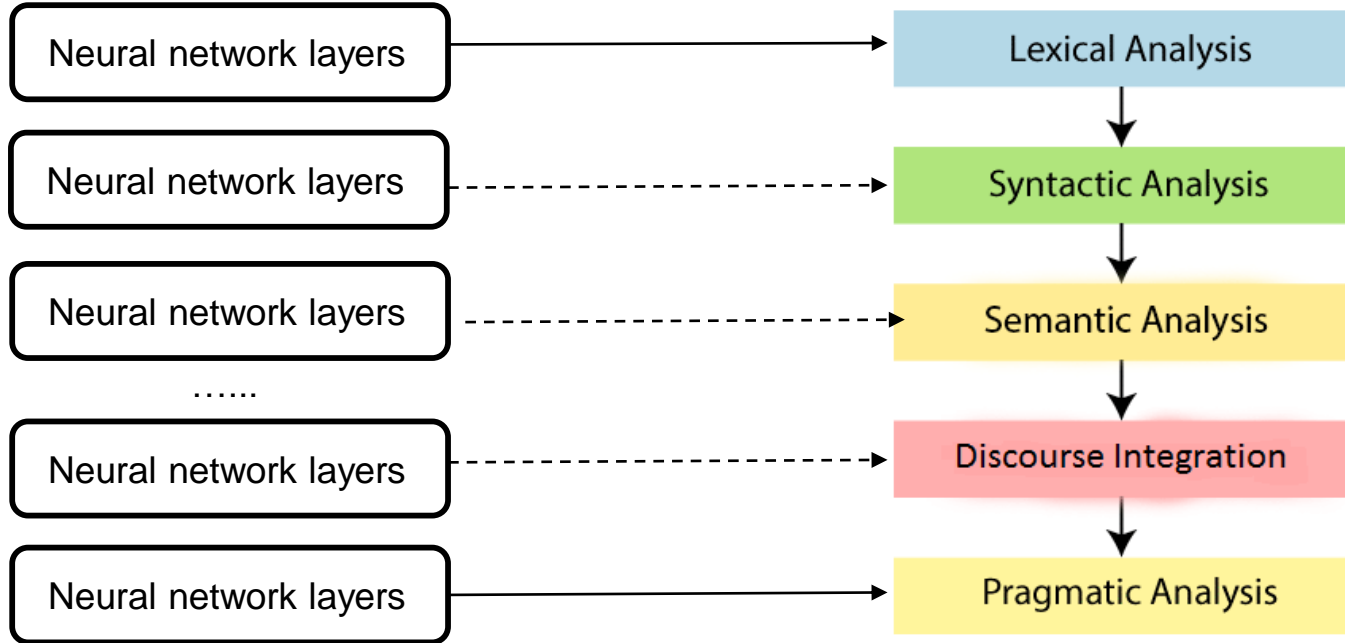# Semantic Abstraction
and Semantic composition

What is Semantic abstraction?
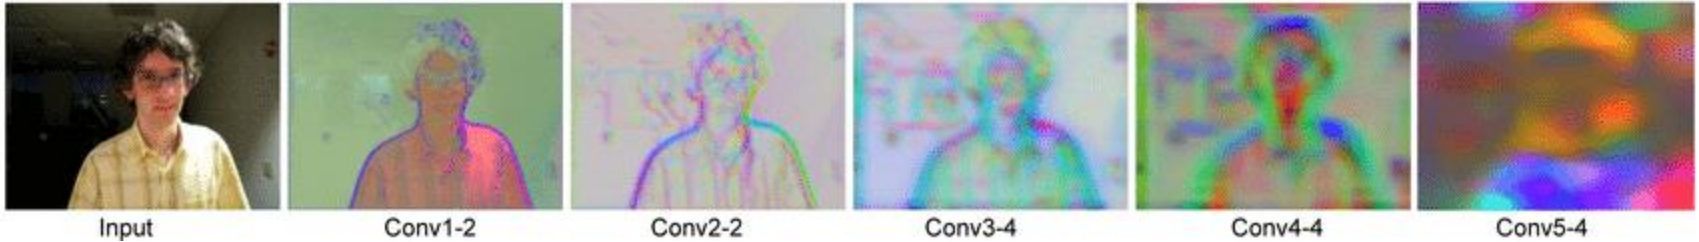
Pixel -> texture -> region -> object -> relation -> semantics->

Higher-level layers deal with higher-degree abstraction

Input: I think therefore I

output: am

What is Semantic composition?

Ivory (象牙)

tower (塔)
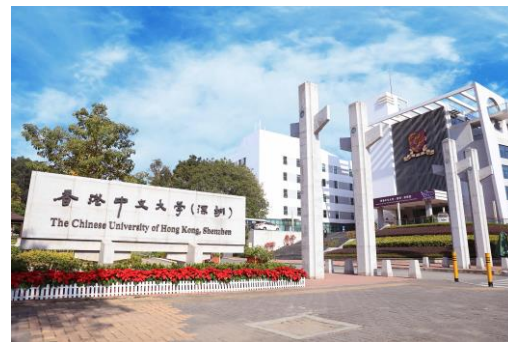
Ivory (象牙塔)

Semantic composition is the task of understanding the meaning of text by composing the meanings of the individual words in the text.

It involves **token interaction**

Semantic composition vs. Semantic Abstraction

Token level: I think therefore I am
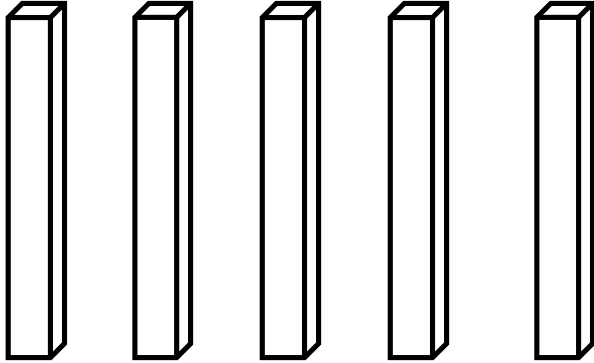
Feature level: word vector
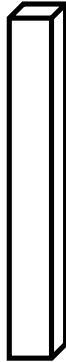
Composition w/ token interaction
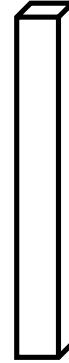
Non-linear Abstraction w/t token interaction

How to combine composition and Abstraction

A flatten solution: MLP (e.g. NNLM)

Complexity: O(D^2L^2)

Yoshua Bengio et.al A Neural Probabilistic Language Model. NIPS 2003
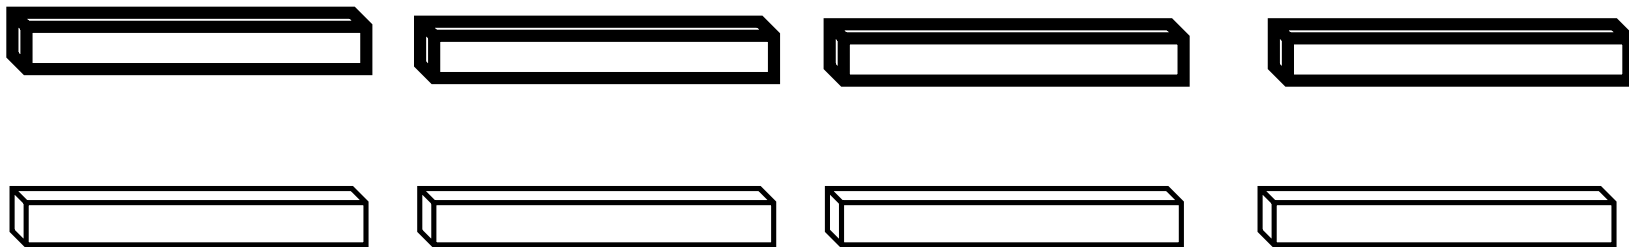
How to combine composition and Abstraction

A variant of MLP (e.g. CBoW**)**
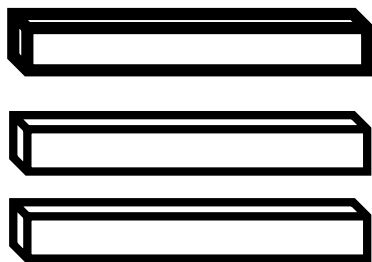
Remove token interaction in deeper layers

Mean pooling (token interaction) in the first layer

Complexity: O(D^2**)**

T Mikolov et.al Efficient Estimation of Word Representations in Vector Space. https://arxiv.org/abs/1301.3781

Inductive bias of composition

How we believe **tokens should be interacted** as the inductive bias, also considering semantic abstraction simultaneously?

Definition: The inductive bias (a.k.a learning bias) of a learning algorithm is the set of assumptions that a machine learning algorithm makes about the relationship between input variables (features) and output variables (labels) based on the training data.

Inductive bias of composition

CNN: **local** composition within a window
RNN: **recurrently** compose tokens from left to right or right to left.
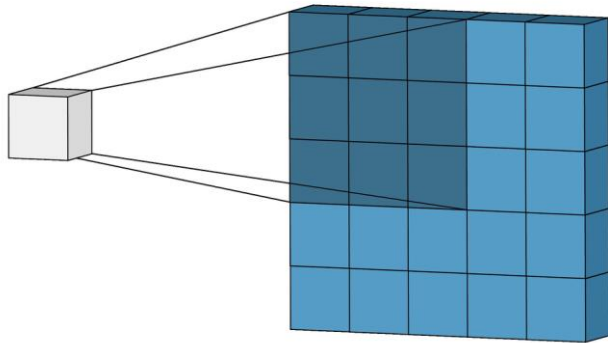
Issues of CNN and RNN

CNN: **local** composition：

     How to make long-term token interaction that is longer than the window size?

RNN: **recurrent** composition

     What if we forget tokens checked 10 timestamp ago?

How can we freely composition tokens without constraints (weaker inductive bias) ?

The modern deep learning is just using weaker inductive biases and make more data-driven instead of prior-driven.

Make each token to see every other token

Attn Scores
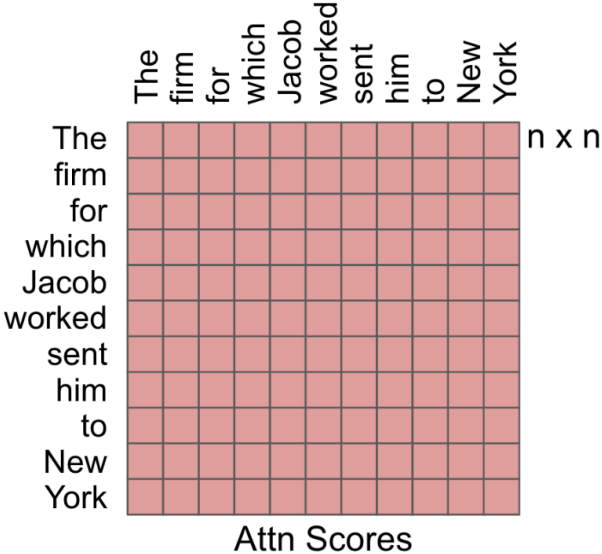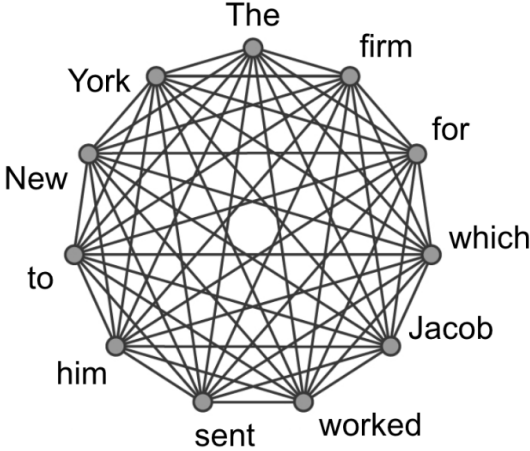
Efficiency: Decompose abstraction and composition

FFN: abstraction

SAN: composition

......

FFN: abstraction

SAN: composition

Feature level: word vector

FFN

SAN

Composition w/ token interaction

Token level: I   think   therefore   I   am

Today's lecture

- **MLP**
  +: Strongest inductive bias: if all words are concatenated
  +: Weakest inductive bias: if all words are averaged
  - : The interaction at the token-level is too weak
- CNN & RNN
  +: The interaction at the token-level is slightly better.
        CNN: Bringing the global token-level interaction to the window-level
              - : Make simplifications, its global dependencies are limited
        RNN: An ideal method for processing token sequences
              - : Its recursive nature has the problem of disaster forgetting.
- Transformer
  +: Achieve **global dependence** at the **token-level** by **decoupling** token-level interaction and feature-level abstraction into two components, in **SAN** and **FNN**.

- Scaling law and emergent ability

Multilayer Perceptron (MLP)

- Introduction of MLP

- Forward Propagation

- Backward Propagation

- Code Implementation

Multilayer Perceptron (MLP)

Definition: The Multilayer Perceptron (MLP) is a type of artificial neural network (ANN) that consists of multiple layers of interconnected artificial neurons or perceptrons.

A **perceptron** can be seen as a single neuron (one output unit with a vector or **layer** of input units):

Output unit: scalar $y = f(\mathbf{x})$

Input layer: vector $\mathbf{x}$
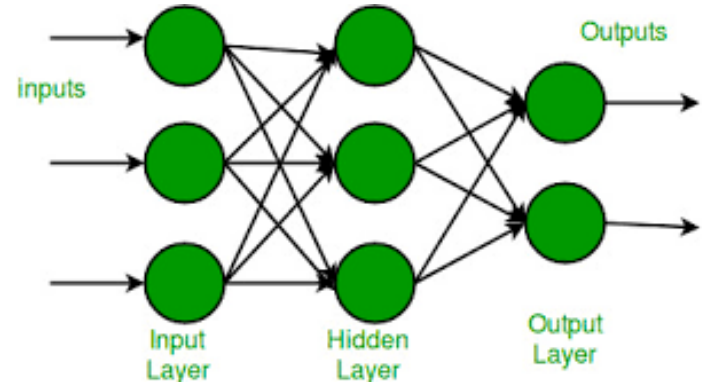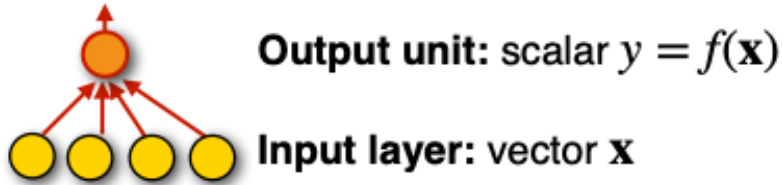
Feed-forward NNs

- The units are connected with no cycles
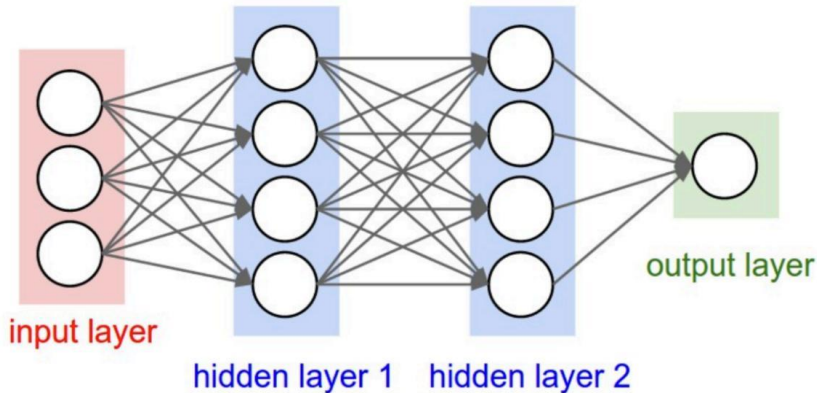- The outputs from units in each layer are passed to units in the next higher layer. No outputs are passed back to lower layers

Fully-connected (FC) layers:

All the units from one layer are fully connected to every unit of the next layer.

input layer

hidden layer 1   hidden layer 2

output layer

```
# forward-pass of a 3-layer neural network:
f = lambda x: 1.0/(1.0 + np.exp(-x)) # activation function (use sigmoid)
x = np.random.randn(3, 1) # random input vector of three numbers (3x1)
h1 = f(np.dot(W1, x) + b1) # calculate first hidden layer activations (4x1)
h2 = f(np.dot(W2, h1) + b2) # calculate second hidden layer activations (4x1)
out = np.dot(W3, h2) + b3 # output neuron (1x1)
```

Feedforward neural language models

A Neural Probabilistic Language Model

(Bengio et al., 2003)

Yoshua Bengio   BENGIOY@IRO.UMONTREAL.CA
Réjean Ducharme   DUCHARME@IRO.UMONTREAL.CA
Pascal Vincent   VINCENTP@IRO.UMONTREAL.CA
Christian Jauvin   JAUVINC@IRO.UMONTREAL.CA

Yoshua Bengio

<u>Probabilistic models of sequences</u>: In the 1990s, Bengio combined neural networks with probabilistic models of sequences, such as hidden Markov models. These ideas were incorporated into a system used by AT&T/NCR for reading handwritten checks, were considered a pinnacle of neural network research in the 1990s, and modern deep learning speech recognition systems are extending these concepts.

<u>High-dimensional word embeddings and attention</u>: In 2000, Bengio authored the landmark paper, "A Neural Probabilistic Language Model," that introduced high-dimension word embeddings as a representation of word meaning. Bengio's insights had a huge and lasting impact on natural language processing tasks including language translation, question answering, and visual question answering. His group also introduced a form of attention mechanism which led to breakthroughs in machine translation and form a key component of sequential processing with deep learning.

<u>Generative adversarial networks</u>: Since 2010, Bengio's papers on generative deep learning, in particular the Generative Adversarial Networks (GANs) developed with Ian Goodfellow, have spawned a revolution in computer vision and computer graphics. In one fascinating application of this work, computers can actually create original images, reminiscent of the creativity that is considered a hallmark of human intelligence.

https://awards.acm.org/about/2018-turing

Feedforward neural language models

A Neural Probabilistic Language Model (Bengio et al., 2003)

Yoshua Bengio   BENGIOY@IRO.UMONTREAL.CA
Réjean Ducharme   DUCHARME@IRO.UMONTREAL.CA
Pascal Vincent   VINCENTP@IRO.UMONTREAL.CA
Christian Jauvin   JAUVINC@IRO.UMONTREAL.CA

Key idea: Instead of estimating raw probabilities, let's use a

neural network to fit the **probabilistic distribution of language**!

$P(w \mid$ I am a good$)$ $P(w \mid$ I am a great$)$

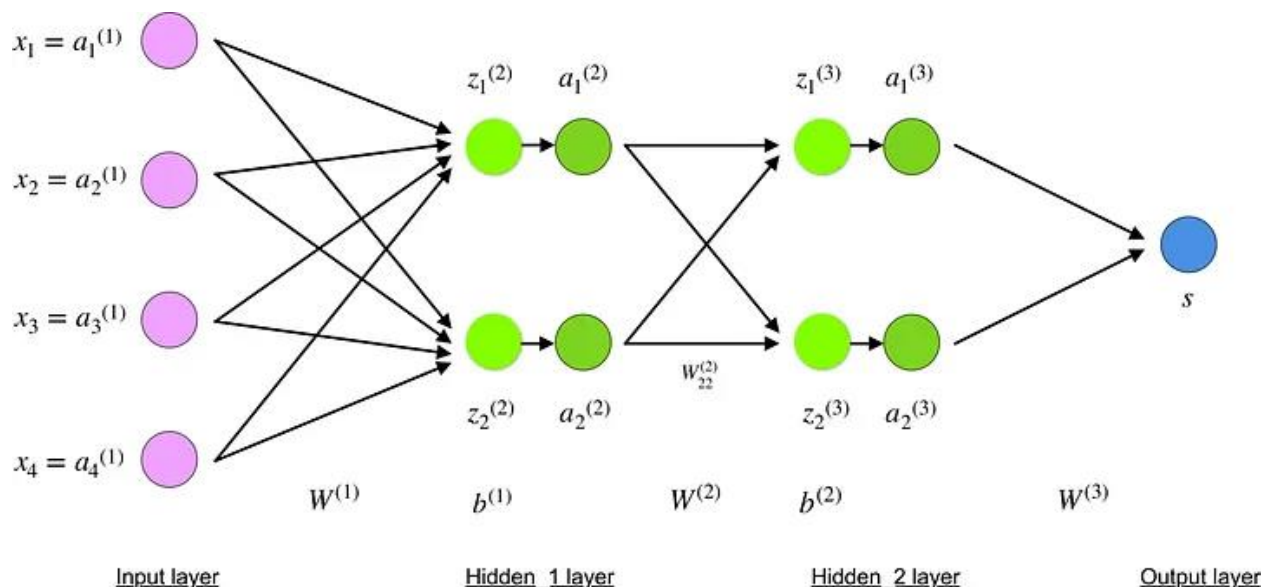**Key ingredient**: word embeddings   **e(good) ≈ e(great)**

Hope: this would give us similar distributions for similar contexts!

Backpropagation

Definition:

Backpropagation, short for "backward propagation of errors," is a supervised learning algorithm used for training artificial neural networks, including deep learning models like Multilayer Perceptrons (MLPs).
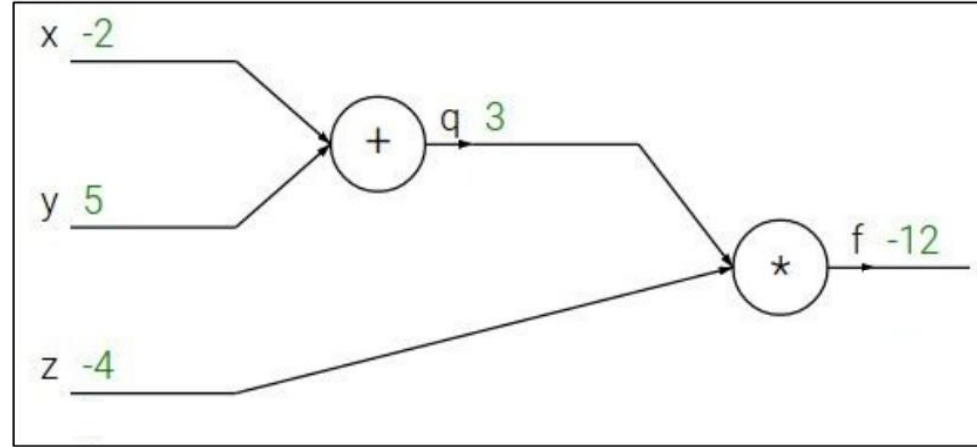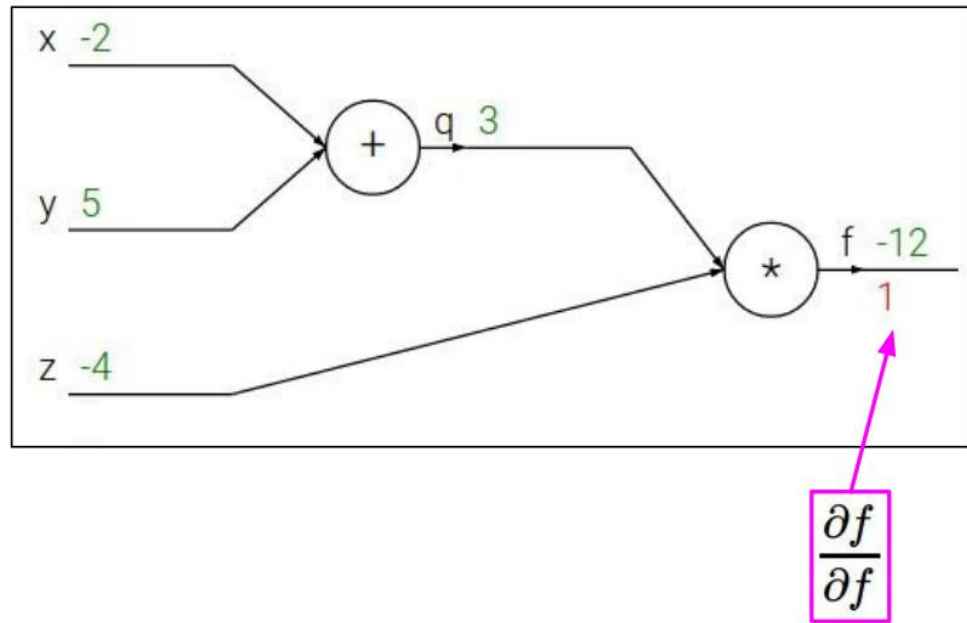
Backpropagation: a simple example

$$f(x, y, z) = (x + y)z$$

e.g. x = -2, y = 5, z = -4

$$q = x + y \qquad \frac{\partial q}{\partial x} = 1, \frac{\partial q}{\partial y} = 1$$

$$f = qz \qquad \frac{\partial f}{\partial q} = z, \frac{\partial f}{\partial z} = q$$

Want: $\dfrac{\partial f}{\partial x}, \dfrac{\partial f}{\partial y}, \dfrac{\partial f}{\partial z}$
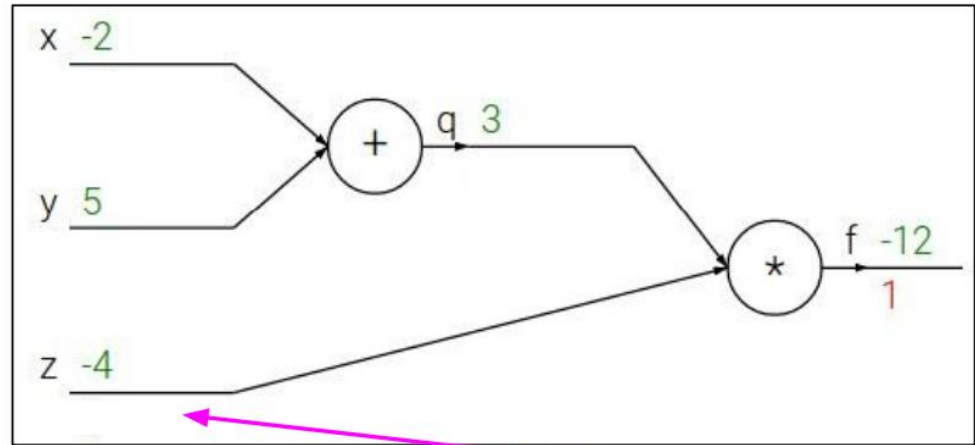
Backpropagation: a simple example

$$f(x, y, z) = (x + y)z$$

e.g. x = -2, y = 5, z = -4

$$q = x + y \qquad \frac{\partial q}{\partial x} = 1, \frac{\partial q}{\partial y} = 1$$

$$f = qz \qquad \frac{\partial f}{\partial q} = z, \frac{\partial f}{\partial z} = q$$

Want: $\dfrac{\partial f}{\partial x}, \dfrac{\partial f}{\partial y}, \dfrac{\partial f}{\partial z}$

Backpropagation: a simple example

$$f(x, y, z) = (x + y)z$$

e.g. x = -2, y = 5, z = -4

$$q = x + y \qquad \frac{\partial q}{\partial x} = 1, \frac{\partial q}{\partial y} = 1$$

$$f = qz \qquad \frac{\partial f}{\partial q} = z, \frac{\partial f}{\partial z} = q$$

Want: $\dfrac{\partial f}{\partial x}, \dfrac{\partial f}{\partial y}, \dfrac{\partial f}{\partial z}$

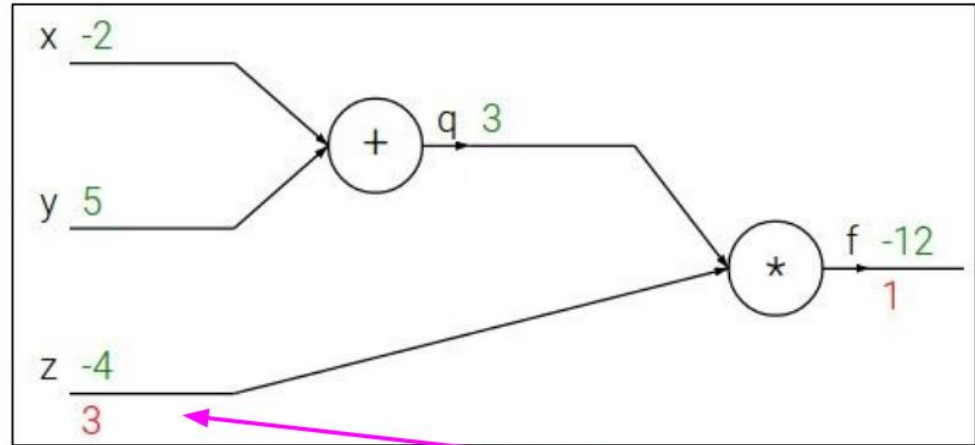$$\frac{\partial f}{\partial z}$$

Backpropagation: a simple example

$$f(x, y, z) = (x + y)z$$

e.g. x = -2, y = 5, z = -4

$$q = x + y \qquad \frac{\partial q}{\partial x} = 1, \frac{\partial q}{\partial y} = 1$$

$$f = qz \qquad \frac{\partial f}{\partial q} = z, \frac{\partial f}{\partial z} = q$$

Want: $\dfrac{\partial f}{\partial x}, \dfrac{\partial f}{\partial y}, \dfrac{\partial f}{\partial z}$

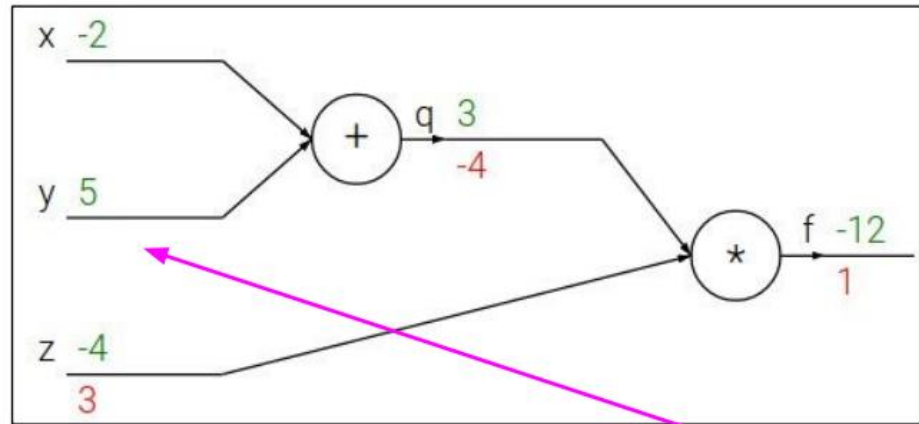$$\frac{\partial f}{\partial z}$$

Backpropagation: a simple example

$$f(x, y, z) = (x + y)z$$

e.g. x = -2, y = 5, z = -4

$$q = x + y \qquad \frac{\partial q}{\partial x} = 1, \frac{\partial q}{\partial y} = 1$$

$$f = qz \qquad \frac{\partial f}{\partial q} = z, \frac{\partial f}{\partial z} = q$$

Want: $\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, \frac{\partial f}{\partial z}$

$$\frac{\partial f}{\partial y}$$

Chain rule:

$$\frac{\partial f}{\partial y} = \frac{\partial f}{\partial q} \frac{\partial q}{\partial y}$$

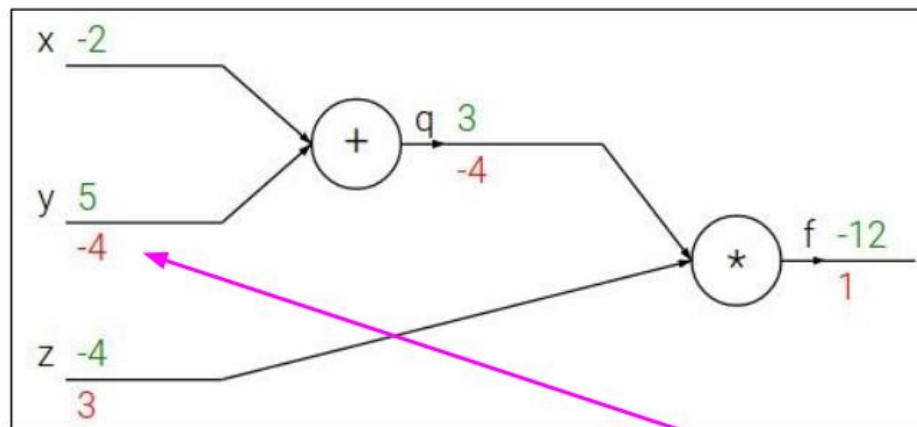Upstream gradient   Local gradient

Backpropagation: a simple example

$$f(x, y, z) = (x + y)z$$

e.g. x = -2, y = 5, z = -4

$$q = x + y \qquad \frac{\partial q}{\partial x} = 1, \frac{\partial q}{\partial y} = 1$$

$$f = qz \qquad \frac{\partial f}{\partial q} = z, \frac{\partial f}{\partial z} = q$$

Want: $\dfrac{\partial f}{\partial x}, \dfrac{\partial f}{\partial y}, \dfrac{\partial f}{\partial z}$

$$\frac{\partial f}{\partial y}$$

Chain rule:

$$\frac{\partial f}{\partial y} = \frac{\partial f}{\partial q} \frac{\partial q}{\partial y}$$

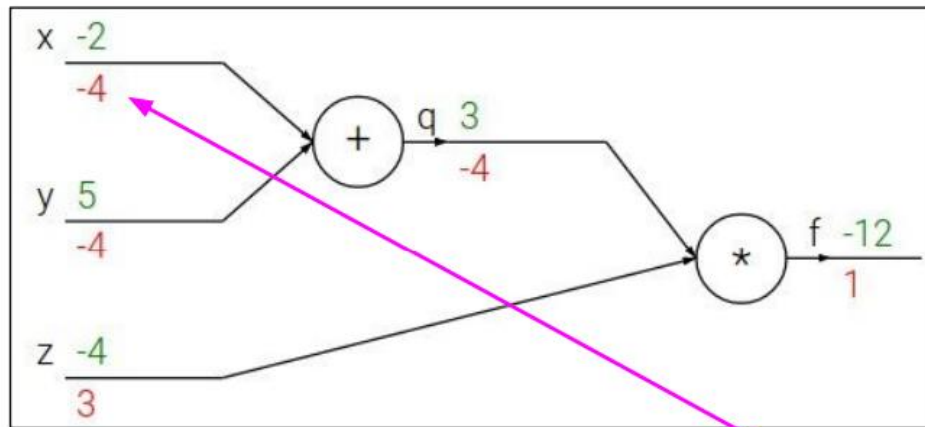Upstream gradient   Local gradient

Backpropagation: a simple example

$$f(x, y, z) = (x + y)z$$

e.g. x = -2, y = 5, z = -4

$$q = x + y \qquad \frac{\partial q}{\partial x} = 1, \frac{\partial q}{\partial y} = 1$$

$$f = qz \qquad \frac{\partial f}{\partial q} = z, \frac{\partial f}{\partial z} = q$$

Want: $\frac{\partial f}{\partial x}, \frac{\partial f}{\partial y}, \frac{\partial f}{\partial z}$
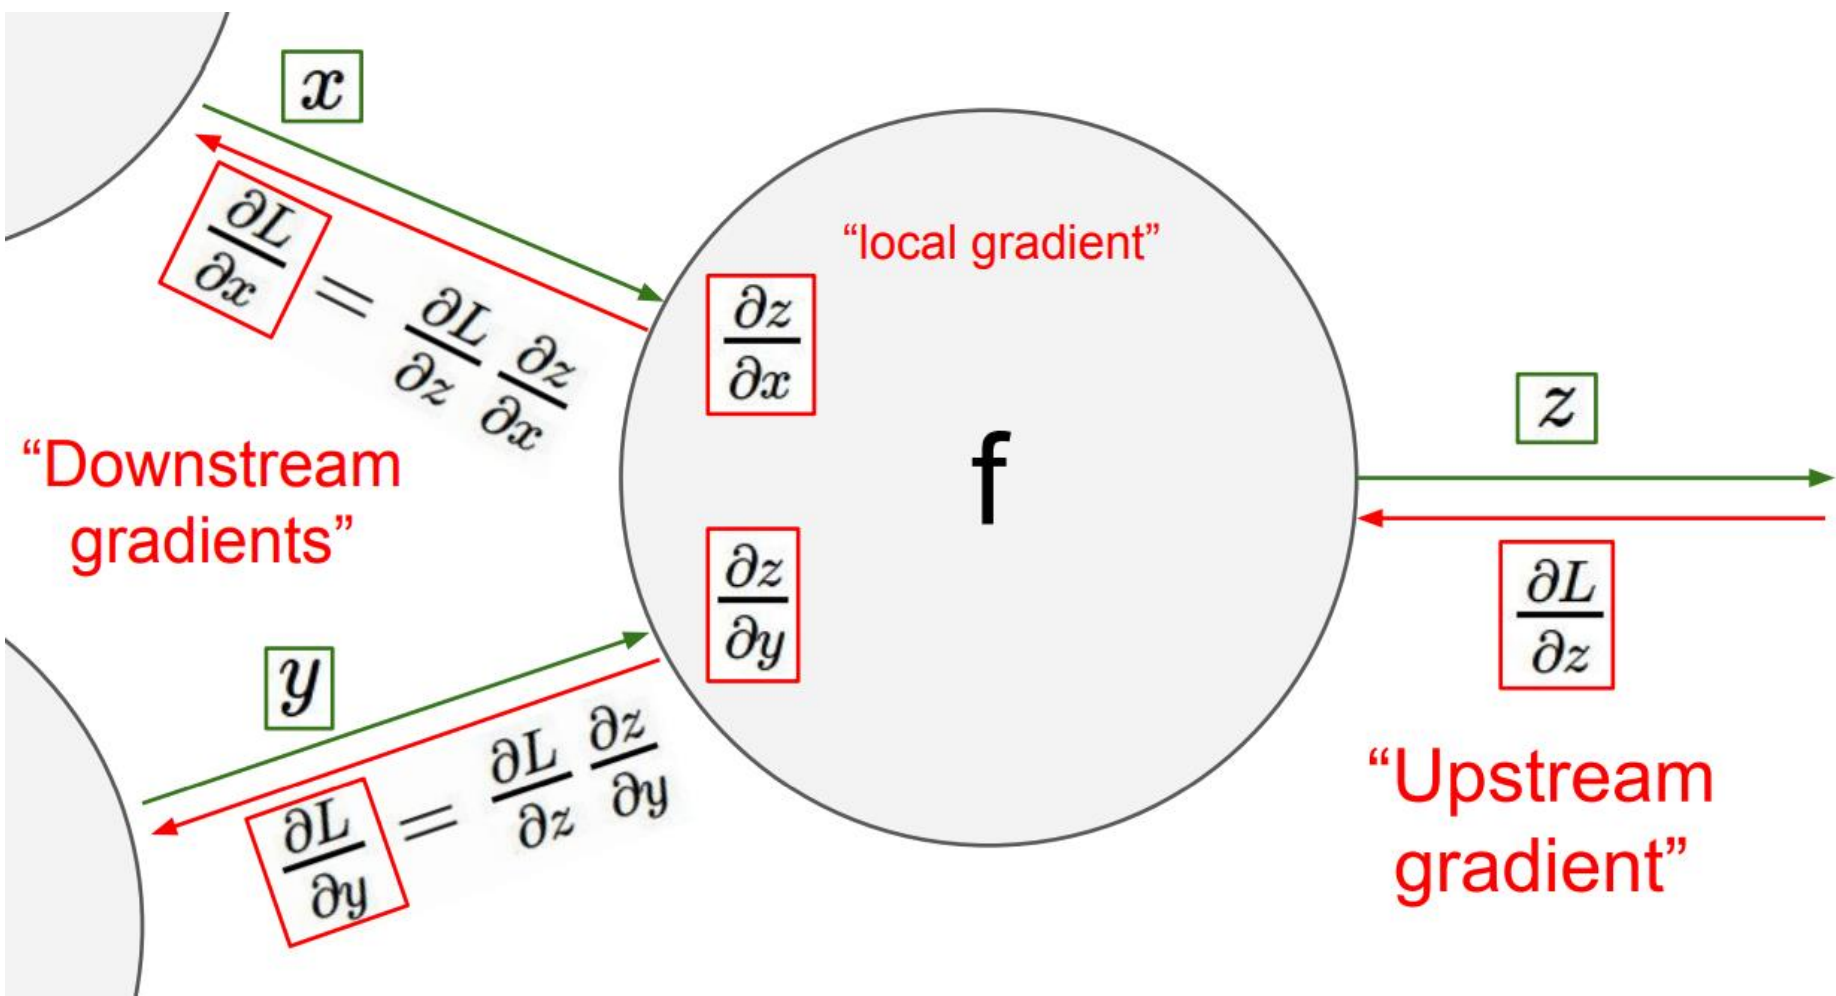

$$\frac{\partial f}{\partial x}$$

Chain rule:

$$\frac{\partial f}{\partial x} = \frac{\partial f}{\partial q} \frac{\partial q}{\partial x}$$
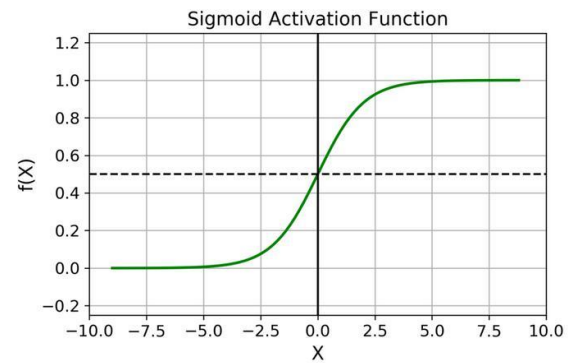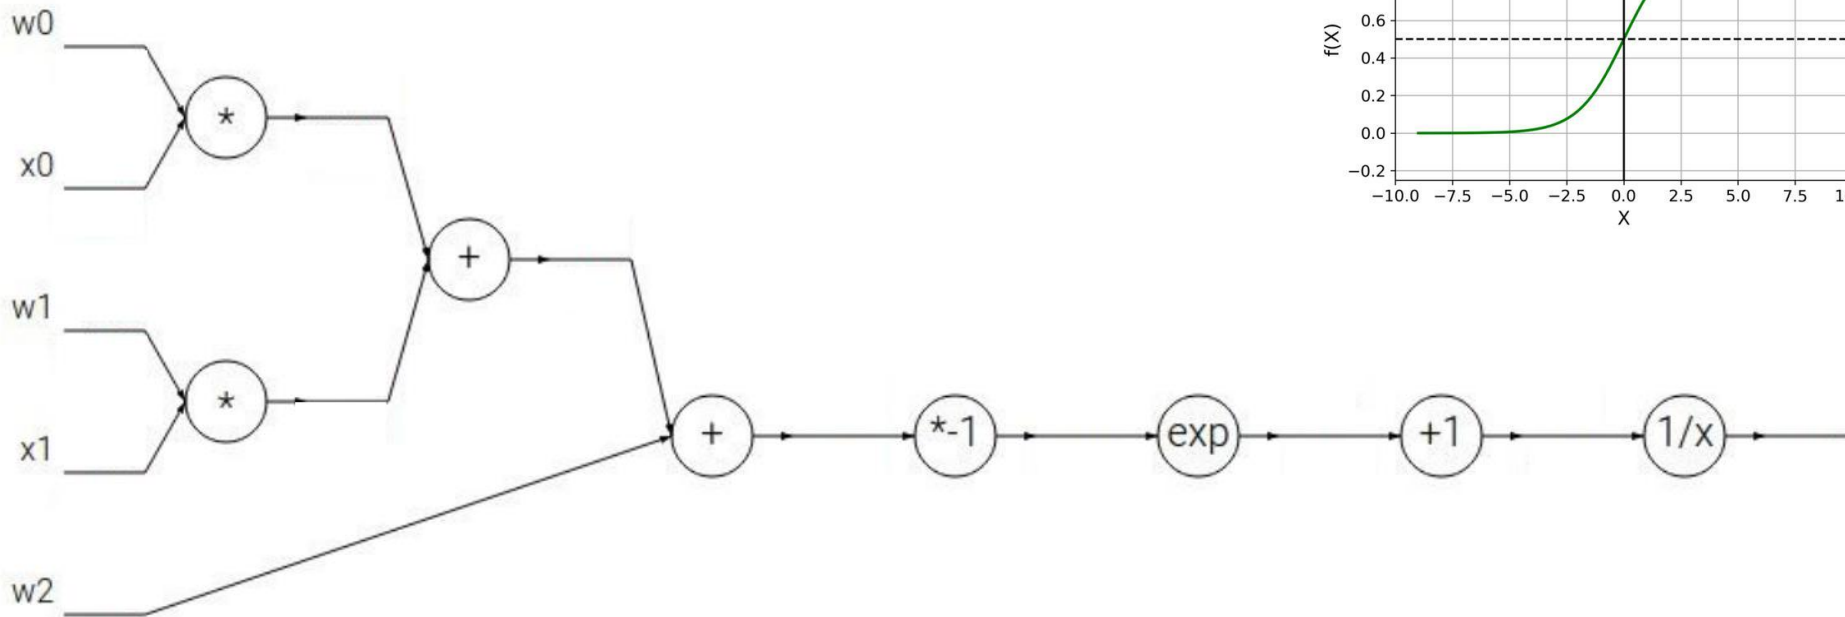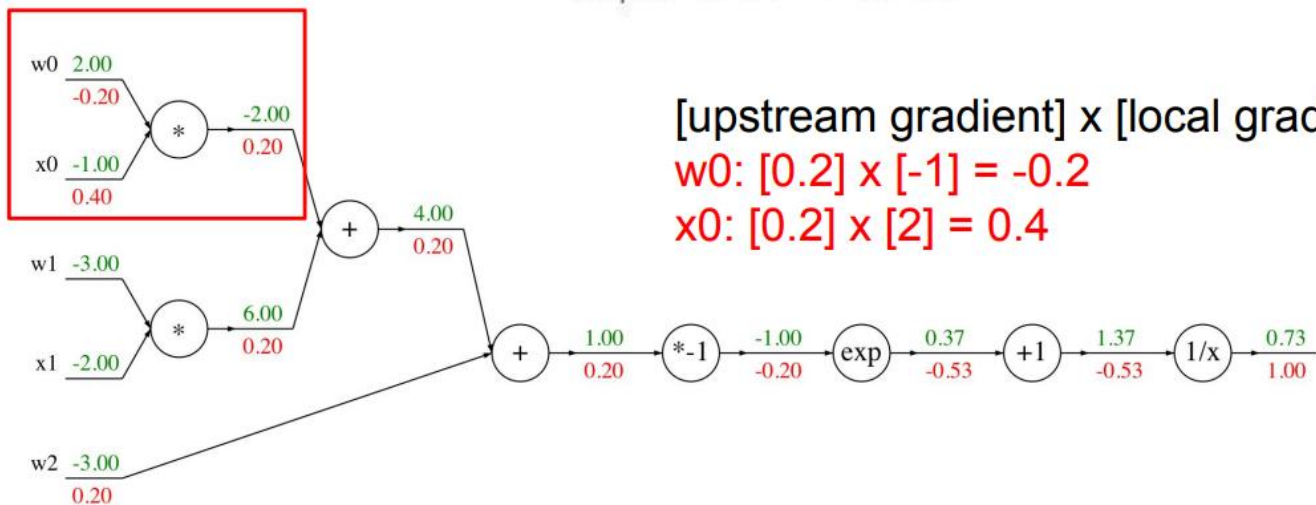
Upstream gradient   Local gradient

x

$$\boxed{\frac{\partial L}{\partial x}} = \frac{\partial L}{\partial z} \frac{\partial z}{\partial x}$$

"Downstream gradients"

y

$$\boxed{\frac{\partial L}{\partial y}} = \frac{\partial L}{\partial z} \frac{\partial z}{\partial y}$$

"local gradient"

$$\boxed{\frac{\partial z}{\partial x}}$$

f

$$\boxed{\frac{\partial z}{\partial y}}$$

z

$$\boxed{\frac{\partial L}{\partial z}}$$

"Upstream gradient"

Another example

$$f(w, x) = \frac{1}{1 + e^{-(w_0 x_0 + w_1 x_1 + w_2)}}$$

Another example

$$f(w, x) = \frac{1}{1 + e^{-(w_0 x_0 + w_1 x_1 + w_2)}}$$

[upstream gradient] x [local gradient]
w0: [0.2] x [-1] = -0.2
x0: [0.2] x [2] = 0.4

w0 2.00
     -0.20
x0 -1.00
     0.40

w1 -3.00
x1 -2.00

w2 -3.00
     0.20

-2.00
 0.20

6.00
 0.20

4.00
 0.20

1.00
 0.20

-1.00
-0.20

0.37
-0.53

1.37
-0.53

0.73
 1.00

$f(x) = e^x \qquad \rightarrow \qquad \dfrac{df}{dx} = e^x$

$f_a(x) = ax \qquad \rightarrow \qquad \dfrac{df}{dx} = a$

$f(x) = \dfrac{1}{x} \qquad \rightarrow \qquad \dfrac{df}{dx} = -1/x^2$

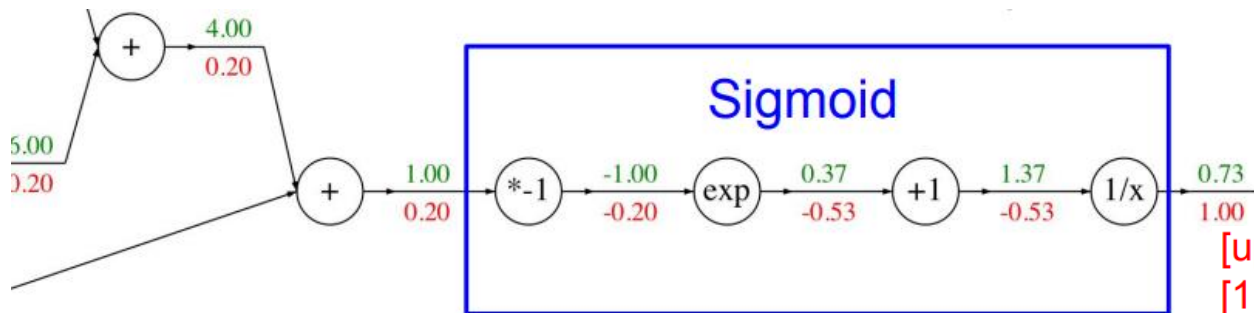$f_c(x) = c + x \qquad \rightarrow \qquad \dfrac{df}{dx} = 1$

Another example

$$f(w,x) = \frac{1}{1 + e^{-(w_0x_0 + w_1x_1 + w_2)}}$$

Sigmoid function

$$\sigma(x) = \frac{1}{1 + e^{-x}}$$

Computational graph representation may not be unique. Choose one where local gradients at each node can be easily expressed!
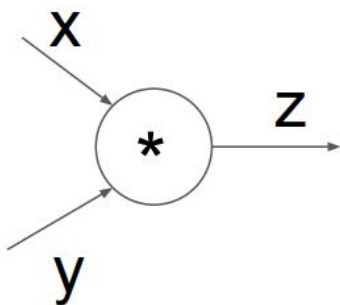

[upstream gradient] x [local gradient]
[1.00] x [(1 - 0.73) (0.73)] = 0.2

Sigmoid local gradient:

$$\frac{d\sigma(x)}{dx} = \frac{e^{-x}}{(1 + e^{-x})^2} = \left(\frac{1 + e^{-x} - 1}{1 + e^{-x}}\right)\left(\frac{1}{1 + e^{-x}}\right) = (1 - \sigma(x))\,\sigma(x)$$

Modularized implementation: forward / backward API

Gate / Node / Function object: Actual PyTorch code


```python
class Multiply(torch.autograd.Function):
  @staticmethod
  def forward(ctx, x, y):
    ctx.save_for_backward(x, y)
    z = x * y
    return z
  @staticmethod
  def backward(ctx, grad_z):
    x, y = ctx.saved_tensors
    grad_x = y * grad_z    # dz/dx * dL/dz
    grad_y = x * grad_z    # dz/dy * dL/dz
    return grad_x, grad_y
```

x

z

*

y

(x,y,z are scalars)

Need to cache some values for use in backward

Upstream gradient

Multiply upstream and local gradients

PyTorch sigmoid layer

Forward actually defined elsewhere...

```
#ifndef TH_GENERIC_FILE
#define TH_GENERIC_FILE "THNN/generic/Sigmoid.c"
#else

void THNN_(Sigmoid_updateOutput)(
        THNNState *state,
        THTensor *input,
        THTensor *output)
{
  THTensor_(sigmoid)(output, input);
}
```

Forward

$$\sigma(x) = \frac{1}{1 + e^{-x}}$$

```
static void sigmoid_kernel(TensorIterator& iter) {
  AT_DISPATCH_FLOATING_TYPES(iter.dtype(), "sigmoid_cpu", [&]() {
    unary_kernel_vec(
        iter,
        [=](scalar_t a) -> scalar_t { return (1 / (1 + std::exp((-a)))); },
        [=](Vec256<scalar_t> a) {
          a = Vec256<scalar_t>((scalar_t)(0)) - a;
          a = a.exp();
          a = Vec256<scalar_t>((scalar_t)(1)) + a;
          a = a.reciprocal();
```

```
void THNN_(Sigmoid_updateGradInput)(
        THNNState *state,
        THTensor *gradOutput,
        THTensor *gradInput,
        THTensor *output)
{
  THNN_CHECK_NELEMENT(output, gradOutput);
  THTensor_(resizeAs)(gradInput, output);
  TH_TENSOR_APPLY3(scalar_t, gradInput, scalar_t, gradOutput, scalar_t, output,
    scalar_t z = *output_data;
    *gradInput_data = *gradOutput_data * (1. - z) * z;
  );
}

#endif
```

Backward

$$(1 - \sigma(x))\,\sigma(x)$$

Common Challenges in Backward Propagation

- Vanishing Gradients

- Exploding Gradient

- Overfitting

- Local Minima

- Gradient Descent Variants

- Training Time

- Poor Initialization

Summary:
- Backward propagation is a critical but challenging step in training neural networks
- Addressing these issues requires a combination of architectural choices, optimization techniques, and regularization methods.

Trap1: Vanishing gradients on sigmoids
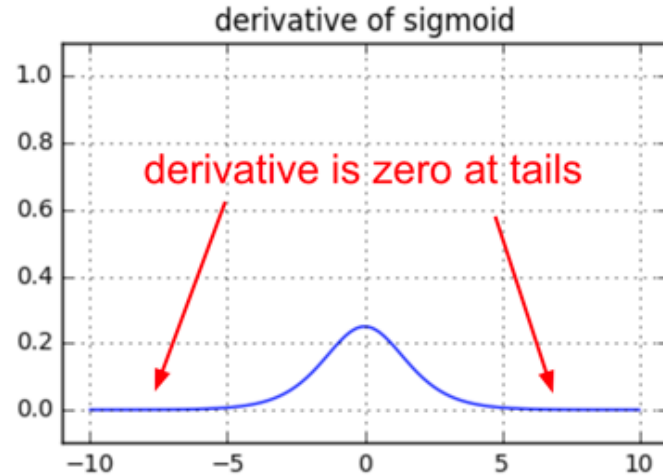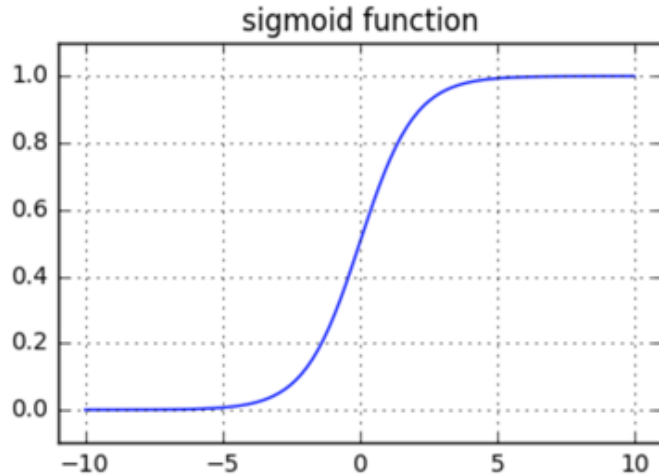
if you're using **sigmoids** or **tanh** non-linearities in your network and you understand backpropagation you should always be nervous about making sure that the initialization doesn't cause them to be fully saturated.

Trap2: Dying ReLUs
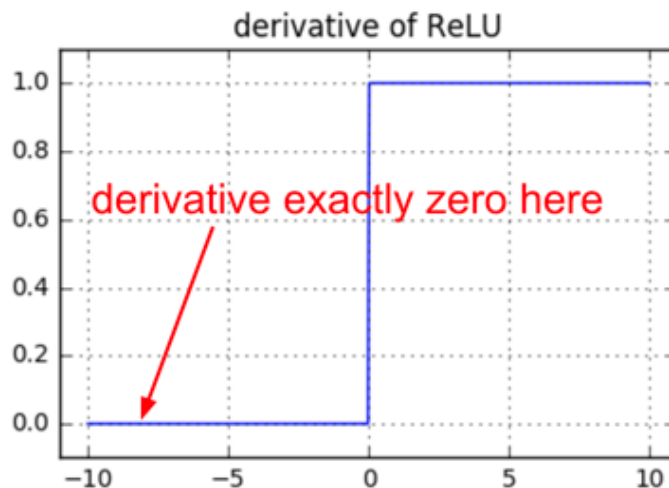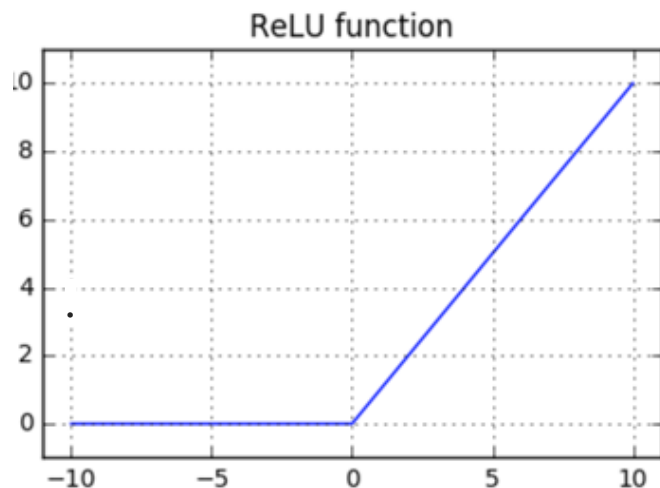
If you understand backpropagation and your network has ReLUs, you're always nervous about dead ReLUs. These are neurons that never turn on for any example in your entire training set and will remain permanently dead. Neurons can also die during training, usually as a symptom of aggressive learning rates.

Trap3: Exploding gradients in RNNs

```python
H = 5      # dimensionality of hidden state
T = 50     # number of time steps
Whh = np.random.randn(H,H)

# forward pass of an RNN (ignoring inputs x)
hs = {}
ss = {}
hs[-1] = np.random.randn(H)
for t in xrange(T):
    ss[t] = np.dot(Whh, hs[t-1])
    hs[t] = np.maximum(0, ss[t])

# backward pass of the RNN
dhs = {}
dss = {}
dhs[T-1] = np.random.randn(H) # start off the chain with random gradient
for t in reversed(xrange(T)):
    dss[t] = (hs[t] > 0) * dhs[t] # backprop through the nonlinearity
    dhs[t-1] = np.dot(Whh.T, dss[t]) # backprop into previous hidden state
```

if the largest eigenvalue is > 1, gradient will explode
if the largest eigenvalue is < 1, gradient will vanish

If you understand backpropagation and you're using RNNs you are nervous about having to do gradient clipping, or you prefer to use an LSTM.

Review: MLP

1. Breif introduction of MLP;
2. Forward propagation and backward propagation;
3. Common Challenges in Backward Propagation
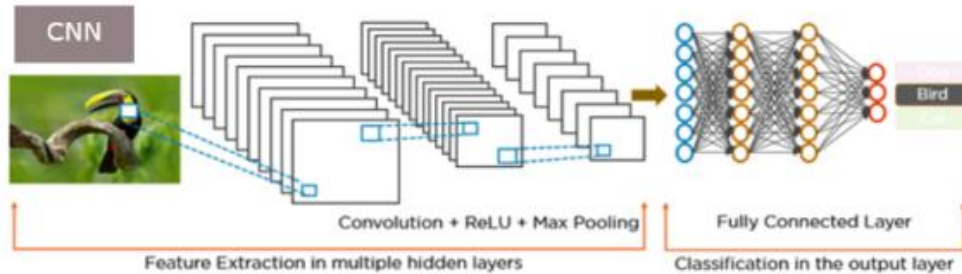
Limitations of MLP:

1. Limited Spatial Invariance (vs. CNNs)

2. Sequential Information Handling (vs. RNNs)

3. Positional Encoding (vs. Transformers)

4. Attention Mechanism (vs. Transformers)

5. Hierarchical Feature Extraction (vs. CNNs and Transformers)

6. Parameter Efficiency (vs. Transformers)

7. Pre-training Efficiency (vs. Transformers)

8. Structured Input Bias (vs. CNNs and Transformers)
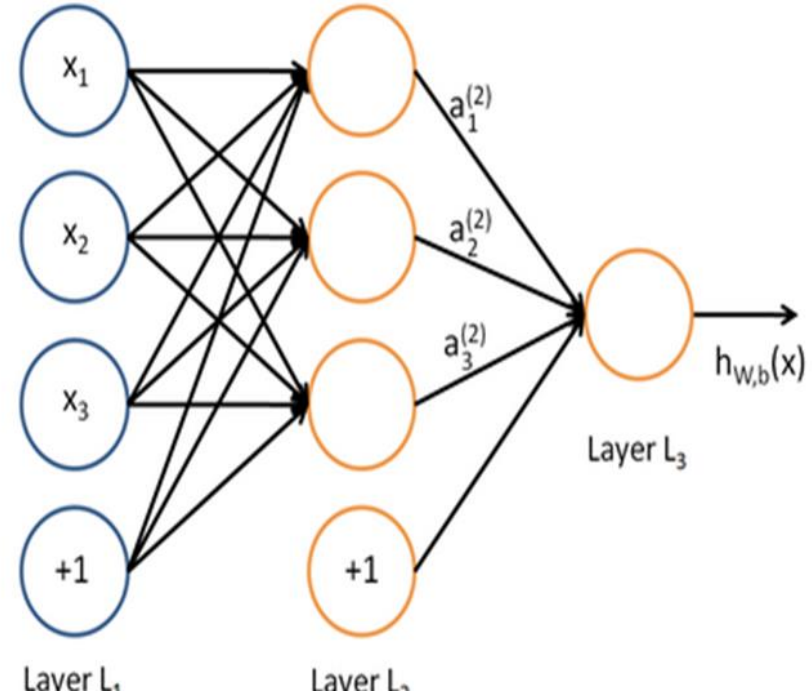
CNN&RNN

- Convolutional Neural Network (CNN)
- Recurrent Neural Network (RNN)

Today's lecture

- MLP
  - +: Strongest inductive bias: if all words are concated
  - +: Weakest inductive bias: if all words are averaged
  - - : The interaction at the token-level is too weak
- **CNN & RNN**
  - +: The interaction at the token-level is slightly better.
    - CNN: Bringing the global token-level interaction to the window-level
      - - : Make simplifications, its global dependencies are limited
    - RNN: An ideal method for processing token sequences
      - - : Its recursive nature has the problem of disaster forgetting.
- Transformer
  - +: Achieve **global dependence** at the **token-level** by **decoupling** token-level interaction and feature-level abstraction into two components, in **SAN** and **FNN**.

- Scaling law and emergent ability

CNN&RNN

CNN
Convolutional Neural Network

- What is CNN?
- Motivation: Image Processing
- Key Components
    - Convolutional Layers
    - Pooling Layers
    - Fully Connected Layers
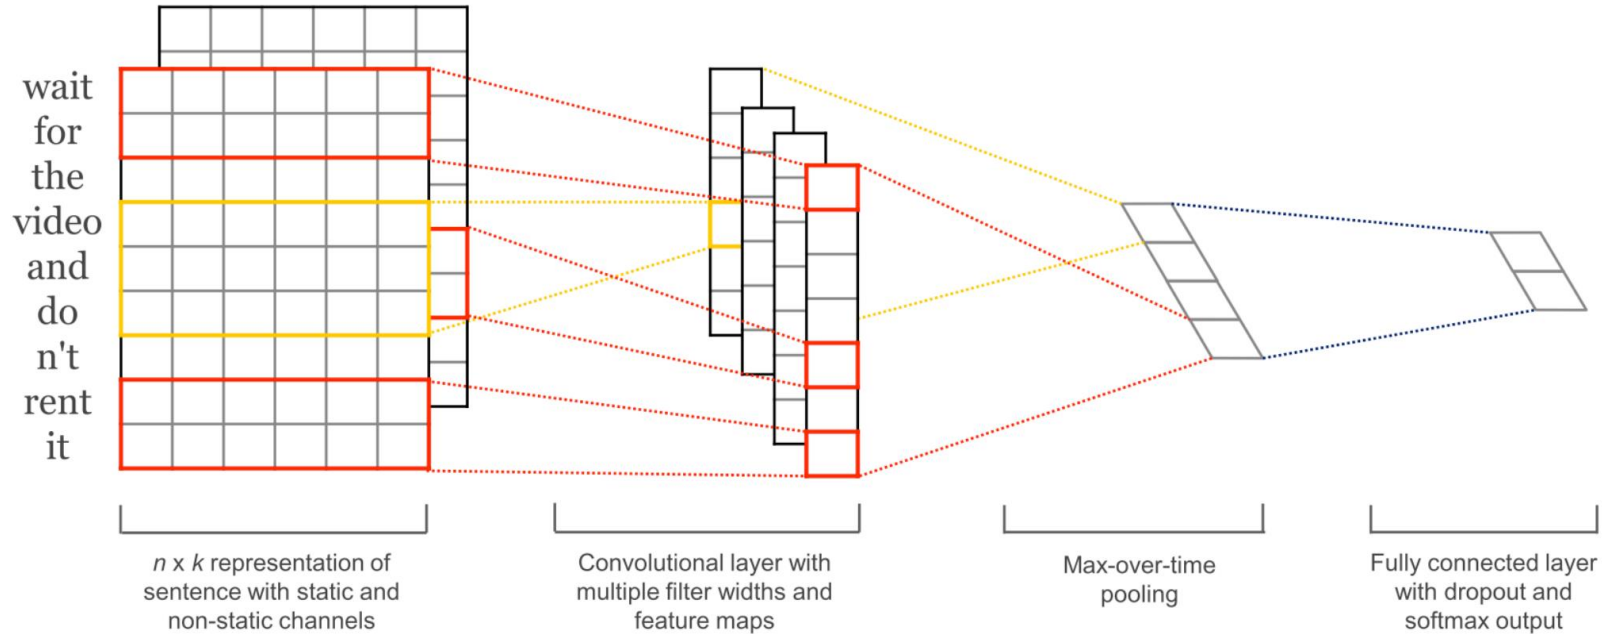- Hierarchical Feature Extraction

A Typical Convolutional Neural Network (CNN)
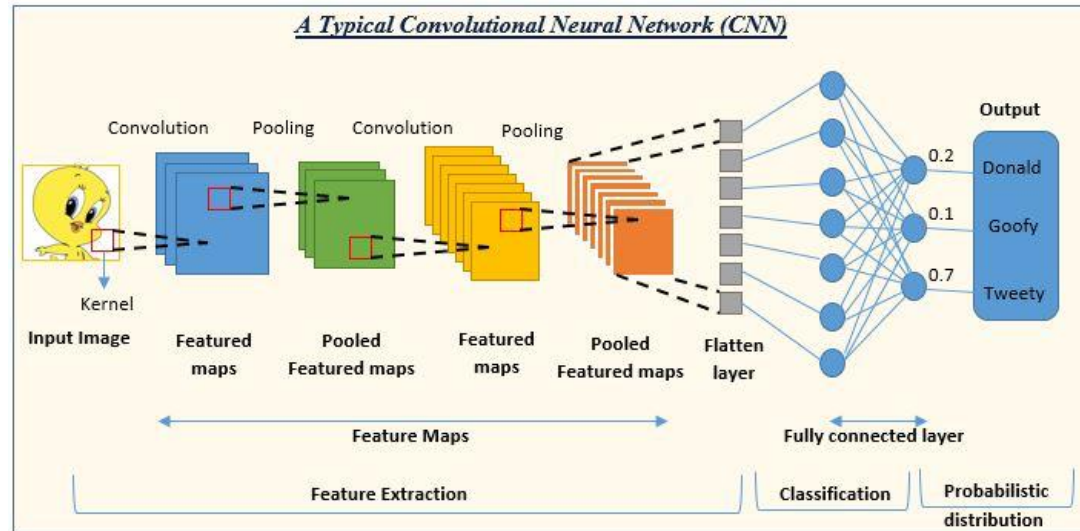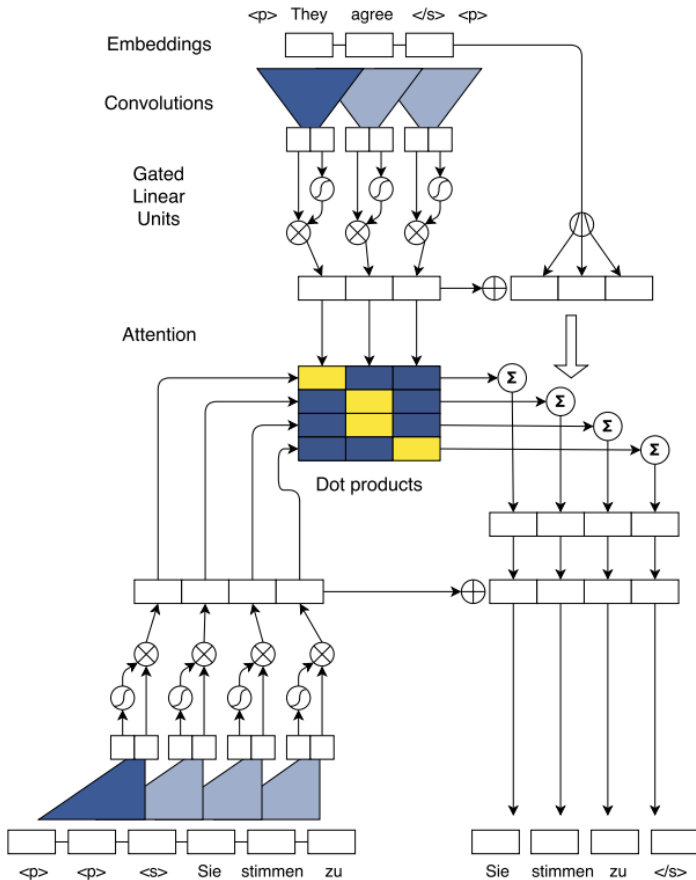
Convolutional NNs in image classification
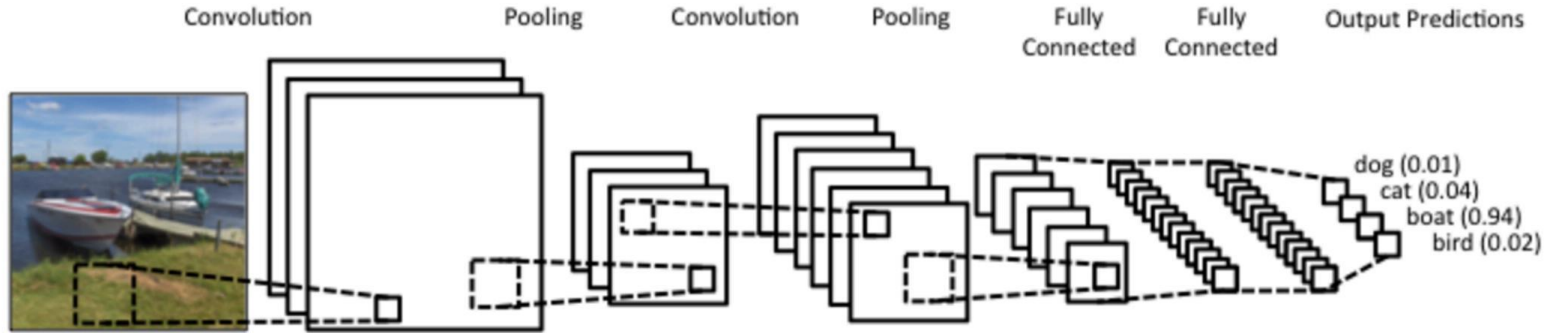
Key components: 1) convolution; 2) pooling; 3) multiple channels (feature maps)

Convolutional NNs for text classification

(Kim 2014): Convolutional Neural Networks for Sentence Classification

Convolutional Sequence to Sequence Learning

❖ Encoder and decoder are simple blocks of convolution operation followed by nonlinearity on fixed size of input.

❖ Introduce a concept of order preservation as a positional vectors p = (p_1,p_2 ...,p_m). In combination of both input elements are represented as E = (e_1=w_1+p_1, e_2=w_2+p_2, ...,e_m=w_m+p_m).

❖ Adds a linear mapping to project between the embedding size f and the convolution outputs that are size 2d.

❖ Computes a distribution over the T possible next target elements y_i+1 by transforming the top decoder output h_i_l via a linear layer with weights and bias.

https://browse.arxiv.org/pdf/1705.03122.pdf

RNN
Recurrent Neural Network

Core idea: Apply the same weights W *repeatedly*

outputs (optional) $\{$

$\hat{y}^{(1)}$ $\hat{y}^{(2)}$ $\hat{y}^{(3)}$ $\hat{y}^{(4)}$ ...

$h^{(1)}$ $h^{(2)}$ $h^{(3)}$ $h^{(4)}$

hidden states $\{$

W W W W ...

input sequence (any length) $\{$

$x^{(1)}$ $x^{(2)}$ $x^{(3)}$ $x^{(4)}$ ...
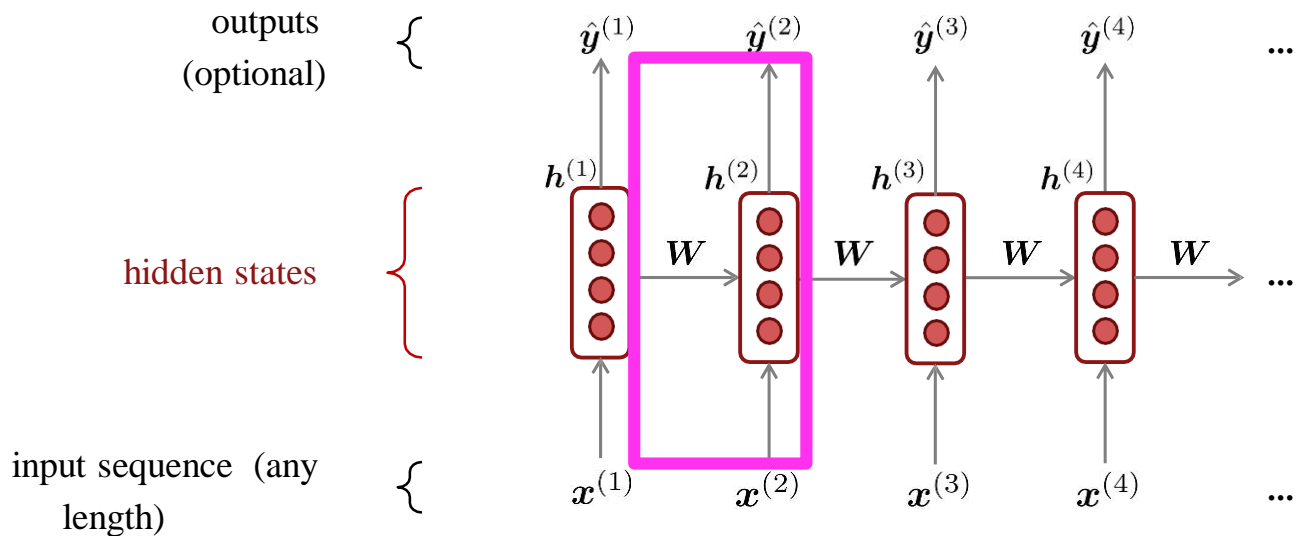
A Simple RNN Language Model

$$\hat{y}^{(4)} = P(x^{(5)}|\text{the students opened their})$$

output distribution

$$\hat{y}^{(t)} = \text{softmax}\left(Uh^{(t)} + b_2\right) \in \mathbb{R}^{|V|}$$
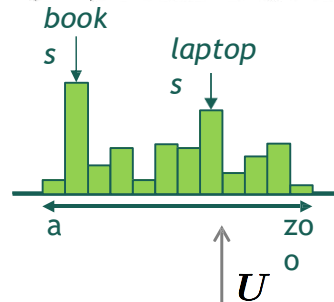
hidden states

$$h^{(t)} = \sigma\left(W_h h^{(t-1)} + W_e e^{(t)} + b_1\right)$$

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

word embeddings

$$e^{(t)} = Ex^{(t)}$$

words / one-hot vectors

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

the $x^{(1)}$ *students* $x^{(2)}$ *opened* $x^{(3)}$ *their* $x^{(4)}$
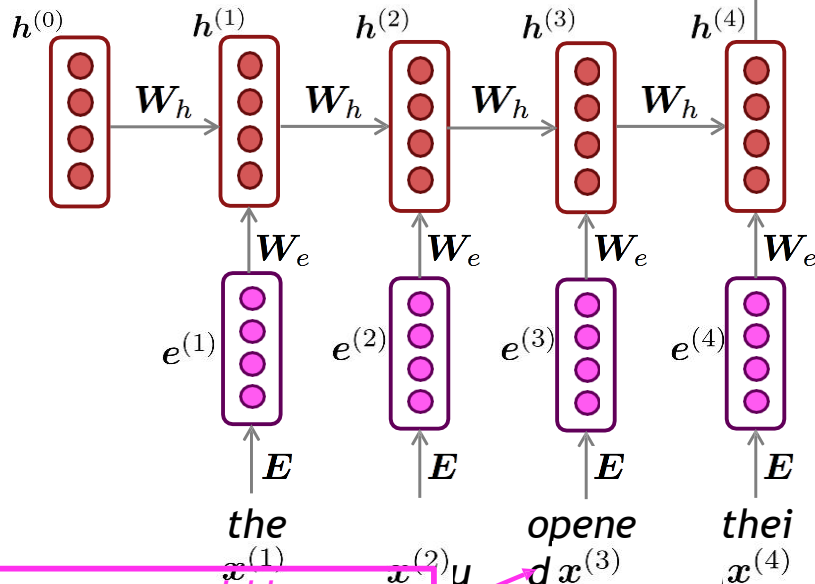
Note: this input sequence could be much longer now!

RNN Language Models

RNN **Advantages**:
- Can process any length input
- Computation for step t can (in theory) use information from many steps back
- Model size doesn't increase for longer input context
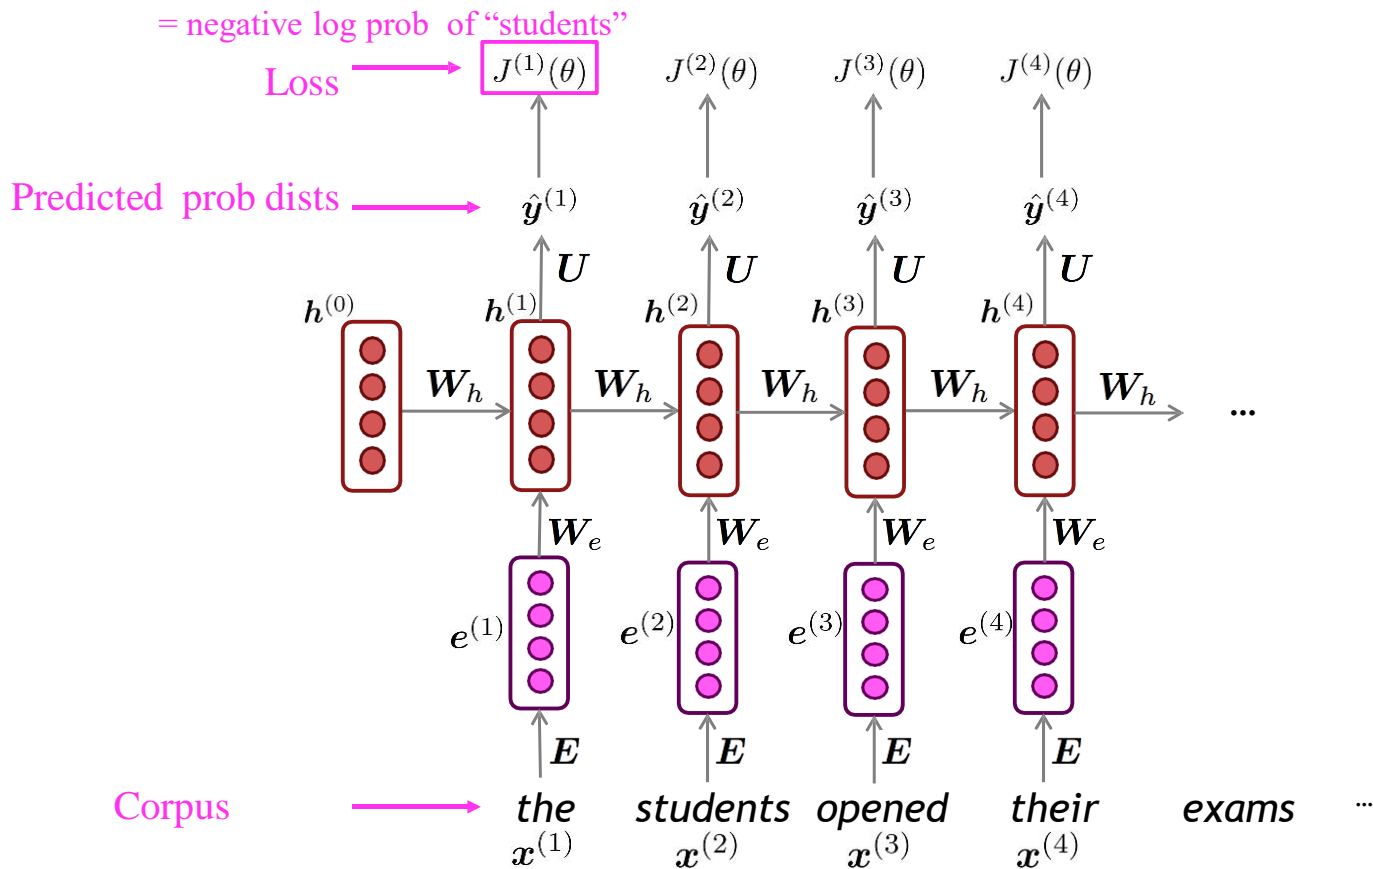- Same weights applied on every timestep, so there is symmetry in how inputs are processed.

RNN **Disadvantages**:
- Recurrent computation is slow
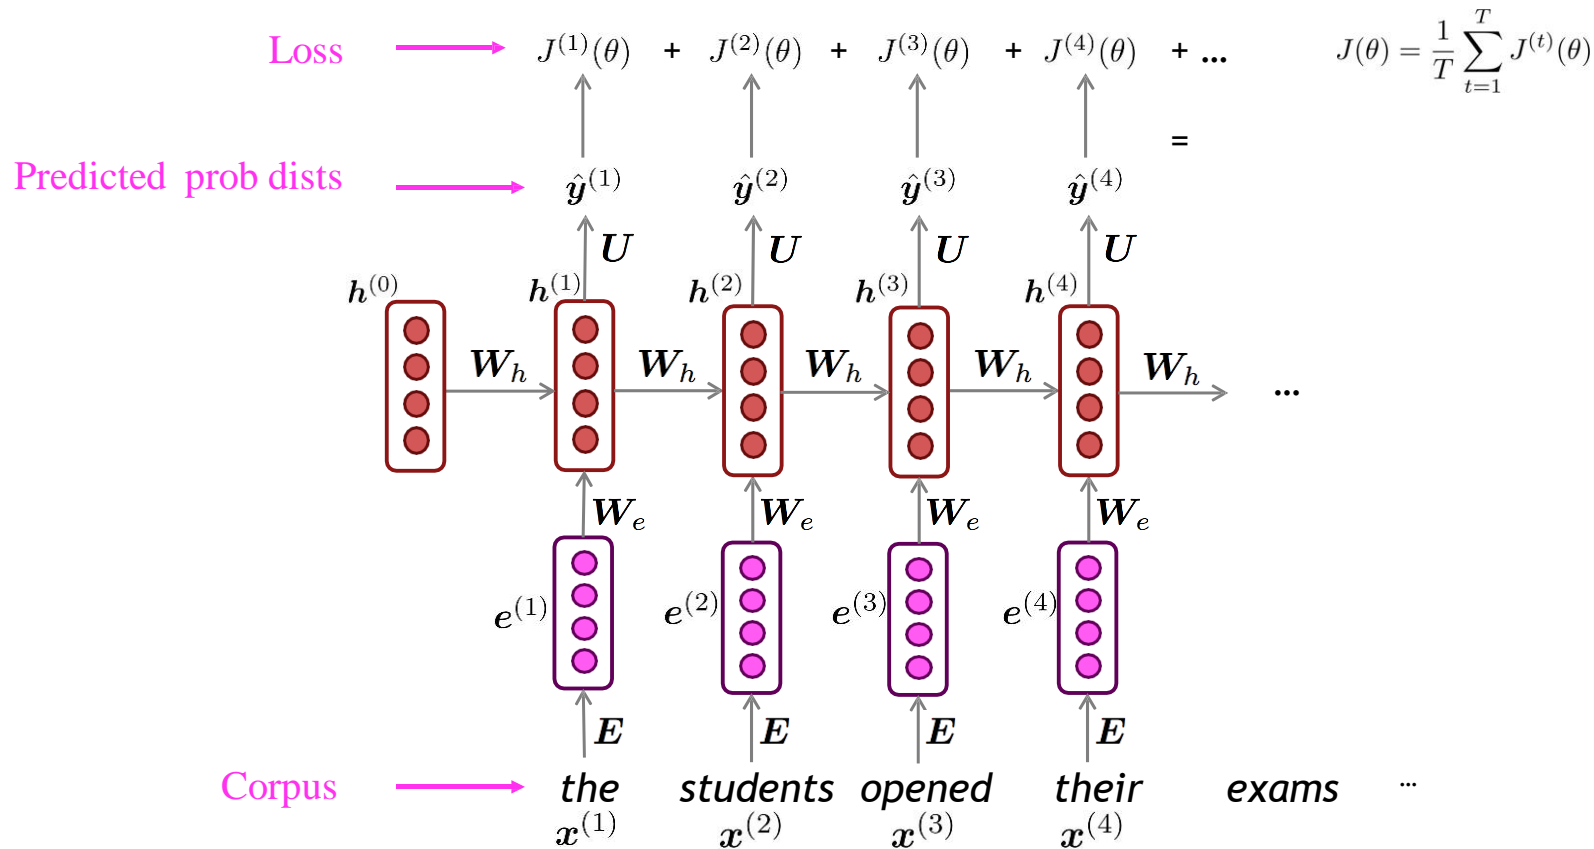- In practice, difficult to access information from many steps back
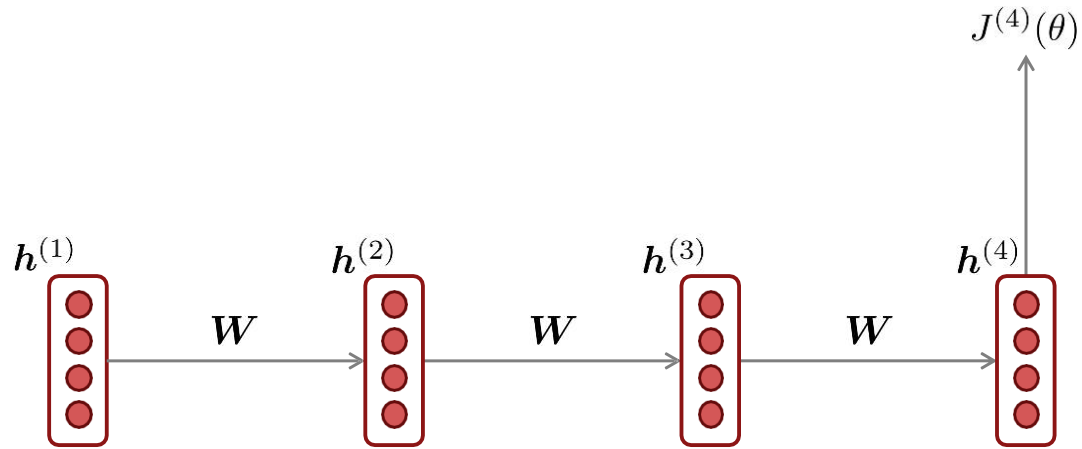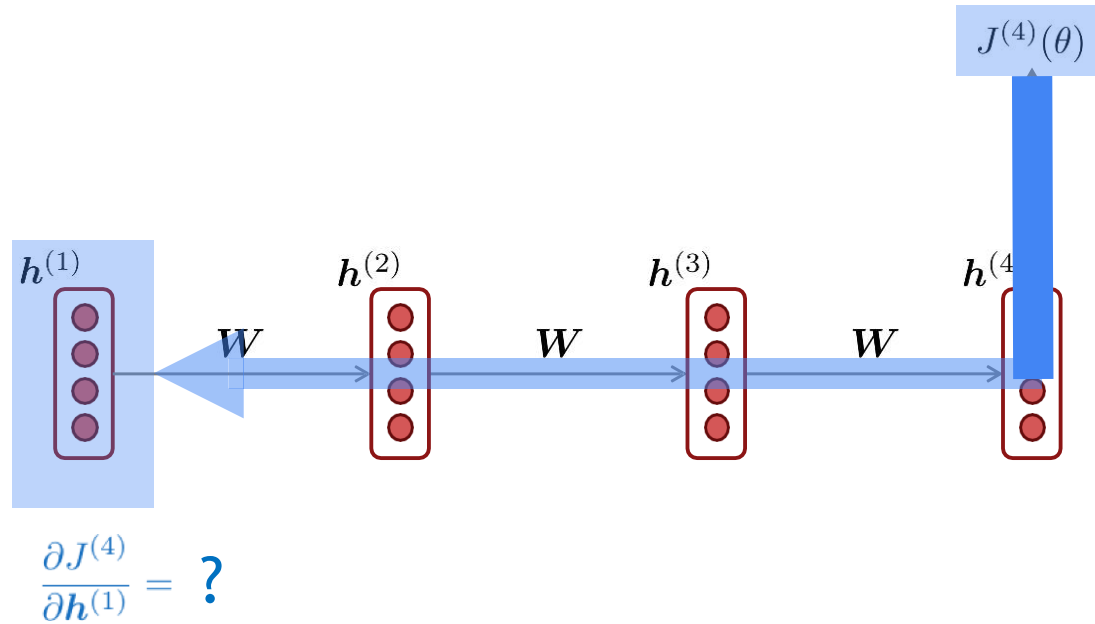
More on these later
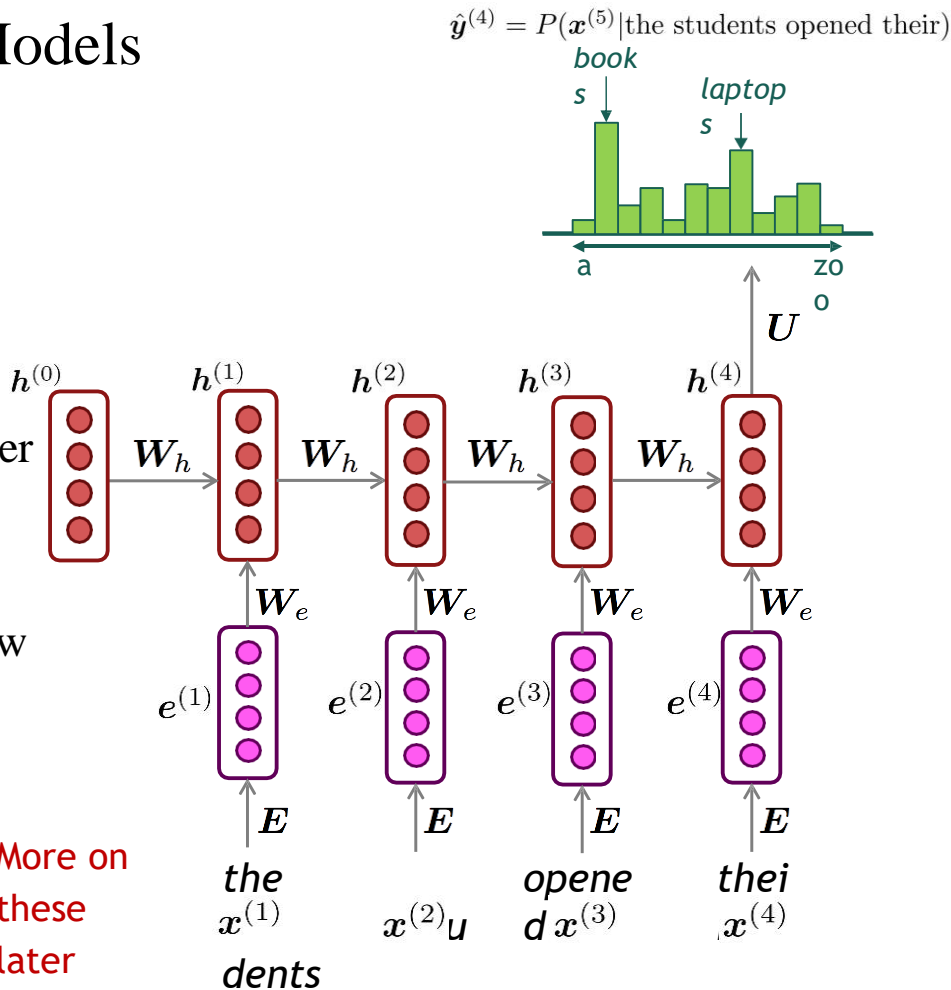
Training an RNN Language Model

Training an RNN Language Model

"Teacher forcing"

Loss $\longrightarrow$ $J^{(1)}(\theta)$ + $J^{(2)}(\theta)$ + $J^{(3)}(\theta)$ + $J^{(4)}(\theta)$ + ...

$$J(\theta) = \frac{1}{T}\sum_{t=1}^{T} J^{(t)}(\theta)$$

=

Predicted prob dists $\longrightarrow$ $\hat{\boldsymbol{y}}^{(1)}$ $\hat{\boldsymbol{y}}^{(2)}$ $\hat{\boldsymbol{y}}^{(3)}$ $\hat{\boldsymbol{y}}^{(4)}$

$\boldsymbol{U}$ $\boldsymbol{U}$ $\boldsymbol{U}$ $\boldsymbol{U}$

$\boldsymbol{h}^{(0)}$ $\boldsymbol{h}^{(1)}$ $\boldsymbol{h}^{(2)}$ $\boldsymbol{h}^{(3)}$ $\boldsymbol{h}^{(4)}$

$\boldsymbol{W}_h$ $\boldsymbol{W}_h$ $\boldsymbol{W}_h$ $\boldsymbol{W}_h$ $\boldsymbol{W}_h$ ...

$\boldsymbol{W}_e$ $\boldsymbol{W}_e$ $\boldsymbol{W}_e$ $\boldsymbol{W}_e$

$\boldsymbol{e}^{(1)}$ $\boldsymbol{e}^{(2)}$ $\boldsymbol{e}^{(3)}$ $\boldsymbol{e}^{(4)}$

$\boldsymbol{E}$ $\boldsymbol{E}$ $\boldsymbol{E}$ $\boldsymbol{E}$

Corpus $\longrightarrow$ the students opened their exams ...

$\boldsymbol{x}^{(1)}$ $\boldsymbol{x}^{(2)}$ $\boldsymbol{x}^{(3)}$ $\boldsymbol{x}^{(4)}$

Problems with RNNs: Vanishing and Exploding Gradients

Vanishing gradient intuition

$J^{(4)}(\theta)$

$h^{(1)}$ $h^{(2)}$ $h^{(3)}$ $h^{(4)}$

W W W

$$\frac{\partial J^{(4)}}{\partial h^{(1)}} = \text{?}$$

Vanishing gradient intuition
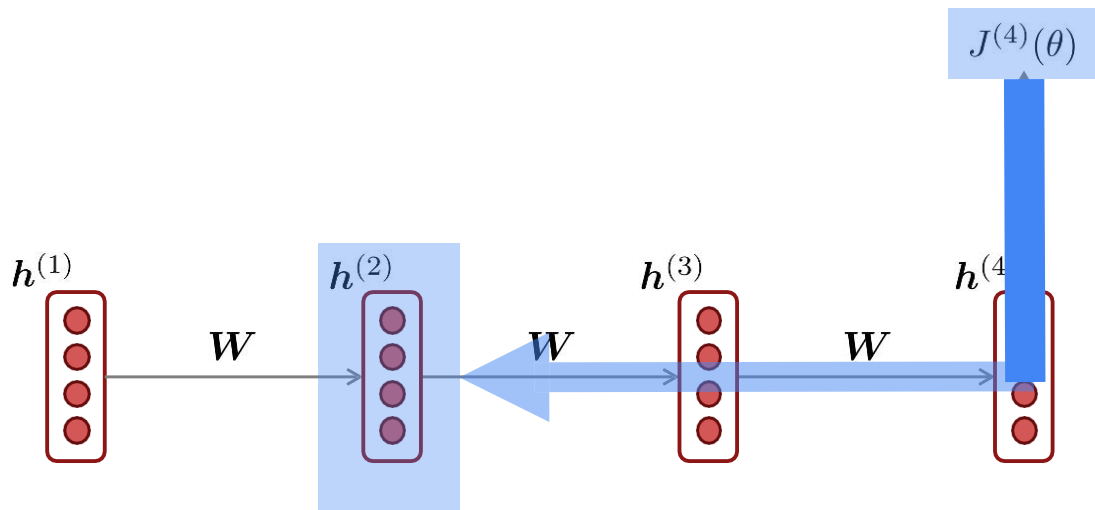
$$\frac{\partial J^{(4)}}{\partial \boldsymbol{h}^{(1)}} = \frac{\partial \boldsymbol{h}^{(2)}}{\partial \boldsymbol{h}^{(1)}} \times \frac{\partial J^{(4)}}{\partial \boldsymbol{h}^{(2)}}$$

chain rule!

Vanishing gradient intuition

$$\frac{\partial J^{(4)}}{\partial \boldsymbol{h}^{(1)}} = \frac{\partial \boldsymbol{h}^{(2)}}{\partial \boldsymbol{h}^{(1)}} \times \quad \frac{\partial \boldsymbol{h}^{(3)}}{\partial \boldsymbol{h}^{(2)}} \times \frac{\partial J^{(4)}}{\partial \boldsymbol{h}^{(3)}}$$

chain rule!

Vanishing gradient intuition

$$J^{(4)}(\theta)$$

$\boldsymbol{h}^{(1)}$ $\quad W \quad$ $\boldsymbol{h}^{(2)}$ $\quad W \quad$ $\boldsymbol{h}^{(3)}$ $\quad W \quad$ $\boldsymbol{h}^{(4)}$

$$\frac{\partial J^{(4)}}{\partial \boldsymbol{h}^{(1)}} = \quad \frac{\partial \boldsymbol{h}^{(2)}}{\partial \boldsymbol{h}^{(1)}} \times \qquad \frac{\partial \boldsymbol{h}^{(3)}}{\partial \boldsymbol{h}^{(2)}} \times \qquad \frac{\partial \boldsymbol{h}^{(4)}}{\partial \boldsymbol{h}^{(3)}} \times \quad \frac{\partial J^{(4)}}{\partial \boldsymbol{h}^{(4)}}$$

chain rule!

Vanishing gradient intuition

$J^{(4)}(\theta)$

$h^{(1)}$ $h^{(2)}$ $h^{(3)}$ $h^{(4)}$

W W W
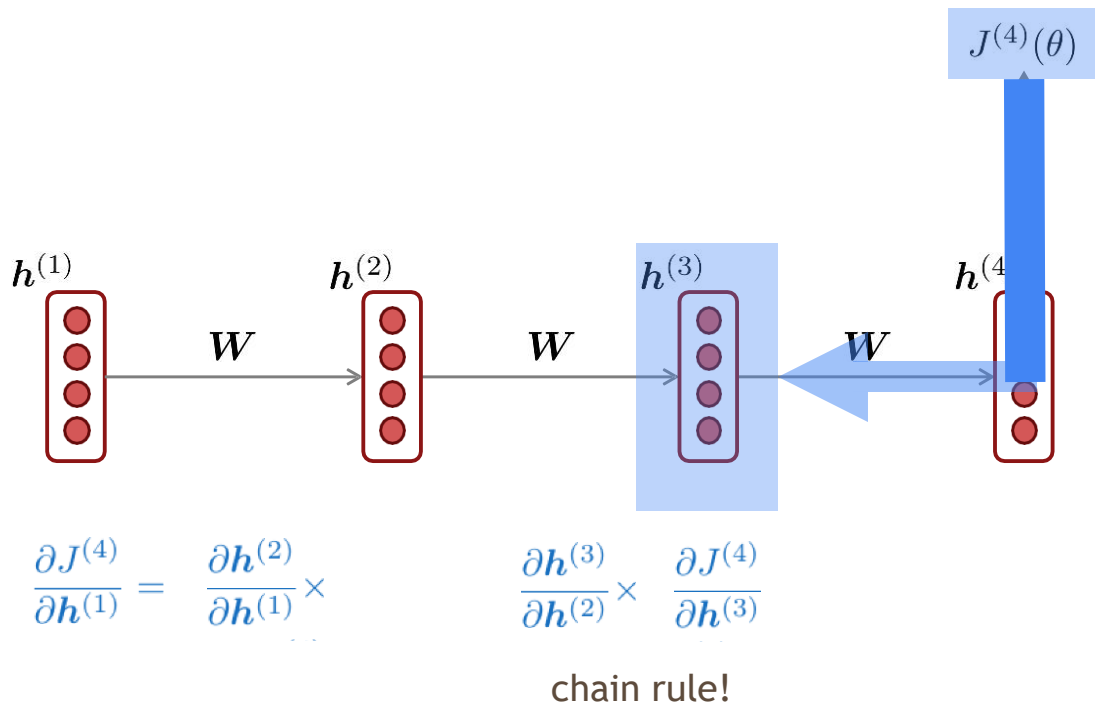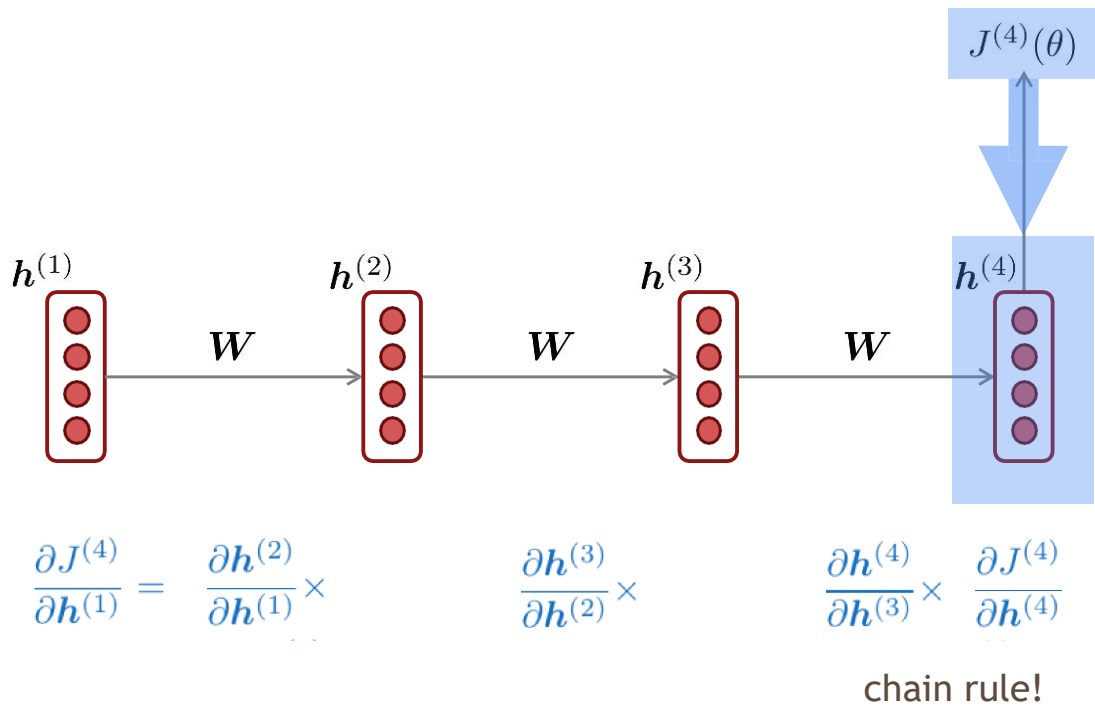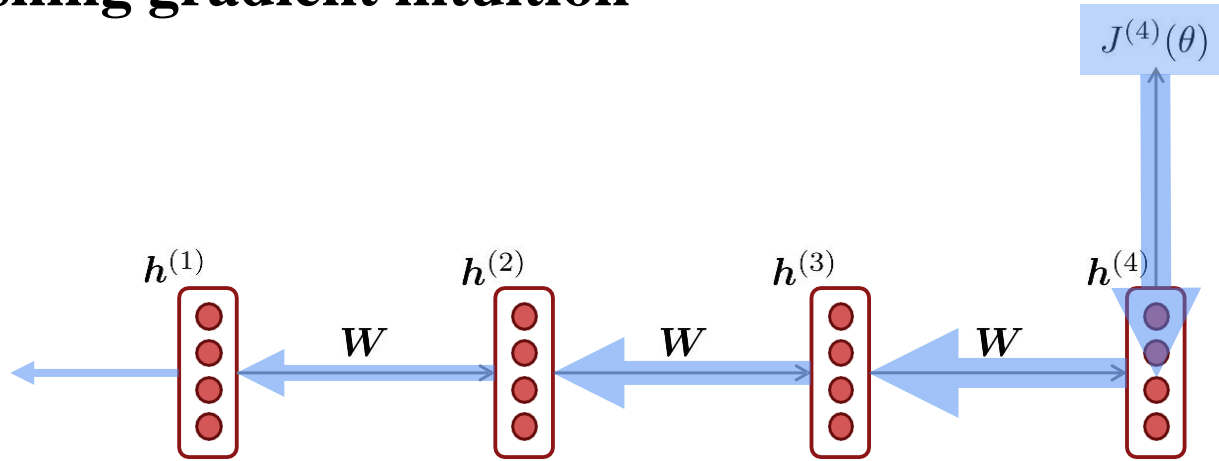
$$\frac{\partial J^{(4)}}{\partial h^{(1)}} = \boxed{\frac{\partial h^{(2)}}{\partial h^{(1)}}} \times \quad \boxed{\frac{\partial h^{(3)}}{\partial h^{(2)}}} \times \quad \boxed{\frac{\partial h^{(4)}}{\partial h^{(3)}}} \times \quad \frac{\partial J^{(4)}}{\partial h^{(4)}}$$

What happens if these are small?

Vanishing gradient problem:
When these are small, the gradient signal gets smaller and smaller as it backpropagates further

Why is exploding gradient a problem?

- If the gradient becomes too big, then the SGD update step becomes too big:

$$\theta^{new} = \theta^{old} - \overbrace{\alpha}^{\text{learning rate}} \underbrace{\nabla_\theta J(\theta)}_{\text{gradient}}$$

- This can cause bad updates: we take too large a step and reach a weird and bad parameter configuration (with large loss)
  - You think you've found a hill to climb, but suddenly you're in Iowa

- In the worst case, this will result in Inf or NaN in your network (then you have to restart training from an earlier checkpoint)

Is vanishing/exploding gradient just an RNN problem?

- No! It can be a problem for all neural architectures (including feed-forward and convolutional), especially very deep ones.
  - Due to chain rule / choice of nonlinearity function, gradient can become vanishingly small as it backpropagates
  - Thus, lower layers are learned very slowly (i.e., are hard to train)
- Another solution: lots of new deep feedforward/convolutional architectures add more direct connections (thus allowing the gradient to flow)

For example:
- Residual connections aka "ResNet"
- Also known as skip-connections
- The identity connection preserves information by default
- This makes deep networks much easier to train
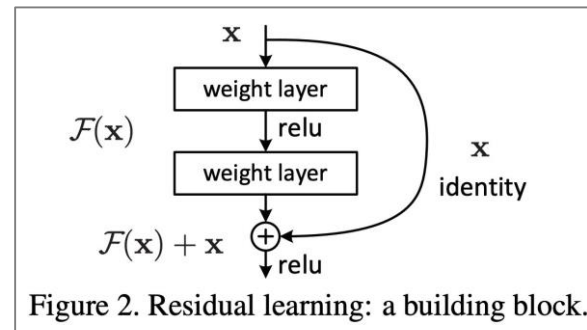
Figure 2. Residual learning: a building block.

"Deep Residual Learning for Image Recognition", He et al, 2015.
https://arxiv.org/pdf/1512.03385.pdf

Transformer

- Encoder

- Decoder

- Self-attention

- Multi-head self-attention

- Positional Encoding

Transformer

- **Encoder**

- **Decoder**

- **Self-attention**

- **Multi-head self-attention**

- **Positional Encoding**

Today's lecture

- MLP
  - +: Strongest inductive bias: if all words are concated
  - +: Weakest inductive bias: if all words are averaged
  - - : The interaction at the token-level is too weak
- CNN & RNN
  - +: The interaction at the token-level is slightly better.
    - CNN: Bringing the global token-level interaction to the window-level
      - - : Make simplifications, its global dependencies are limited
    - RNN: An ideal method for processing token sequences
      - - : Its recursive nature has the problem of disaster forgetting.
- **Transformer**
  - +: Achieve **global dependence** at the **token-level** by **decoupling** token-level interaction and feature-level abstraction into two components, in **SAN** and **FNN**.

- Scaling law and emergent ability

Output
Probabilities

Softmax

Linear

Add & Norm

Feed
Forward

Add & Norm

Multi-Head
Attention

Add & Norm

Masked
Multi-Head
Attention

N×

Positional
Encoding

Output
Embedding

Positional
Encoding

Outputs
(shifted right)

encoder

Add & Norm

Feed
Forward

Add & Norm

Multi-Head
Attention

N×

Positional
Encoding

Input
Embedding

Inputs

Output
Probabilities

Softmax

Linear

decoder

Add & Norm

Feed
Forward

Add & Norm

Multi-Head
Attention

Add & Norm

Masked
Multi-Head
Attention

N×

Positional
Encoding

Output
Embedding

Outputs
(shifted right)

encoder

Add & Norm

Feed
Forward

Add & Norm

Multi-Head
Attention

N×

Positional
Encoding

Input
Embedding

Inputs

Self-attention (in encoder)

Q
K
V

Layer p

Nobel   committee   awards   who   advanced   optics

Self-attention (in encoder)

Q
K
V

Layer p

Nobel   committee   awards   who   advanced

Self-attention (in encoder)

[Vaswani et al. 2017]

Q
K
V

Layer p

Nobel committee awards who advanced

Self-attention (in encoder)

Self-attention (in encoder)

Self-attention (in encoder)

Self-attention (in encoder)

[Vaswani et al. 2017]

Self-attention (in encoder)

Multi-head self-attention
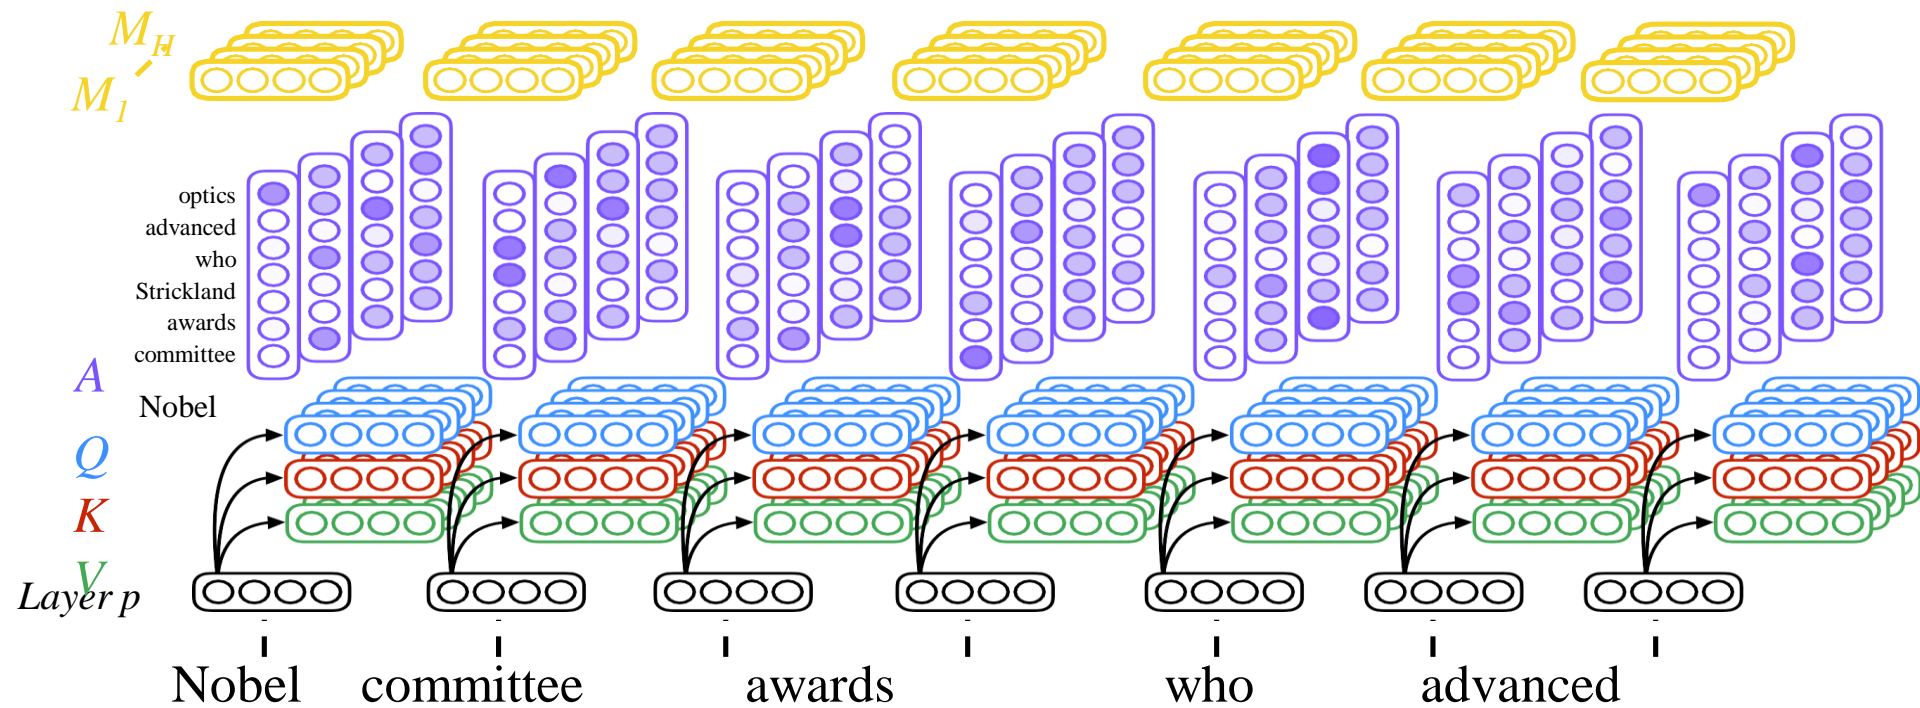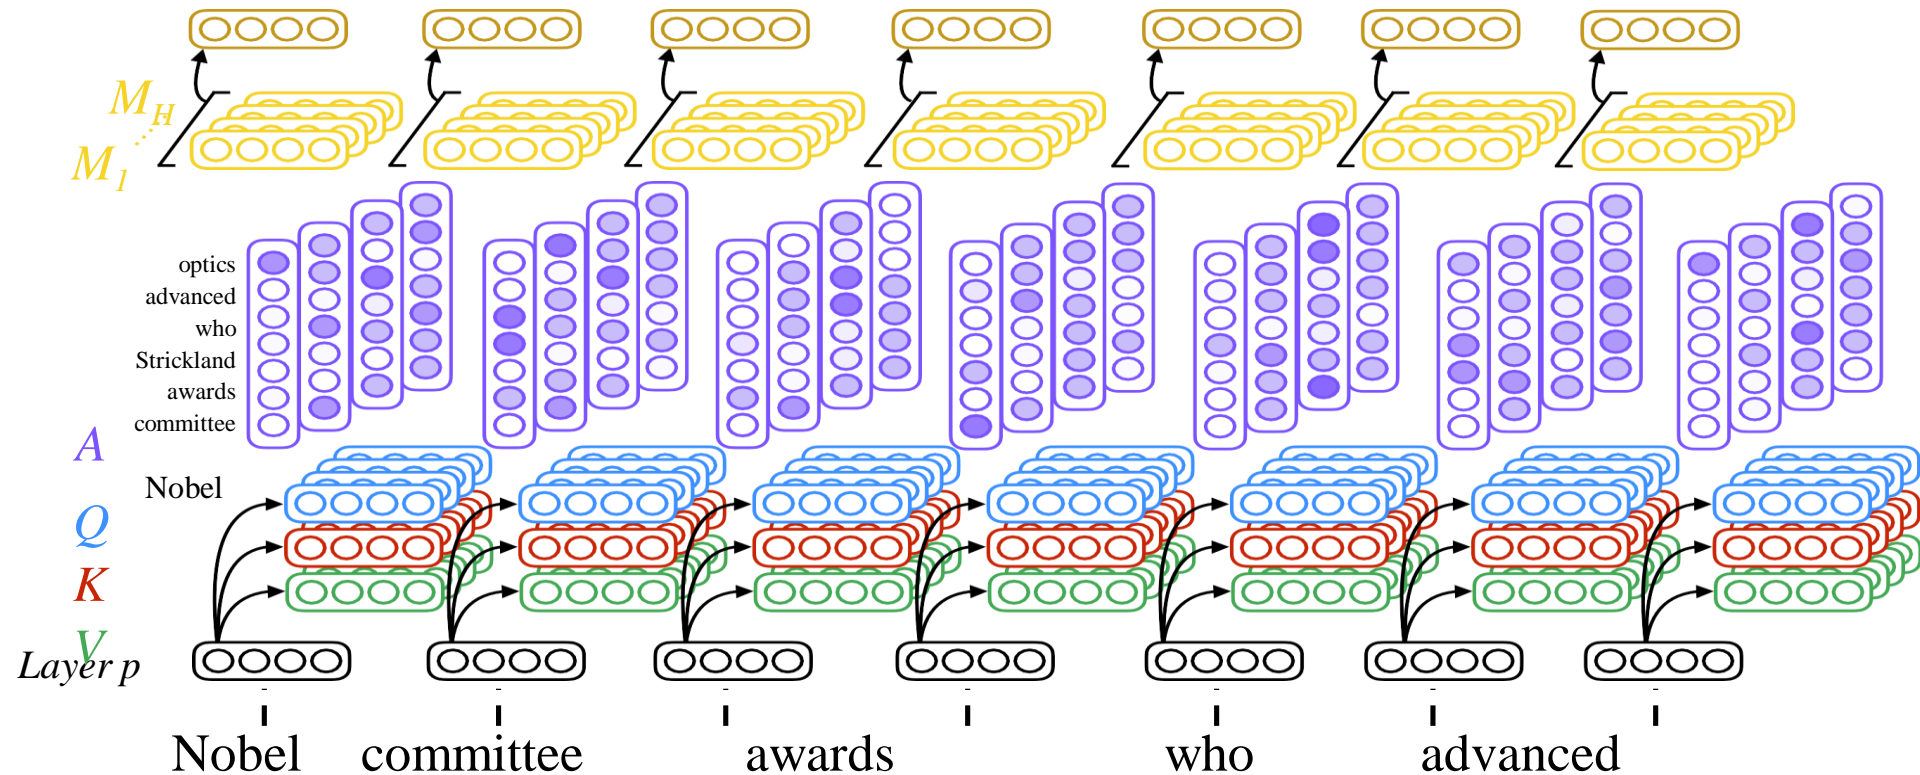
M_H
M_1

optics
advanced
who
Strickland
awards
committee

A

Nobel

Q
K
V

Layer p

Nobel   committee   awards   who   advanced

Multi-head self-attention

Multi-head self-attention

Multi-head self-attention

M_H
M_1

optics
advanced
who
Strickland
awards
committee
A Nobel

Q
K
V

Layer
p+1

Nobel   committee   awards   who   advanced

Multi-head self-attention

Layer J | **Multi-head self-attention + feed forward** |

Layer p | **Multi-head self-attention + feed forward** |

Layer 1 | **Multi-head self-attention + feed forward** |

Nobel   committee   awards   Who   advanced
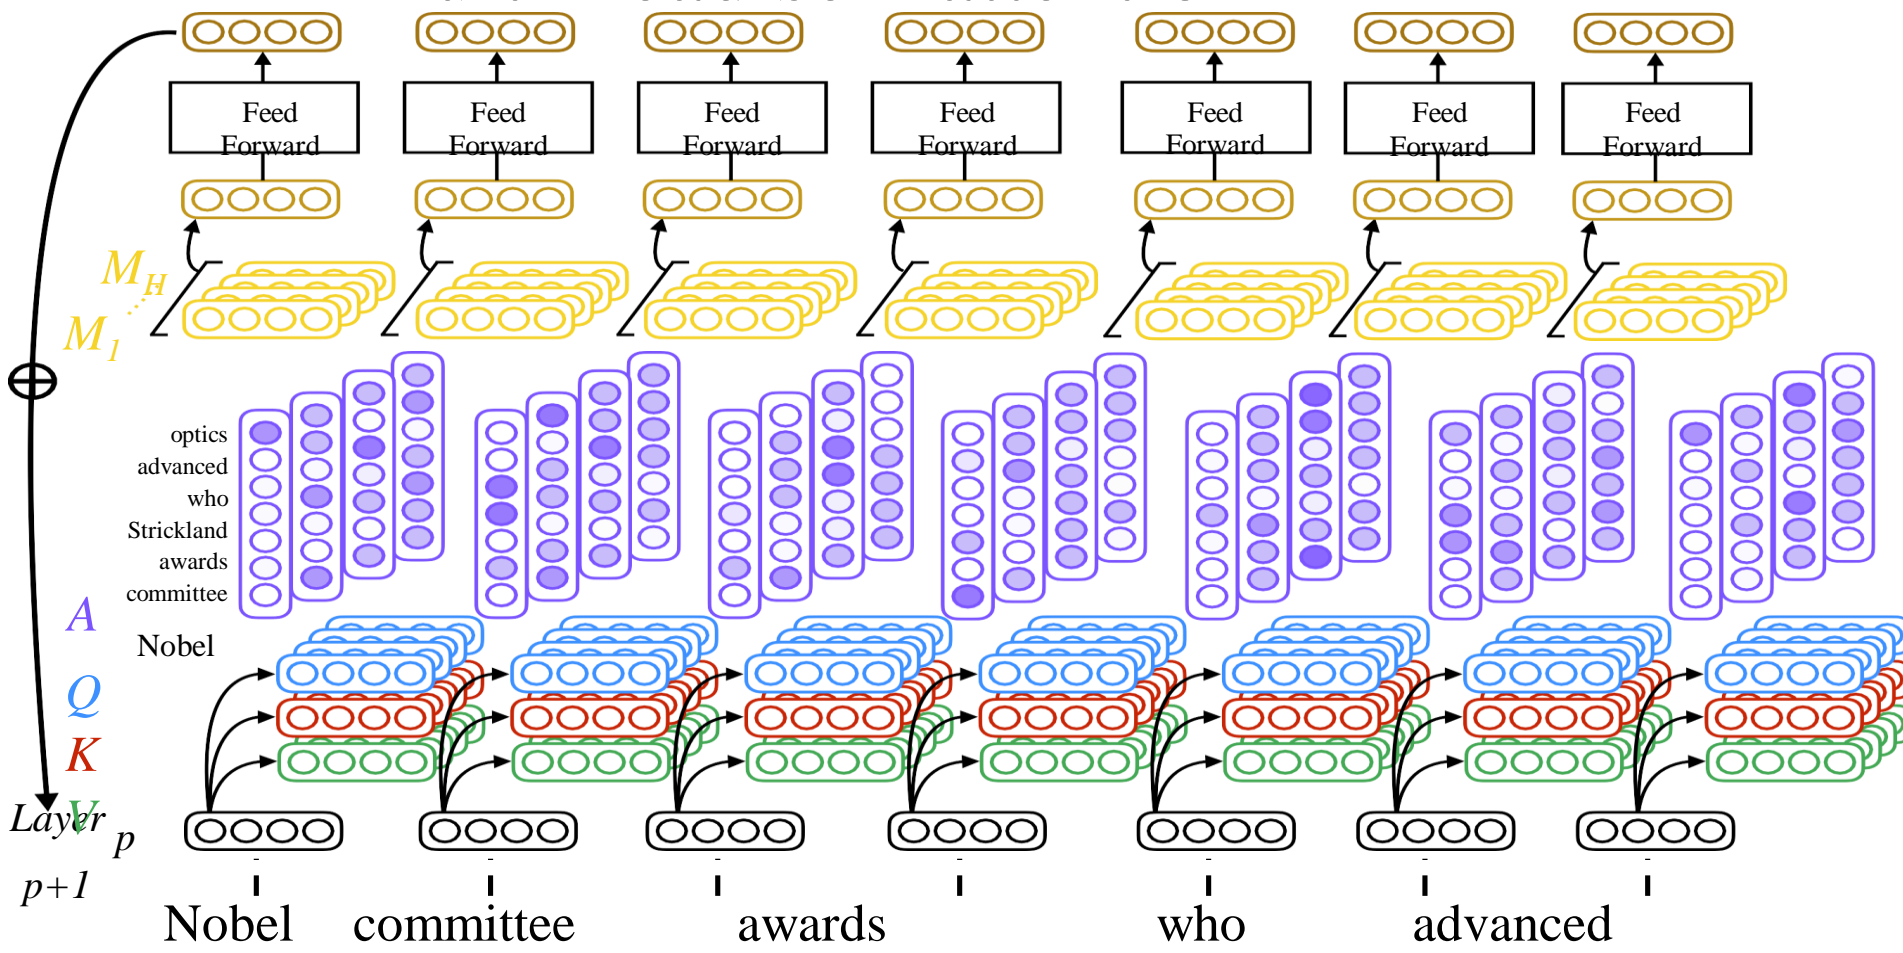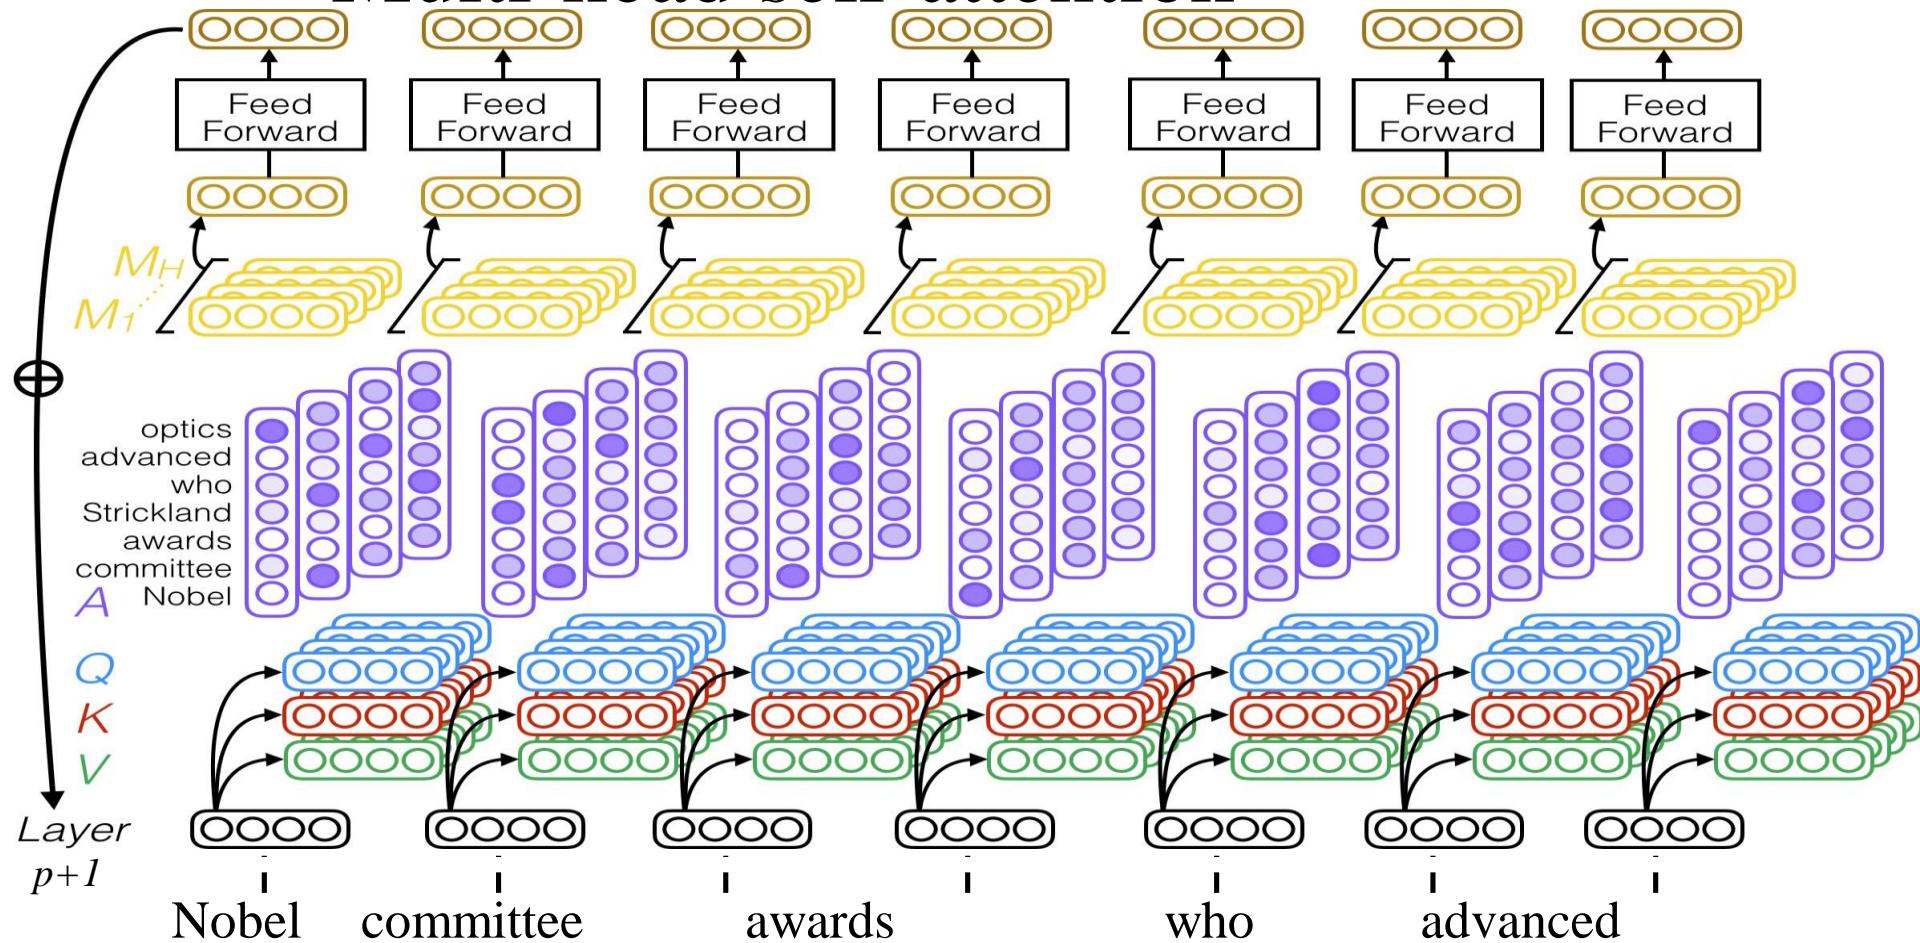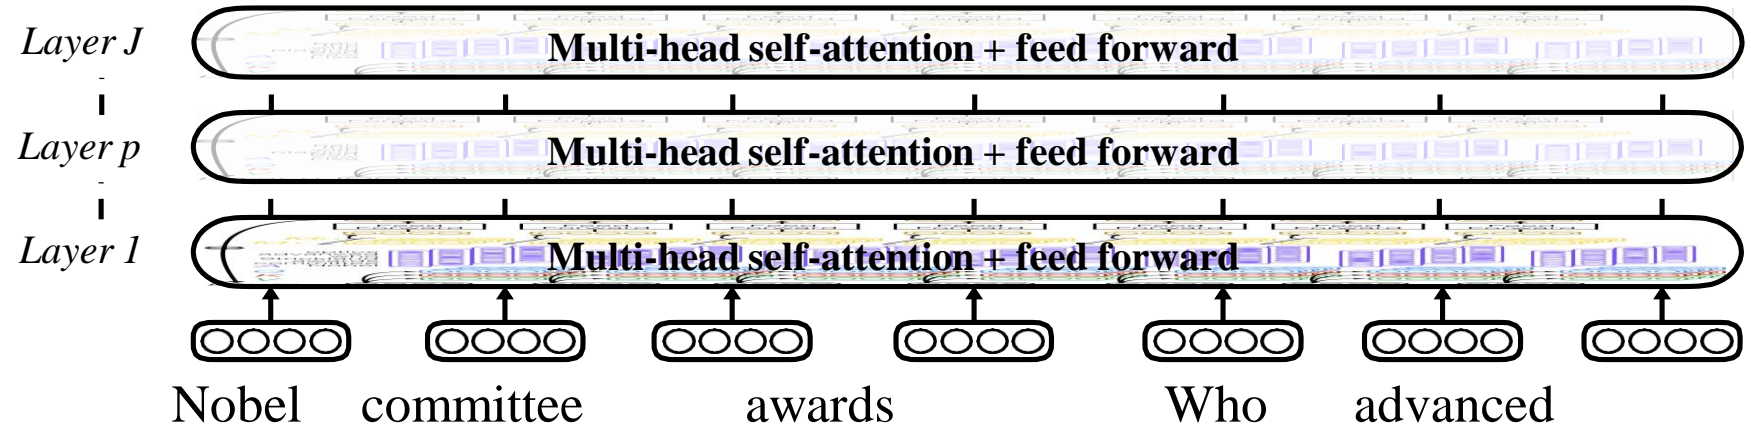
Position embeddings are *added* to each word embedding. Otherwise, since we have no recurrence, our model is unaware of the position of a word in the sequence!

Output Probabilities

Softmax

Add & Norm
Feed Forward

Add & Norm
Multi-Head Attention

Add & Norm
Masked Multi-Head Attention

Add & Norm
Multi-Head Attention

N×

N×

Positional Encoding

Positional Encoding

Input Embedding

Output Embedding

Inputs

Outputs (shifted right)
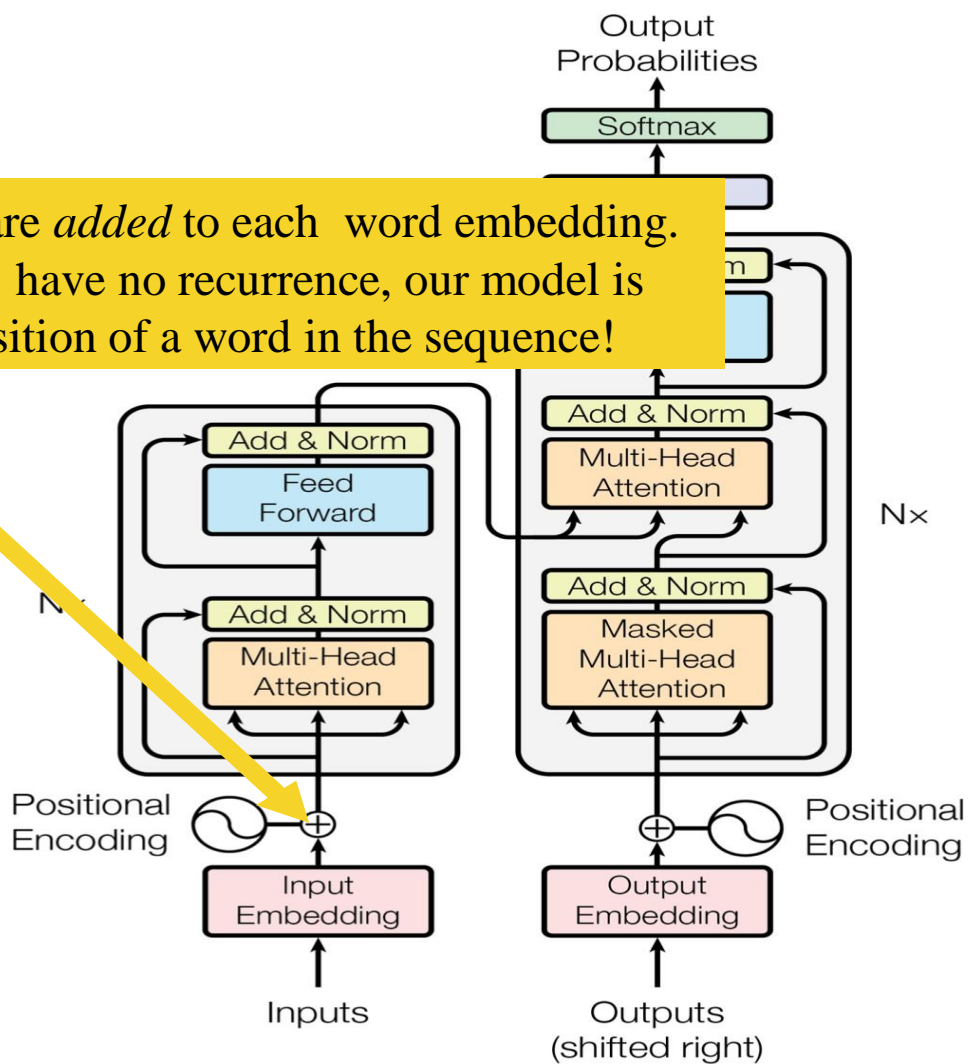
Residual connections, which mean that we add the input to a particular block to its output, help improve gradient flow
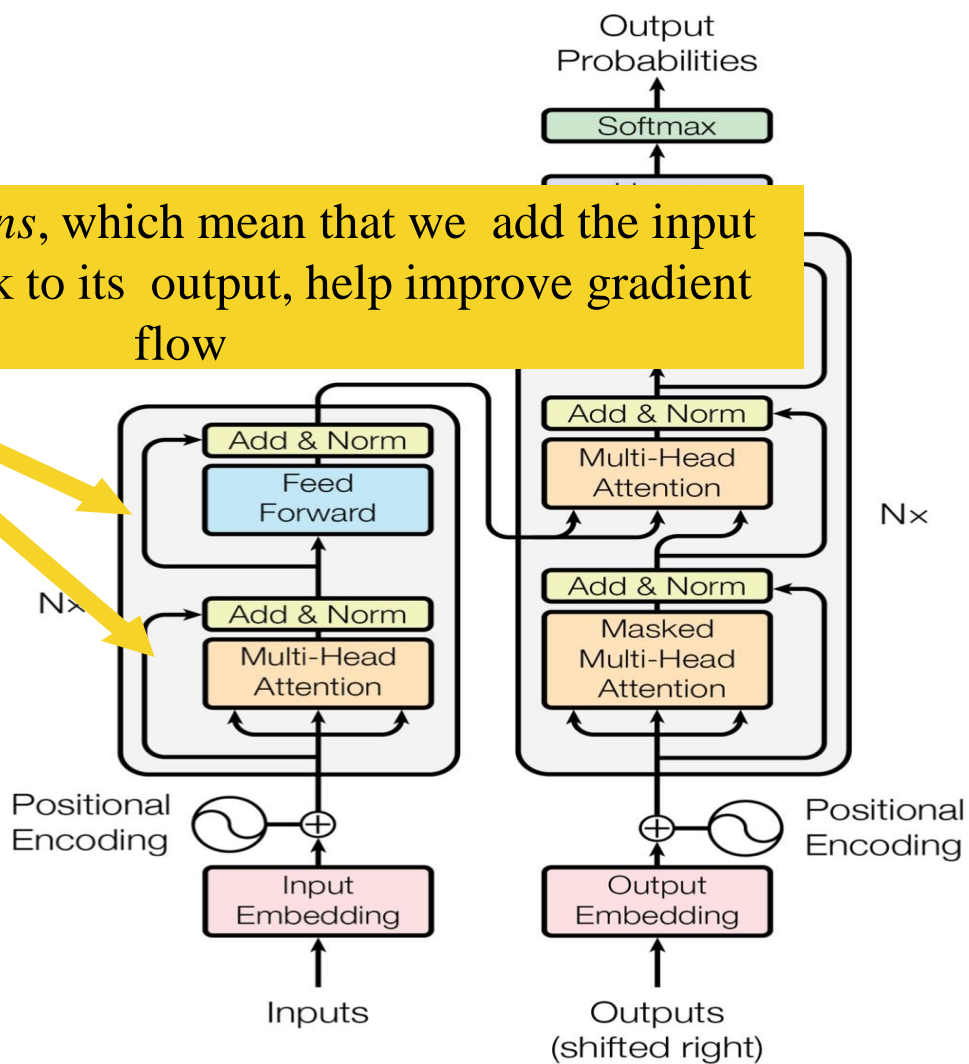
A feed-forward layer on top of the attention- weighted averaged value vectors allows us to add more parameters / nonlinearity
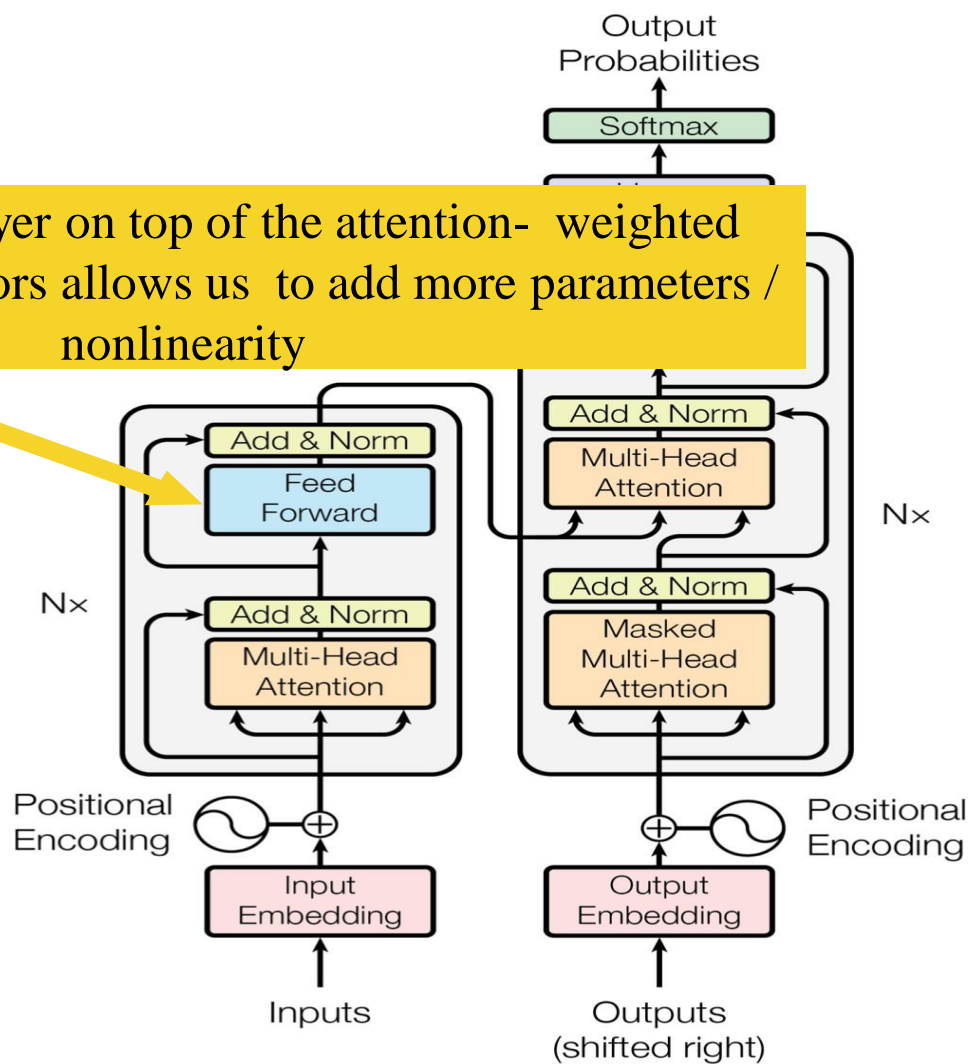
We stack as many of these *Transformer* blocks on top of each other as we can (bigger models are generally better given enough data!)
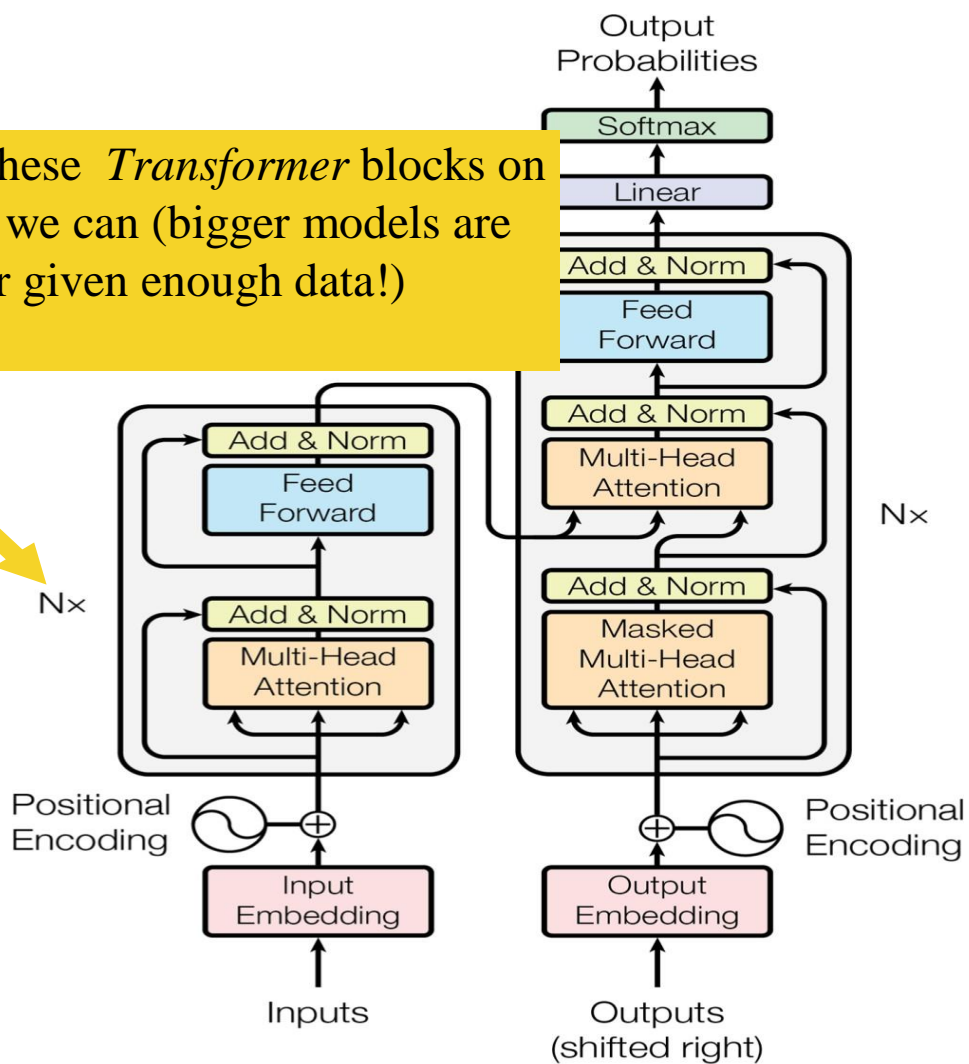
Moving onto the decoder, which takes in English sequences that have been shifted to the right (e.g., *<START> schools opened their*)
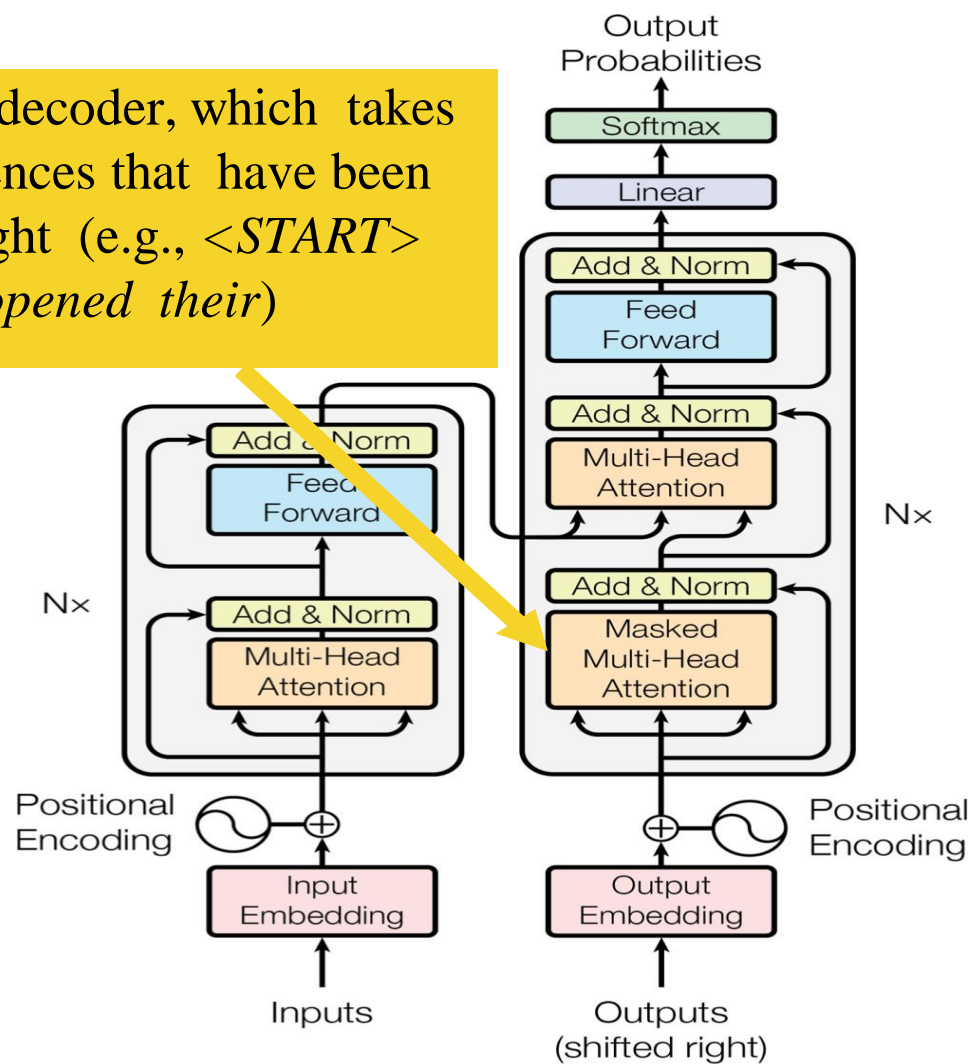
Output Probabilities

Softmax

Linear

Add & Norm

Feed Forward

Add & Norm

Multi-Head Attention

Add & Norm

Masked Multi-Head Attention

Nx

Add & Norm

Feed Forward

Add & Norm

Multi-Head Attention

Nx

Positional Encoding

Input Embedding

Inputs

Positional Encoding

Output Embedding

Outputs (shifted right)

We first have an instance of *masked self attention*. Since the decoder is responsible for predicting the English words, we need to apply masking as we saw before.
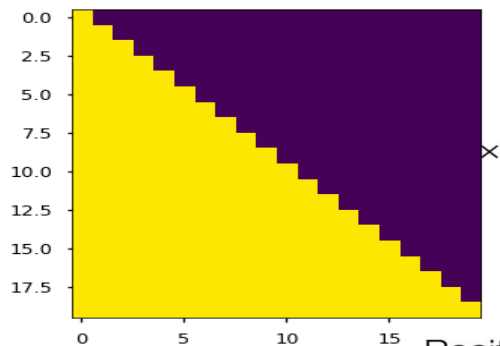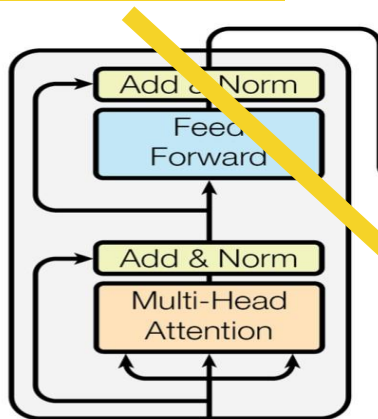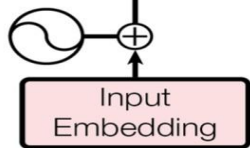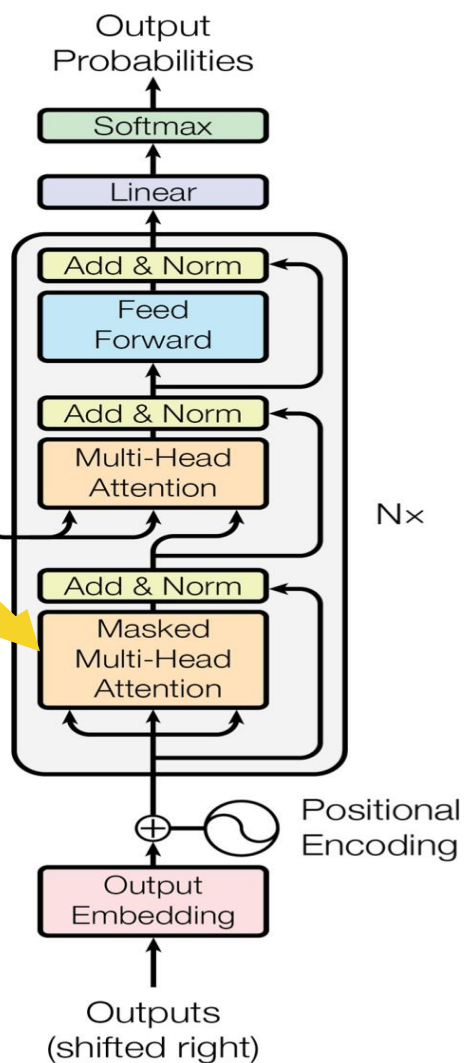
We first have an instance of *masked self attention*. Since the decoder is responsible for predicting the English words, we need to apply masking as we saw before.

Why don't we do masked self-attention in the encoder?

Output Probabilities

Softmax

Linear

Add & Norm

Feed Forward

Add & Norm

Multi-Head Attention

Add & Norm

Masked Multi-Head Attention

Add & Norm

Feed Forward

Add & Norm

Multi-Head Attention

Nx

Positional Encoding

Positional Encoding

Input Embedding

Output Embedding

Inputs

Outputs (shifted right)
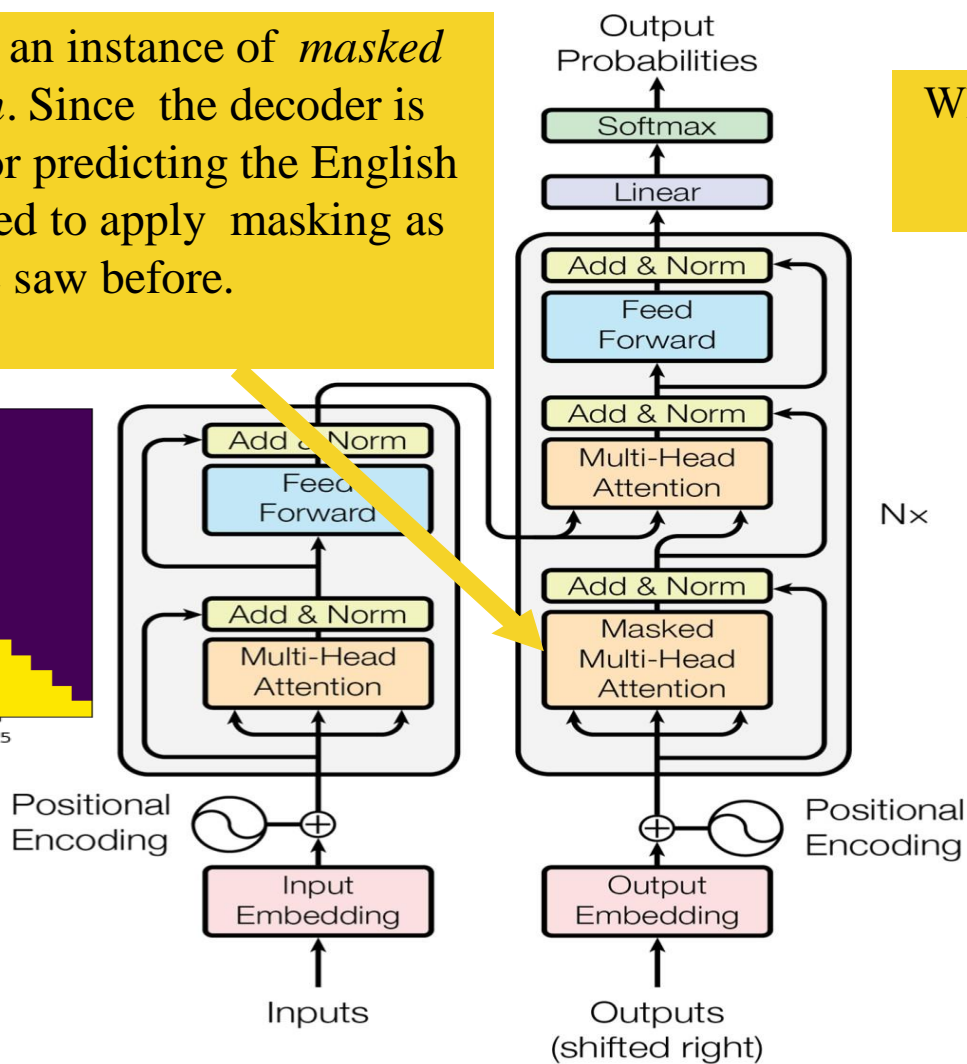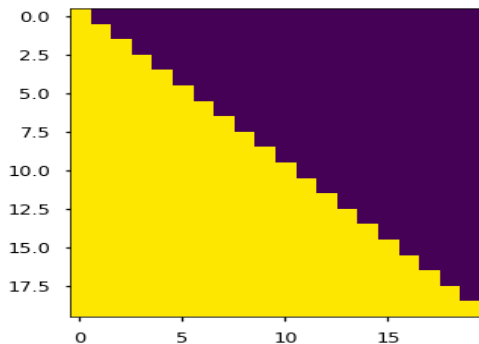
Now, we have *cross attention*, which connects the decoder to the encoder by enabling it to attend over the encoder's final hidden states.
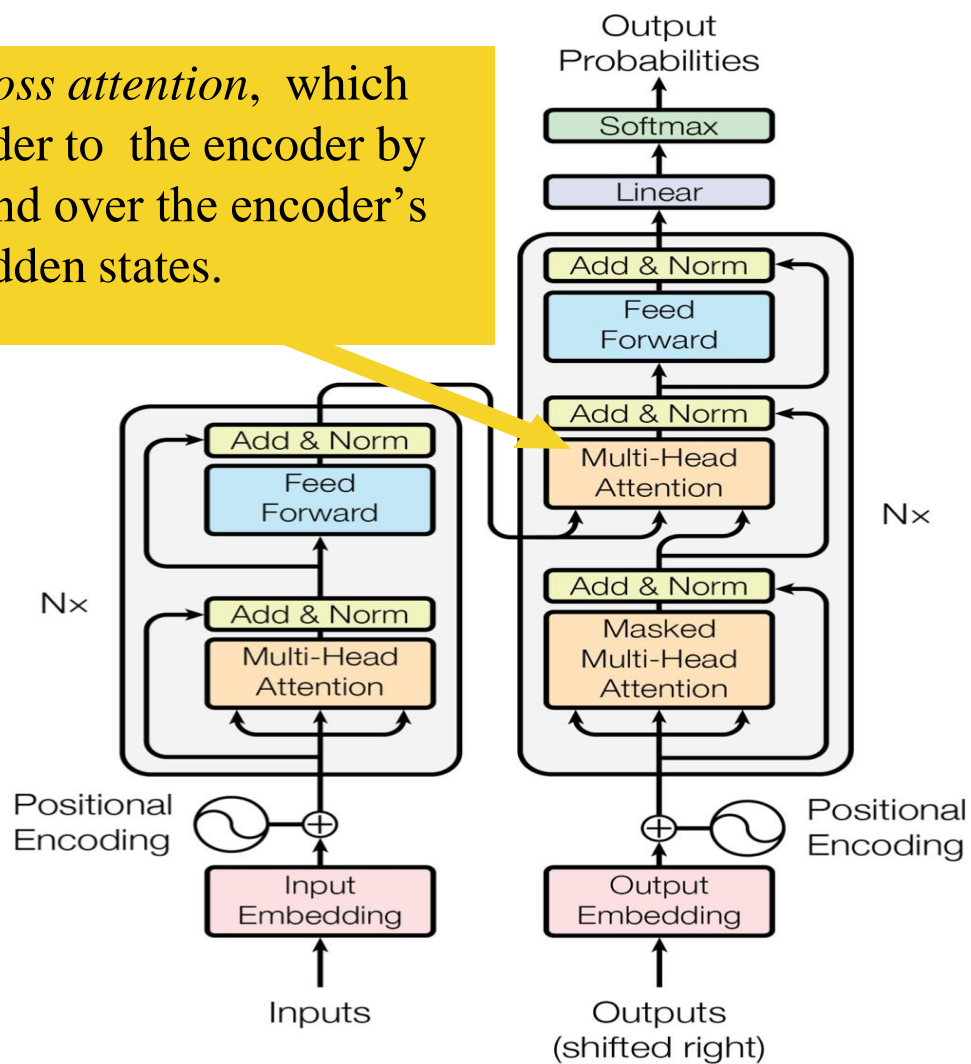
After stacking a bunch of these decoder blocks, we finally have our familiar Softmax layer to predict the next English word

Output Probabilities

Softmax

Linear

Add & Norm

Feed Forward

Add & Norm

Multi-Head Attention

Add & Norm

Masked Multi-Head Attention

Add & Norm

Feed Forward

Add & Norm

Multi-Head Attention

N×

N×

Positional Encoding

Positional Encoding

Input Embedding

Output Embedding

Inputs

Outputs (shifted right)
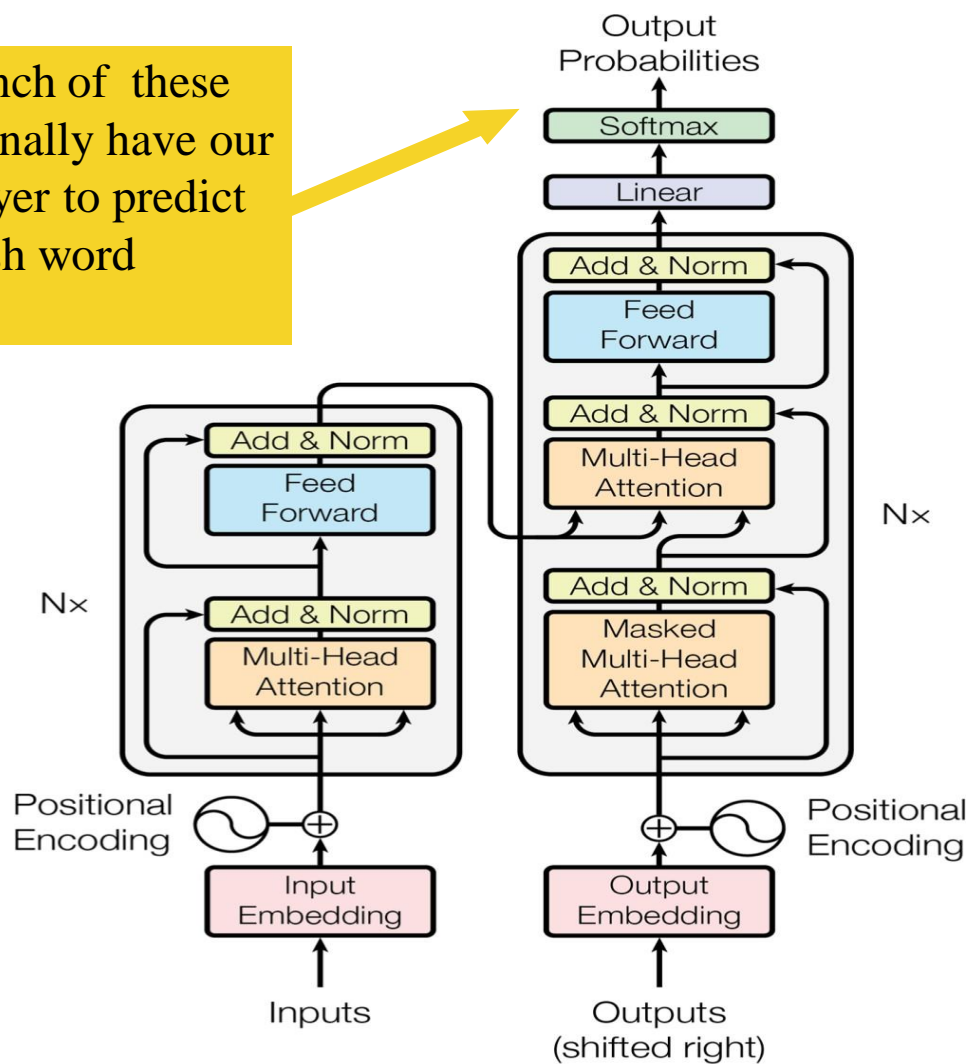
Positional encoding
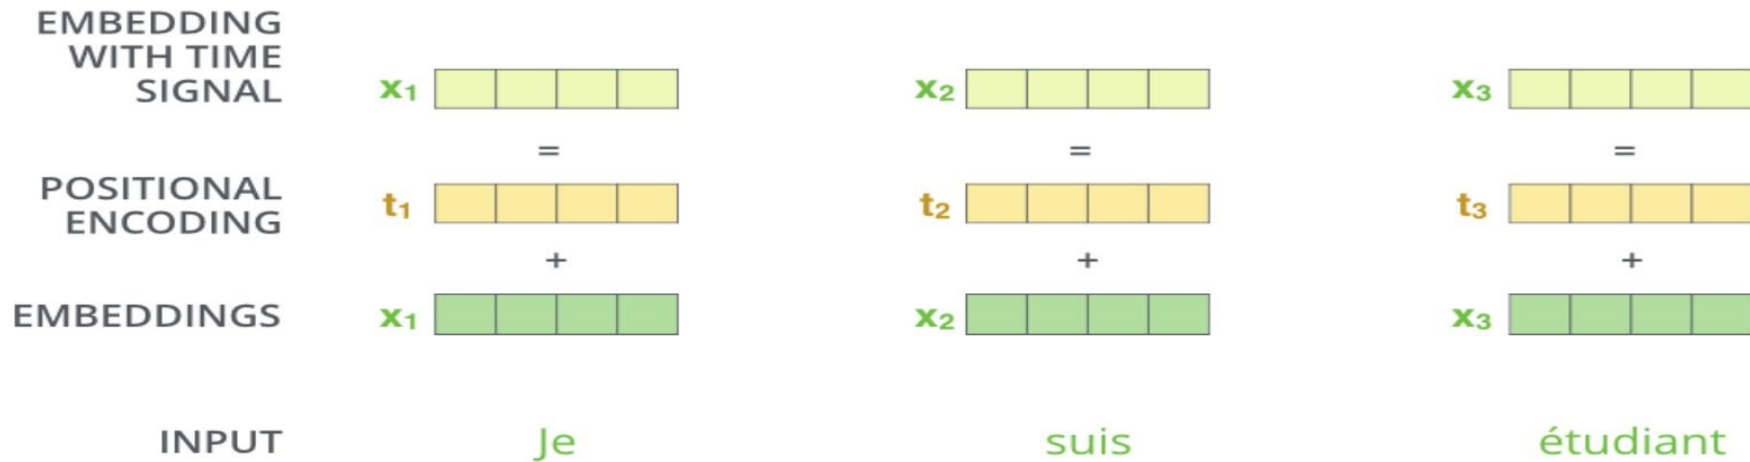
Intuitive example

0 :	0 0 0 0				8 :	1 0 0 0		
1 :	0 0 0 1				9 :	1 0 0 1		
2 :	0 0 1 0				10 :	1 0 1 0		
3 :	0 0 1 1				11 :	1 0 1 1		
4 :	0 1 0 0				12 :	1 1 0 0		
5 :	0 1 0 1				13 :	1 1 0 1		
6 :	0 1 1 0				14 :	1 1 1 0		
7 :	0 1 1 1				15 :	1 1 1 1		

Transformer positional encoding

$$PE_{(pos,2i)} = \sin\left(\frac{pos}{10000^{2i/d_{model}}}\right)$$

$$PE_{(pos,2i+1)} = \cos\left(\frac{pos}{10000^{2i/d_{model}}}\right)$$

Positional encoding is a 512d vector
i = a particular dimension of this vector
pos = dimension of the word
d_model = 512

What does this look like?

(each row is the pos. emb. of a 50-word sentence)
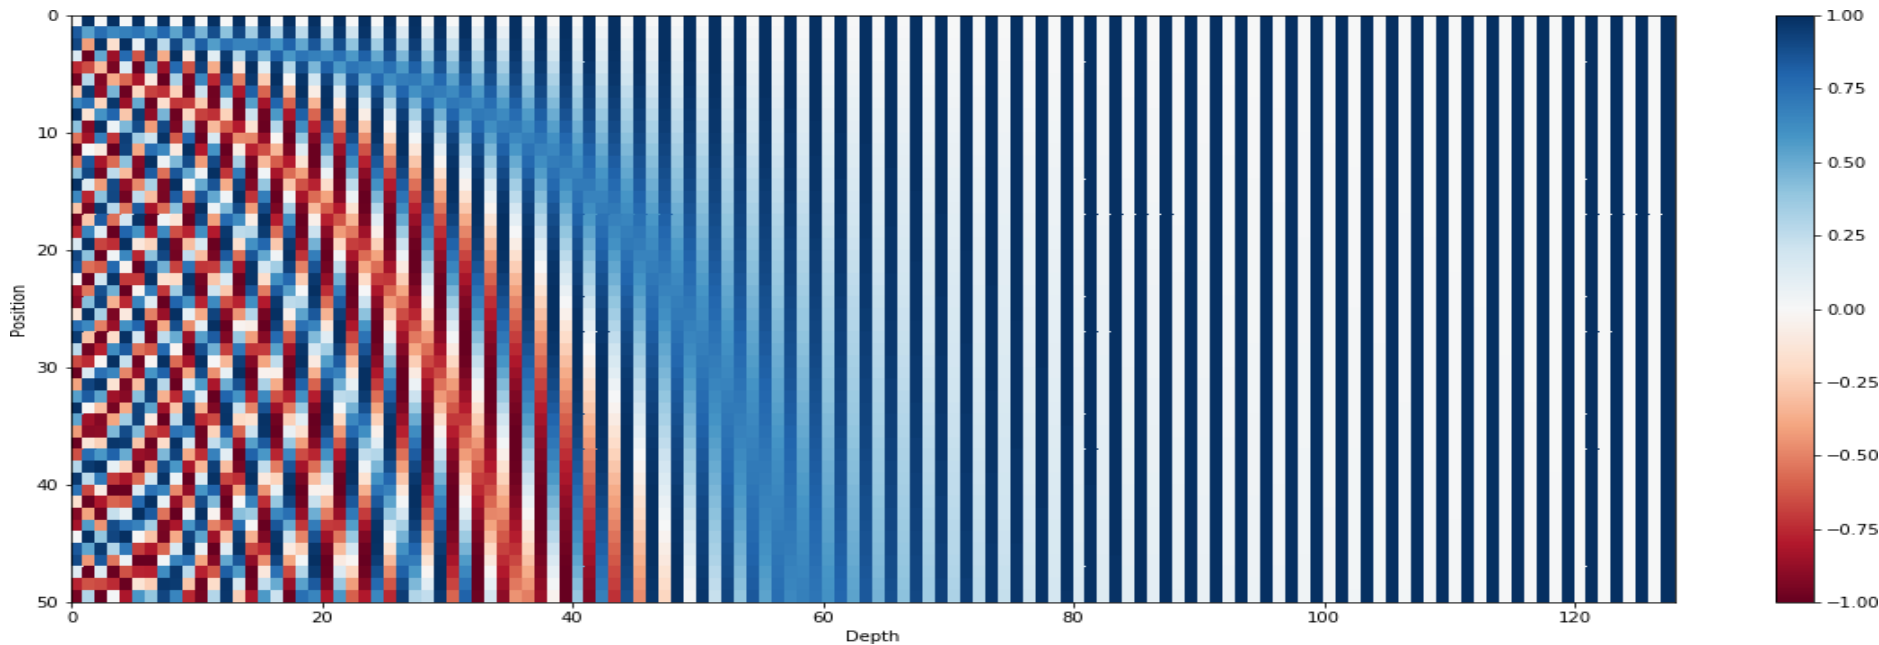
More on new-Transformer

What would we like to fix about the Transformer?

Quadratic compute in self-attention (today):

- Computing all pairs of interactions means our computation grows quadratically with the sequence length!

- For recurrent models, it only grew linearly!
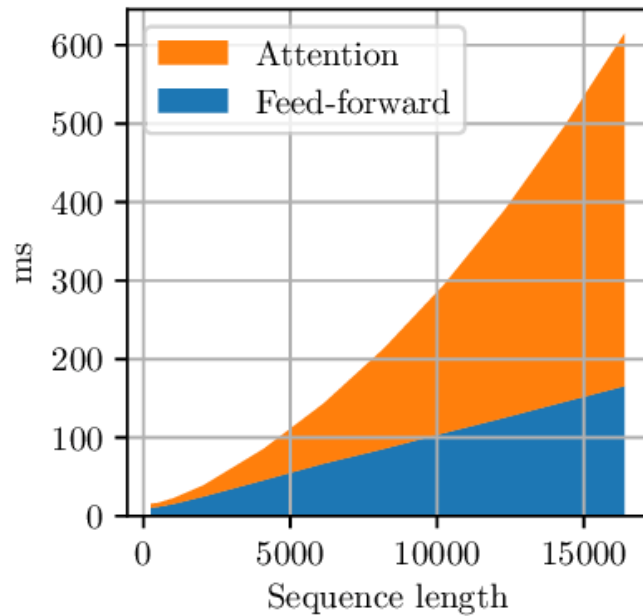
Quadratic computation as a function of sequence length

- One of the benefits of self-attention over recurrence was that it's highly parallelizable.

- However, its total number of operations grows as $O(n^2 d)$, where n is the sequence length, and d is the dimensionality.

$$XQ \quad K^\top X^\top \quad = \quad XQK^\top X^\top \quad \in \mathbb{R}^{n \times n}$$

Need to compute all pairs of interactions!
$O(n^2 d)$

- Think of d as around **1,000** (though for large language models it's much larger!).
  - So, for a single (shortish) sentence, $n \leq 30; n^2 \leq$ **900.**
  - In practice, we set a bound like $n = 512$.
  - **But what if we'd like $n \geq 50,000$?** For example, to work on long documents?

Work on improving on quadratic self-attention cost

Considerable recent work has gone into the question, *Can we build models like Transformers without paying the all-pairs self-attention cost?*
For example, **Linformer** [Wang et al., 2020]

Key Idea:
- Linformer introduces a novel concept called "compressed" or "linearized" self-attention.
- Instead of computing attention scores for all pairs of input elements, it employs linear projections to reduce the complexity.

Sequence length / batch size
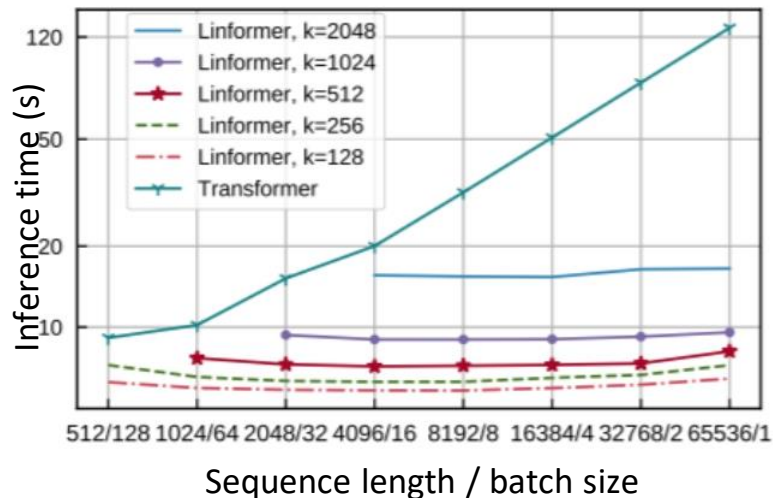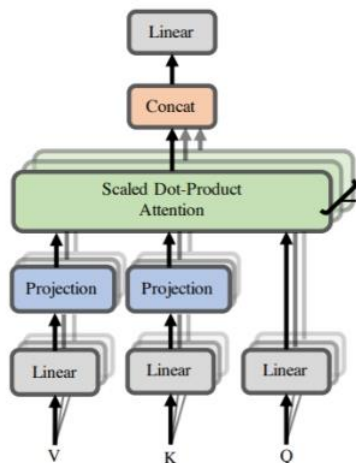
Example: Longformer / Big Bird

Key idea: use sparse attention patterns!
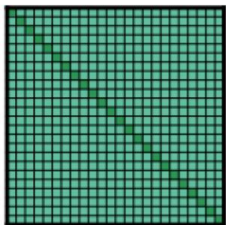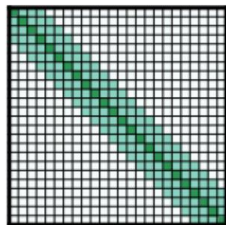
(a) Full n^2 attention   (b) Sliding window attention   (c) Dilated sliding window   (d) Global+sliding window
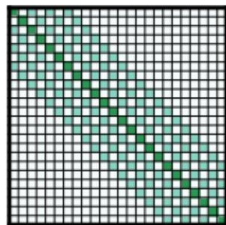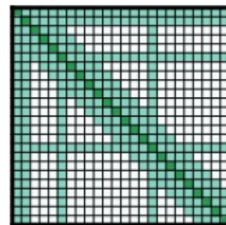
(Beltagy et al., 2020): Longformer: The Long-Document Transformer
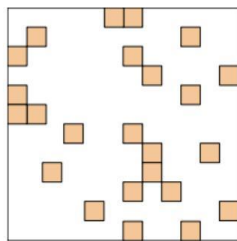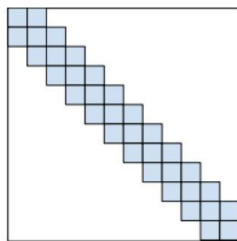
(a) Random attention   (b) Window attention   (c) Global Attention   (d) BIGBIRD
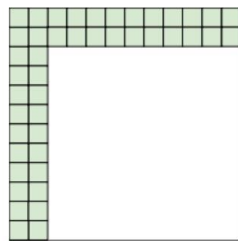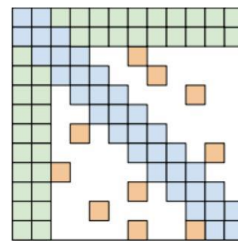
(Zaheer et al., 2021): Big Bird: Transformers for Longer

Do we even need to remove the quadratic cost of attention?

- **As Transformers Scale Up:** When Transformers are scaled to larger sizes, an increasingly significant portion of computational resources is allocated to tasks outside of the self-attention mechanism, despite its quadratic computational cost.

- **Current Practice:** In practice, nearly all large Transformer-based language models continue to rely on the traditional quadratic-cost attention mechanism that has been presented.

- **Challenges with Cost-Efficiency:** Alternative, more computationally efficient methods often do not perform as effectively when applied at a large scale.

- **Exploring Cheaper Alternatives:** Is there value in exploring cost-efficient alternatives to self-attention, or could we unlock the potential for significantly improved models with much longer contextual information (e.g., >100k tokens) if we find the right approach?

Do Transformer Modifications Transfer?

- "Surprisingly, we find that most modifications do not meaningfully improve performance."

Do Transformer Modifications Transfer Across Implementations and Applications?

Sharan Narang* Hyung Won Chung Yi Tay William Fedus

Thibault Fevry[†] Michael Matena[†] Karishma Malkan[†] Noah Fiedel

Noam Shazeer Zhenzhong Lan[†] Yanqi Zhou Wei Li

Nan Ding Jake Marcus Adam Roberts Colin Raffel[†]

Vision Transformer (ViT)
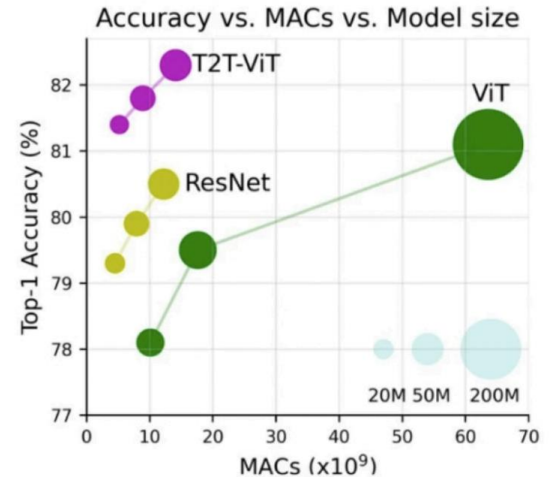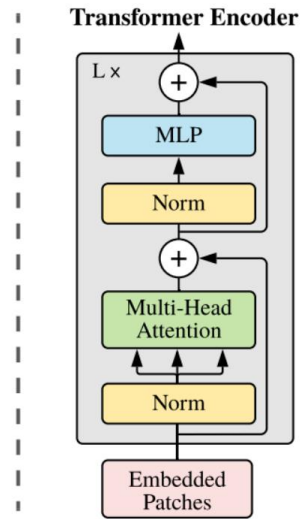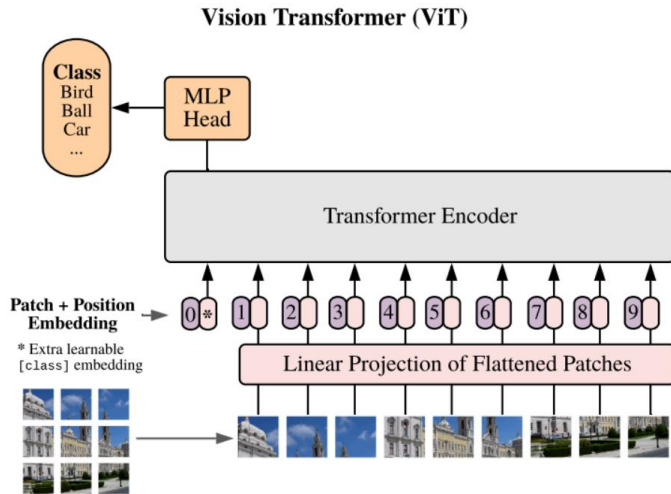
(Dosovitskiy et al., 2021): An Image is Worth 16x16 Words: Transformers for Image Recognition at Scale

Music Transformer

https://magenta.tensorflow.org/music-transformer
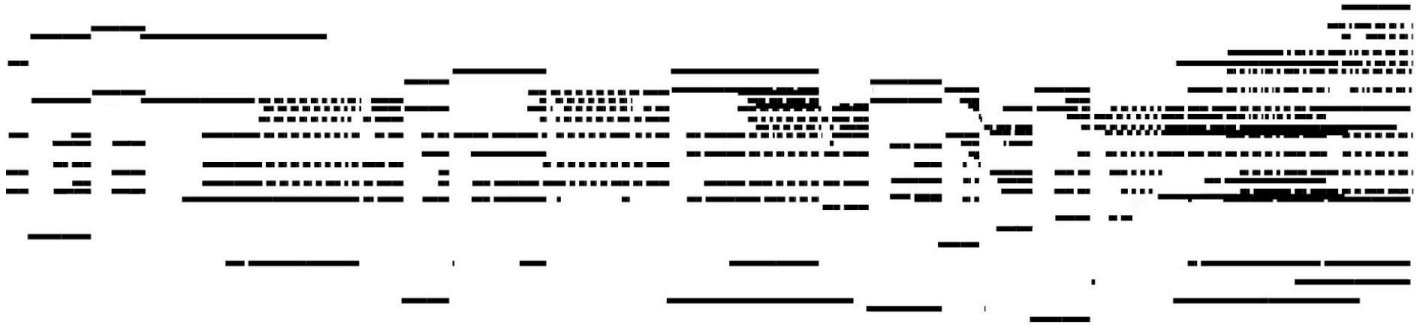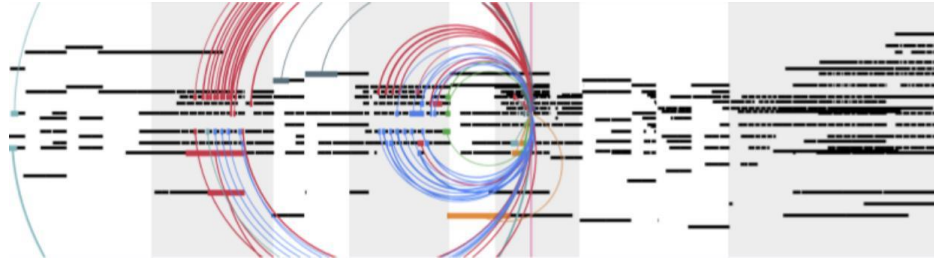
(Huang et al., 2018): Music Transformer: Generating Music with Long-Term Structure

Why transformer

Why Pretraining + Transformers

- 1.Because transformers are more efficient?

Transformers are shower comparing to LSTM with same amount parameters

Why Pretraining + Transformers

- 1.Because transformers are more efficient?

Transformers are shower comparing to LSTM with same amount parameters

- 2. Because transformers are better on machine translation?

RNNs and CNNs are equally good in machine translations

Why Pretraining + Transformers

- 1.Because transformers are more efficient?

Transformers are shower comparing to LSTM with same amount parameters

- 2. Because transformers are better on machine translation?

RNNs and CNNs are equally good in machine translations

- 3. Because transformers use nothing but attention?

So what?

Why Pretraining + Transformers

- 1.Because transformers are more efficient?

Transformers are shower comparing to LSTM with same amount parameters

- 2. Because transformers are better on machine translation?

RNNs and CNNs are equally good in machine translations

- 3. Because transformers use nothing but attention?

So what?

- 4. Because transformers learns contextualised word embeddings?

RNN also can learn contextualised word embeddings

Why Pretraining + Transformers

❖ **Capacity**: The model has sufficient expressive capabilities
❖ **Optimization**: Can optimize and obtain better solutions in a huge expression space
❖ **Generalization**: Better solutions can generalize on test data

"Exploring the Limits of Language Modeling
Jozefowicz et al 2016

LSTM-8192-1024, 1.8 billion params, ppl 30.6
LSTM-8192-2048, 3.3 billion params, ppl 32.2

Dai, Yang et al 2016
Transformer-XL Base, 0.46 billion params, ppl 23.5
Transformer-XL Large, 0.8 billion params, ppl 21.8

ppl=perplexity, the lower the better

Scalability: Transformers scale much better with more parameters

Deep understanding of transformer

What if

- ✓ removing SAN
- ✓ removing FFN
- ✓ removing PE
- ✓ and many others?

Without FFN, pure SAN
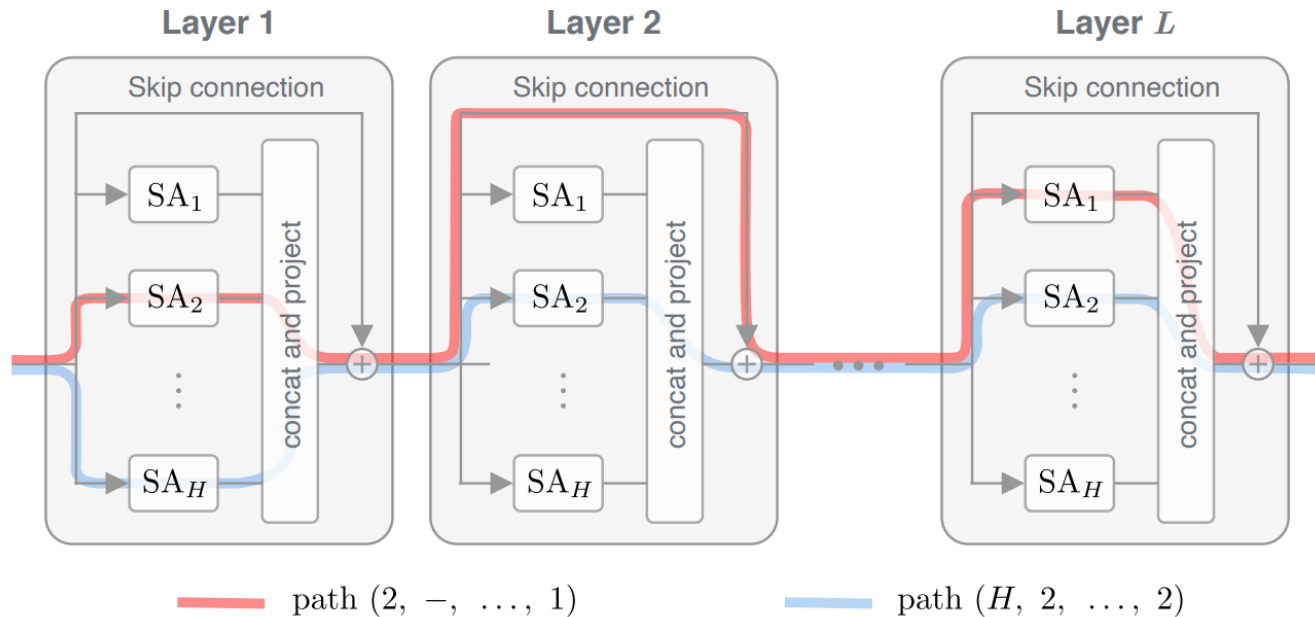
Y Dong, JB Cordonnier, A Loukas. Attention is not all you need: Pure attention loses rank doubly exponentially with depth. https://browse.arxiv.org/pdf/2103.03404.pdf

Without SAN, pure FNN
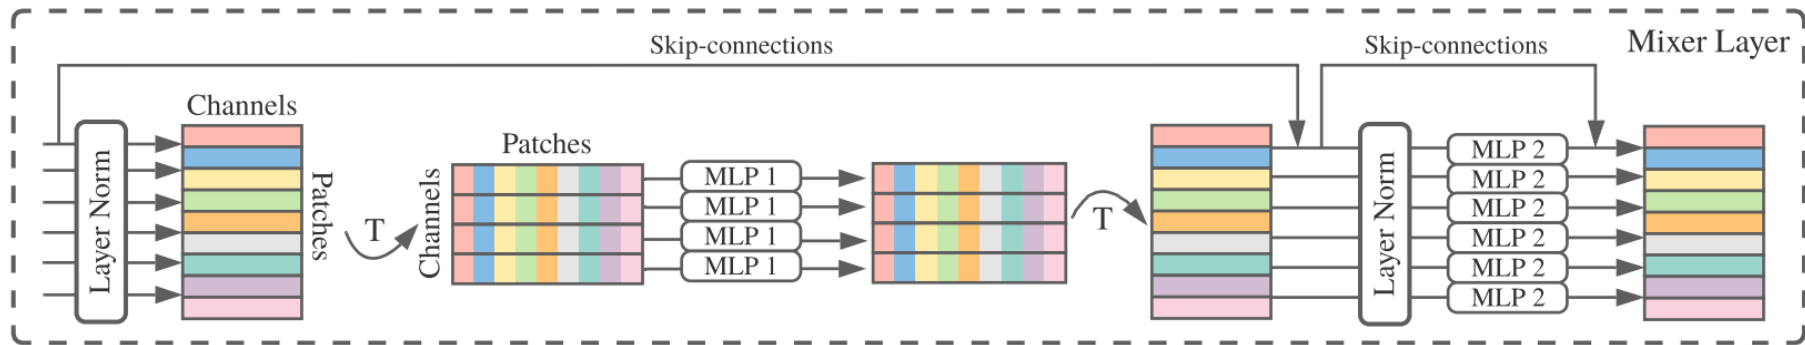
At least it works for computer vision.

Replace SAN with fourier
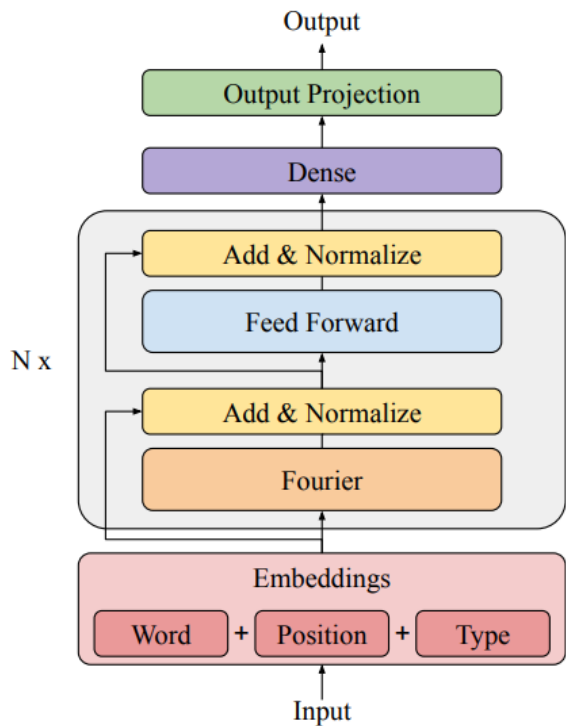
- ❖ Highlight the potential of linear units as a drop-in replacement for the attention mechanism in text classification tasks.
- ❖ FNet will be effective as a lightweight

James Lee-Thorp, Joshua Ainslie, Ilya Eckstein, Santiago Ontanon .
FNet: Mixing Tokens with Fourier Transforms. NAACL 2022

How to place FFN and SAN

(a) Interleaved Transformer

(b) Sandwich Transformer

Figure 1: A transformer model (a) is composed of inter-
leaved self-attention (green) and feedforward (purple)
sublayers. Our sandwich transformer (b), a reordering
of the transformer sublayers, performs better on lan-
guage modeling. Input flows from left to right.

Model	PPL
	20.74
	20.64
	20.33
	20.27
	19.98
	19.92
	19.69
	19.54
	19.13
	19.08
	18.90
	18.83
	18.83
	18.77
	18.68
	18.64
	18.61
	18.60
	18.55
	18.54
	18.49
	18.38
	18.28
	18.25
	18.19

Ofir Press, Noah A. Smith, Omer Levy. Improving Transformer Models by Reordering their Sublayers.
https://browse.arxiv.org/pdf/1911.03864.pdf

What will happen if the position embedding model is removed?

Table 3: Experiments on GLUE. The evaluation metrics are following the official GLUE benchmark (Wang et al., 2018). The best performance of each task is bold.

| PEs | single sentence | | sentence pair | | | | | | | |
	CoLA acc	SST-2 acc	MNLI acc	MRPC F1	QNLI acc	QQP F1	RTE acc	STS-B spear. cor.	WNLI acc	mean ± std
BERT without PE	39.0	86.5	80.1	86.2	83.7	86.5	63.0	87.4	33.8	76.6 ± 0.41
fully learnable (BERT-style) APE	60.2	**93.0**	84.8	89.4	88.7	87.8	**65.1**	**88.6**	37.5	82.2 ± 0.30
fixed sin. APE	57.1	92.6	84.3	89.0	88.1	87.5	58.4	86.9	45.1	80.5 ± 0.71
learnable sin. APE	56.0	92.8	84.8	88.7	88.5	87.7	59.1	87.0	40.8	80.6 ± 0.29
fully-learnable RPE	58.9	92.6	84.9	**90.5**	88.9	**88.1**	60.8	88.6	50.4	81.7 ± 0.31
fixed sin. RPE	60.4	92.2	84.8	89.5	88.8	88.0	62.9	88.1	45.1	81.8 ± 0.53
learnable sin. RPE	60.3	92.6	**85.2**	90.3	**89.1**	**88.1**	63.5	88.3	49.9	82.2 ± 0.40
fully learnable APE + fully-learnable RPE	59.8	92.8	85.1	89.6	88.6	87.8	62.5	88.3	**51.5**	81.8 ± 0.17
fully learnable APE + fixed sin. RPE	59.2	92.4	84.8	89.9	88.8	87.9	61.0	88.3	48.2	81.5 ± 0.20
fully learnable APE+ learnable sin. RPE	**61.1**	92.8	**85.2**	**90.5**	89.5	87.9	**65.1**	88.2	49.6	**82.5 ± 0.44**
learnable sin. APE + fully-learnable RPE	57.2	92.7	84.8	88.9	88.5	87.8	58.6	88.0	51.3	80.8 ± 0.44
learnable sin. APE + fixed sin. RPE	57.6	92.6	84.5	88.8	88.6	87.6	63.1	87.4	48.7	81.3 ± 0.43
learnable sin. APE + learnable sin. RPE	57.7	92.7	85.0	89.6	88.7	87.8	62.3	87.5	50.1	81.4 ± 0.33

Benyou Wang, Lifeng Shang, Christina Lioma, Xin Jiang, Hao Yang, Qun Liu, Jakob Grue Simonsen. On Position Embeddings in BERT. https://openreview.net/pdf?id=onxoVA9FxMw
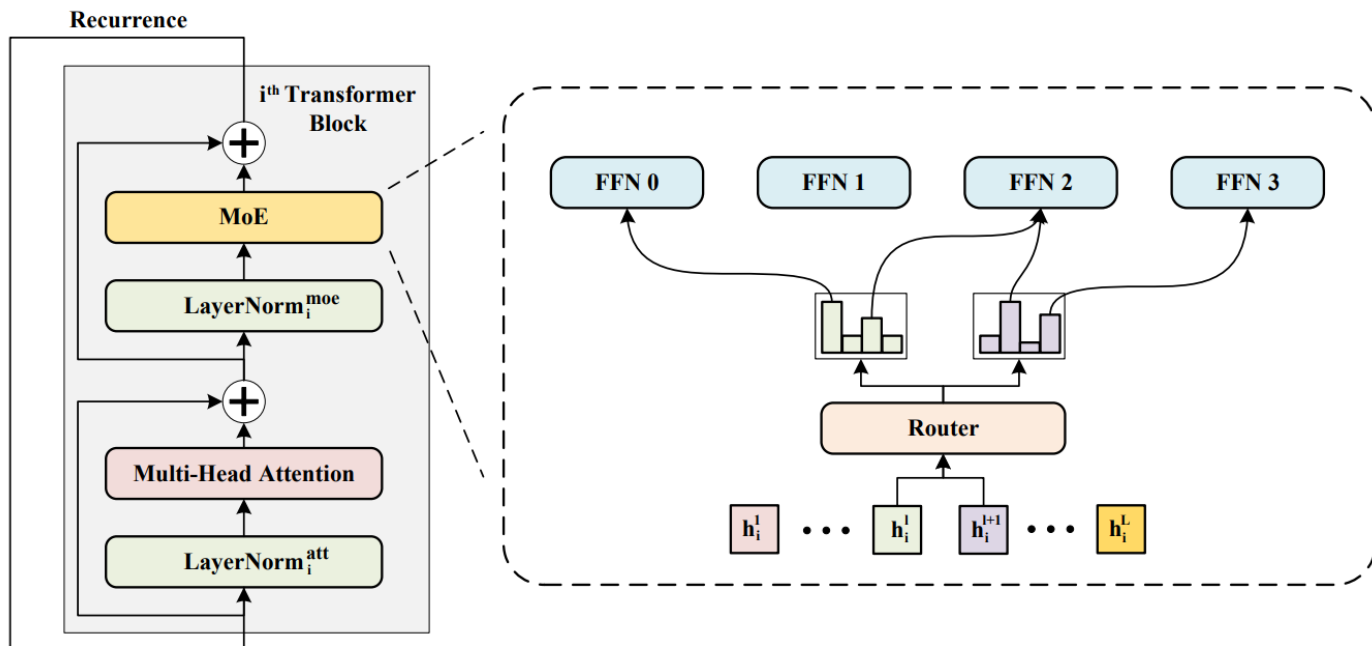
Improvements for Norm

[DeepNet](#) - 1000 layer Transformers

A new normalization function (DEEPNORM) is introduced
[replacing it is not Layer Norm! Instead, modify it similarly to:

layernorm (x + f(x)) ---> layernorm(x*alpha + f(x)).

The proposed method combines the advantages of both schools,
namely the good performance of Post-LN and the stable training of
Pre-LN, making DEEPNORM the preferred alternative.

Hongyu Wang, Shuming Ma, Li Dong, Shaohan Huang, Dongdong Zhang, Furu Wei. DeepNet: Scaling Transformers to 1,000 Layers. https://browse.arxiv.org/pdf/2203.00555.pdf

Is the model deeper or wider?

[Go Wider Instead of Deeper](#)

- ❖ WideNet first compresses trainable parameters along with depth by parameter-sharing across transformer blocks
- ❖ Each expert requires enough tokens to train.

Fuzhao Xue, Ziji Shi, Futao Wei, Yuxuan Lou, Yong Liu, Yang You. **Go Wider Instead of Deeper. https://arxiv.org/abs/2107.11817**

Scaling law?

Scaling Law for Neural Language Models

   Performance depends strongly on scale! We keep getting better performance as we scale the model, data, and compute up!
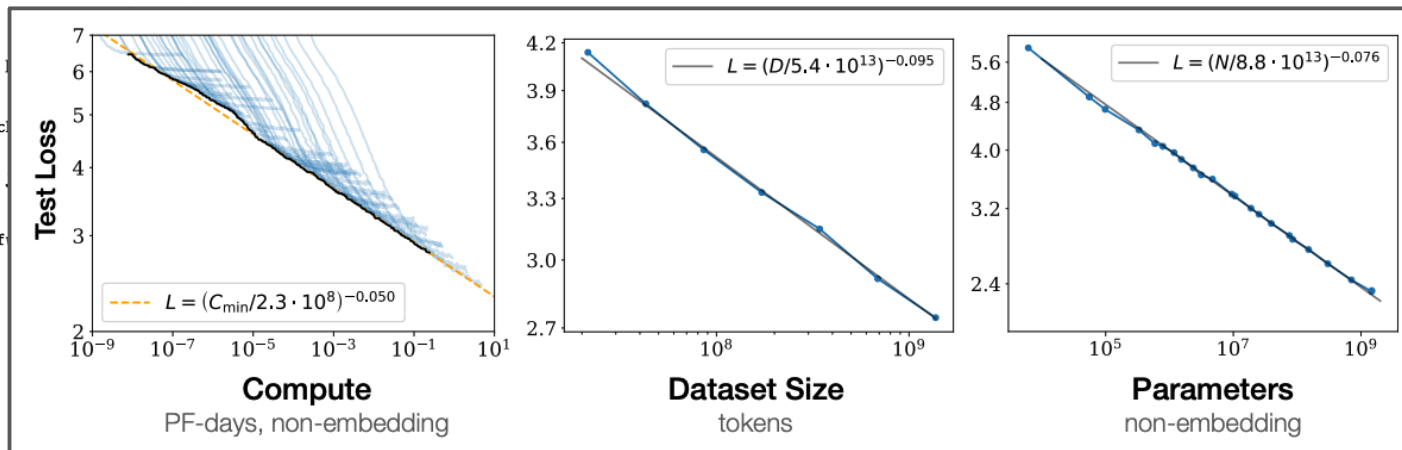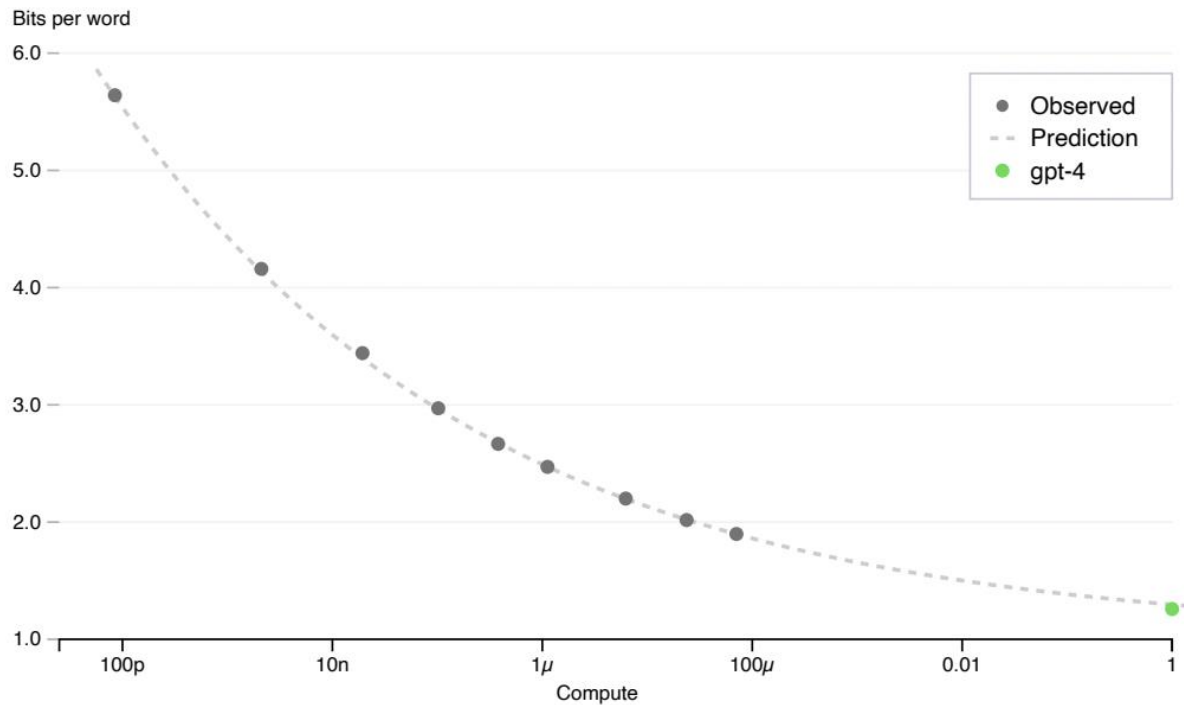

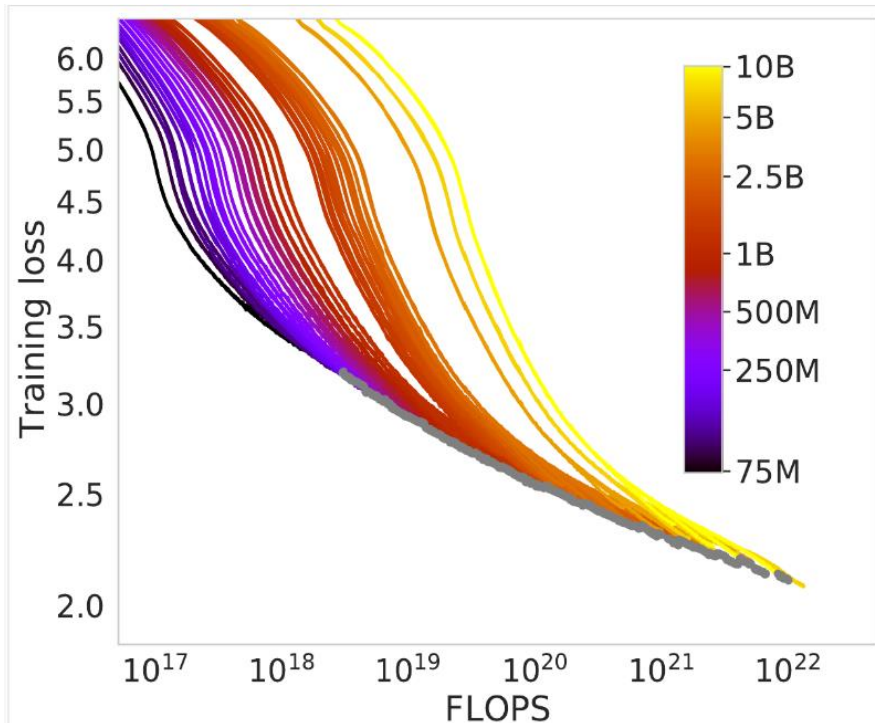Emergent abilities of large language models (TMLR '22).
J. Wei, Y. Tay, R. Bommasani, C. Raffel, B. Zoph, S. Borgeaud, D. Yogatama, M. Bosma, D. Zhou, D. Metzler, E. Chi, T. Hashimoto, O. Vinyals, P. Liang, J. Dean, & W. Fedus.

Scaling laws

OpenAI codebase next word prediction

GPT-4 Technical Report, *OpenAI (2023)*

Challenge to scaling law: Chinchilla's Death

Smaller models eventually reach the limit of their capacity for knowledge, and their learning slows, while that of a larger model, with a larger capacity, will overtake them and reach better performance past a given amount of training time.

While estimating how to get the best bang during training, OpenAI & DeepMind attempted to draw the Pareto frontier.

https://espadrine.github.io/blog/posts/chinchilla-s-death.html

Challenge to scaling law: Chinchilla's Death

## Can Chinchillas picture a Llama's sights?
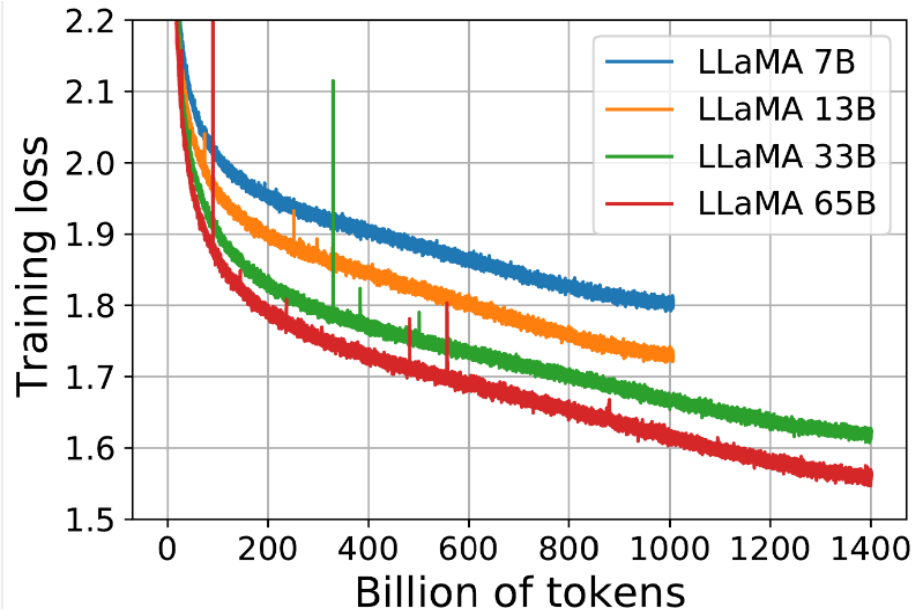
❖ Each curve first plummets in a **power law**,

❖ and then seemingly enters a **nearly-linear** decrease in loss (corresponding to a fairly constant rate of knowledge acquisition).

❖ At the very tip of the curve, they all break this line by **flattening** slightly.

❖ This should consider the cosine LR schedule.

https://espadrine.github.io/blog/posts/chinchilla-s-death.html

Challenge to scaling law: Chinchilla's Death
Can Chinchillas picture a Llama's sights?

Training speed (lower is better)
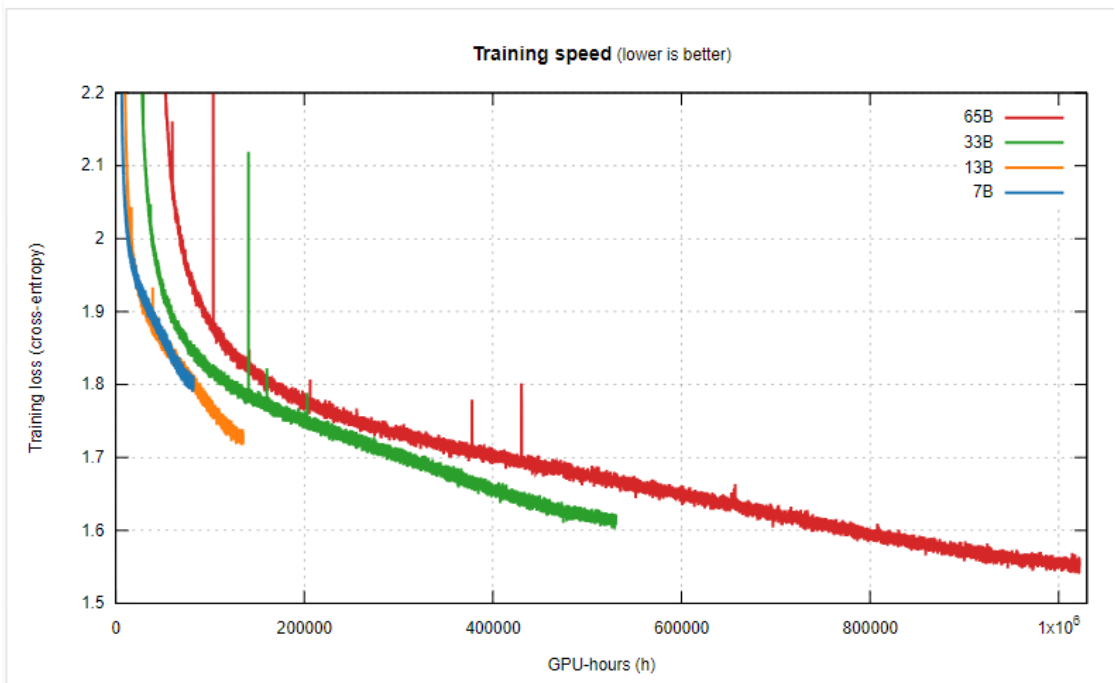
Let's picture instead a race: All those models start at the same time, and we want to know which one crosses the finish line first.

In other words, when throwing a fixed amount of compute at the training, who learns the most in that time?
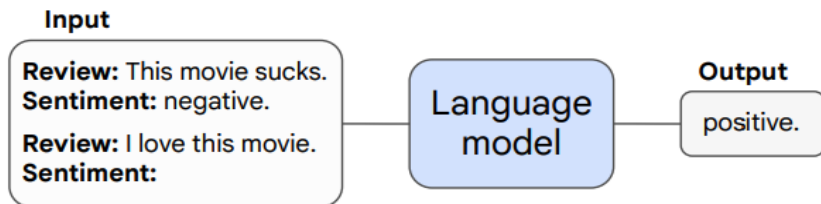
the 7B enters a near-linear regime, with a steep downward trend, and seems on its way to maybe overpass the 13B again?

https://espadrine.github.io/blog/posts/chinchilla-s-death.html
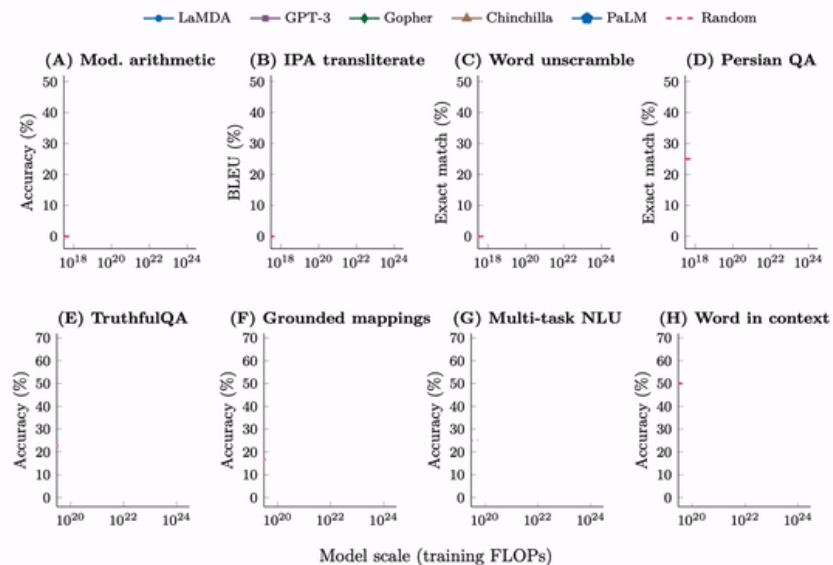
Emergent ability ?

Emergent properties in LLMs:
Some ability of LM is not present in smaller models but is present in larger models
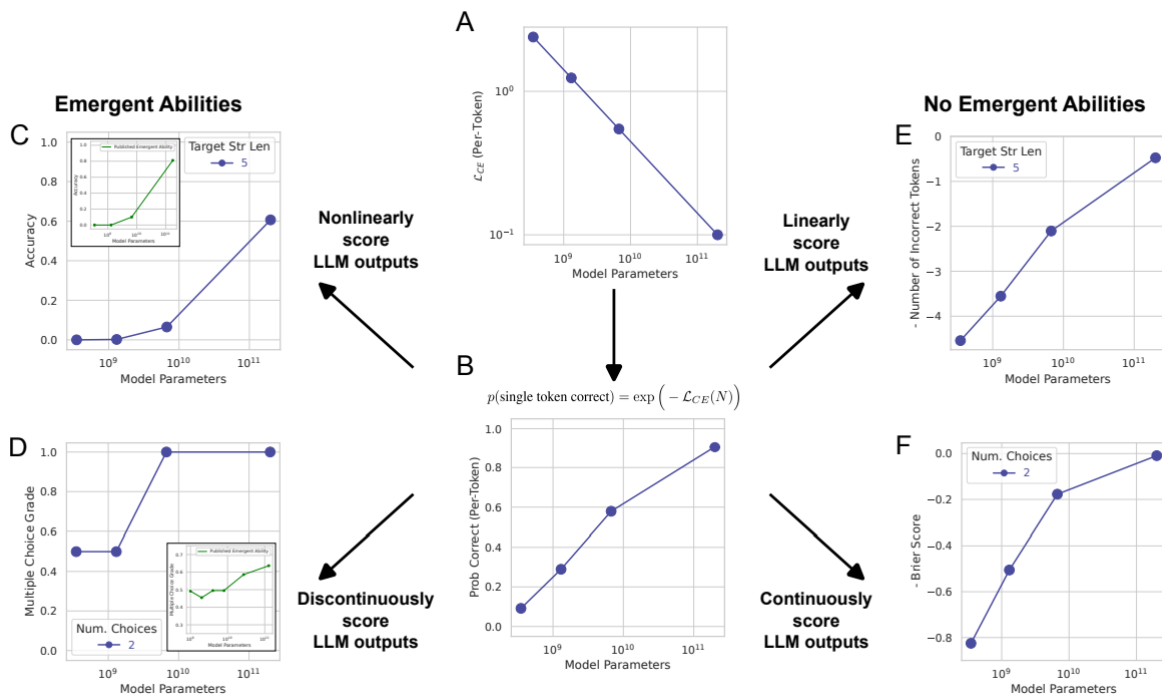
Emergent Capability: Few-shot prompting

> A few-shot prompted task is emergent if it achieves random accuracy for small models and above-random accuracy for large models.

Emergent capabilities may be a consequence of metric choice

It seems that emergent ability of a model only occurs if the measure of per-token error rate of any model is scaled **non-linearly or discontinuously**.

Rylan Schaeffer, Brando Miranda, and Sanmi Koyejo. Are Emergent Abilities of Large Language Models a Mirage? https://browse.arxiv.org/pdf/2304.15004.pdf

Assignment 1: Using LLM for prompt engineering

See updates in our BB system, WeChat and Emails. It will be released before the next lecture.

Tutorial and the homework explanation will be in the next lecture (Oct. 13th)

Acknowledgement

- Princeton COS 484: Natural Language Processing. Contextualized Word Embeddings. Fall 2019

- CS447: Natural Language Processing. Language Models. *http://courses.engr.illinois.edu/cs447*

- *http://cs231n.stanford.edu/*

- *https://medium.com/@gautam.karmakar/summary-seq2seq-model-using-convolutional-neural-network-b1eb100fb4c4*

- Transformers and sequence- to-sequence learning. CS 685, Fall 2021. Mohit Iyyer. College of Information and Computer Sciences. University of Massachusetts Amherst. https://people.cs.umass.edu/~miyyer/cs685_f21/slides/05-transformers.pdf

Challenge to scaling law: Chinchilla's Death

Can Chinchillas picture a Llama's sights?
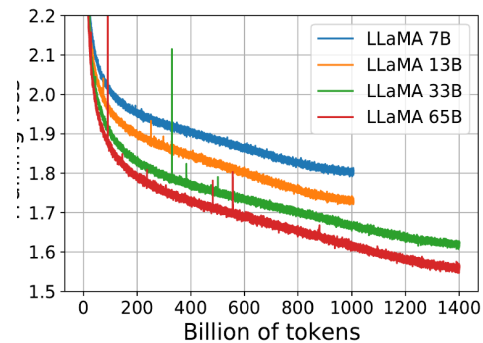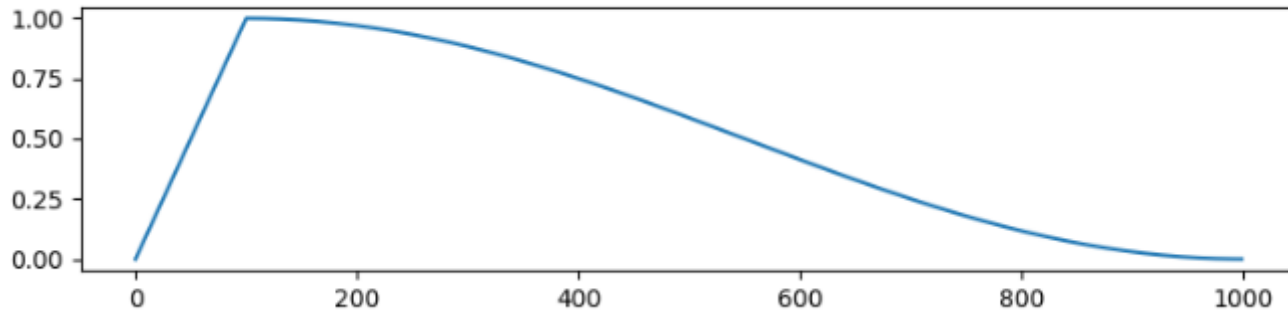
The slowdown in learning is an artefact of cosine schedule. The model does not necessarily cease to have the capacity to learn at the same near-linear rate!