

Recurrent Neural Networks

Gemma Roig
Computer Vision
Goethe University

Today's class objectives


- What are Recurrent Neural Networks
- When to use RNN
- Training RNN
- Long-Short Term Memory Networks (LSTM)

Recurrent Neural Networks

ar...

have seen neural networks to model one to one dependencies

one to one



Input: No sequence - example: image

Output: No sequence - examples: label of object class

Example:


- “standard” classification
- regression problems

Recurrent Neural Networks


How do we model sequences?

Input -> **Output**


one to one




one to many




many to one



many to many




many to many



Recurrent Neural Networks

How do we model sequences?



Input: No sequence


Output: Sequence

Example:

- Image captioning (image -> words)

Recurrent Neural Networks

How do we model sequences?



Input: Sequence

Output: No sequence


Example:

- sentence classification
- sentiment classification (words seq. \rightarrow sentiment)

Recurrent Neural Networks

How do we model sequences?


many to many



Example:

- machine translation,
(words seq-> words seq)

many to many



Example:

- video captioning


Input: Sequence

Output: Sequence

Recurrent Neural Networks

How do we model sequences?

one to many




Input: No sequence

Output: Sequence

Example:

- Image captioning (image -> words)

many to one




Input: Sequence

Output: No sequence

Example:

- sentence classification
- sentiment classification (words seq.->sentiment)

many to many




Input: Sequence

Output: Sequence

Example:

- machine translation, (words seq-> words seq)

many to many




Example:

- video captioning

Recurrent Neural Networks

Recurrent formula at each time step:



$$h_t = f_W(h_{t-1}, x_t)$$

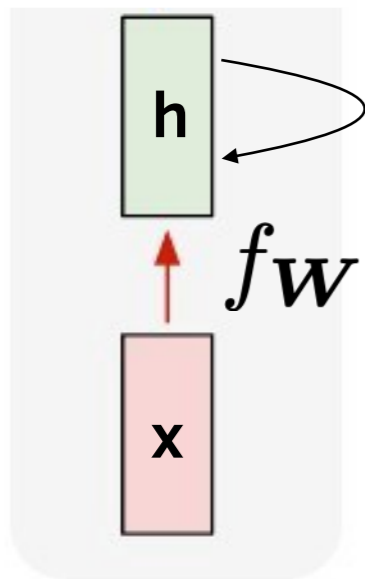
x : input sequence of vectors

h : hidden units, representing the state of the network

f_W : function with parameters W

Recurrent Neural Networks

Recurrent formula at each time step:



$$h_t = f_W(h_{t-1}, x_t)$$

new state

old state

input vector
at time t

x : input sequence of vectors

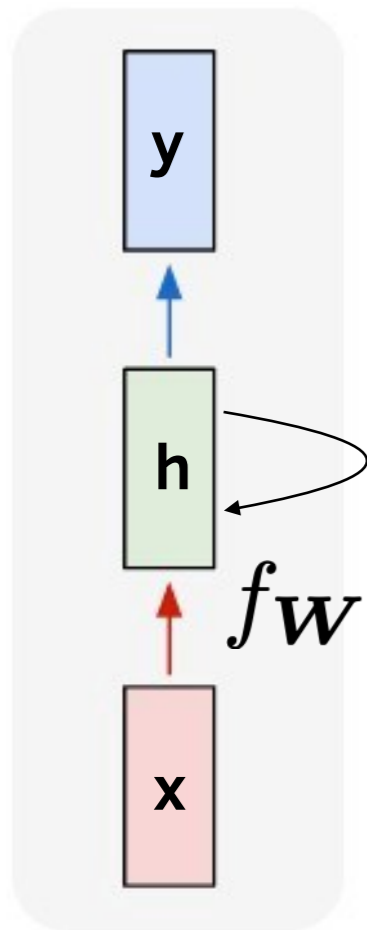
h : hidden units, representing the state of the network

f_W : function with parameters W .

The same function and same W are used at every time step

Recurrent Neural Networks

Example - RNN with one hidden vector h :



$$h_t = f_{\mathbf{W}}(h_{t-1}, \mathbf{x}_t)$$

$$f_{\mathbf{W}}(h_{t-1}, \mathbf{x}_t) = \tanh(\mathbf{W}_{hh}h_{t-1} + \mathbf{W}_{xh}\mathbf{x}_t)$$

$$y_t = \mathbf{W}_{hy}h_t$$

x : input sequence of vectors


h : hidden units, representing the state of the network

$f_{\mathbf{W}}$: function with parameters \mathbf{W}

y : output sequence

Recurrent Neural Networks

Computational Graph-
Unrolled recurrent neural network:




$$h_t = f_w(h_{t-1}, x_t)$$

- ▶ A recurrent neural network can be thought of as multiple copies of the same network, each passing a message to a successor

Recurrent Neural Networks

Computational Graph-
Unrolled recurrent neural network:




$$h_t = f_W(h_{t-1}, x_t)$$

Re-use the same weight matrix W at every time step

Recurrent Neural Networks

Computational Graph-
Unrolled recurrent neural network:




$$h_t = f_w(h_{t-1}, x_t)$$

Re-use the same weight matrix W at every time step

Recurrent Neural Networks


Computational Graph-
Unrolled recurrent neural network:



If the relevant input information is close to where is needed, the RNNs can learn to use the past information

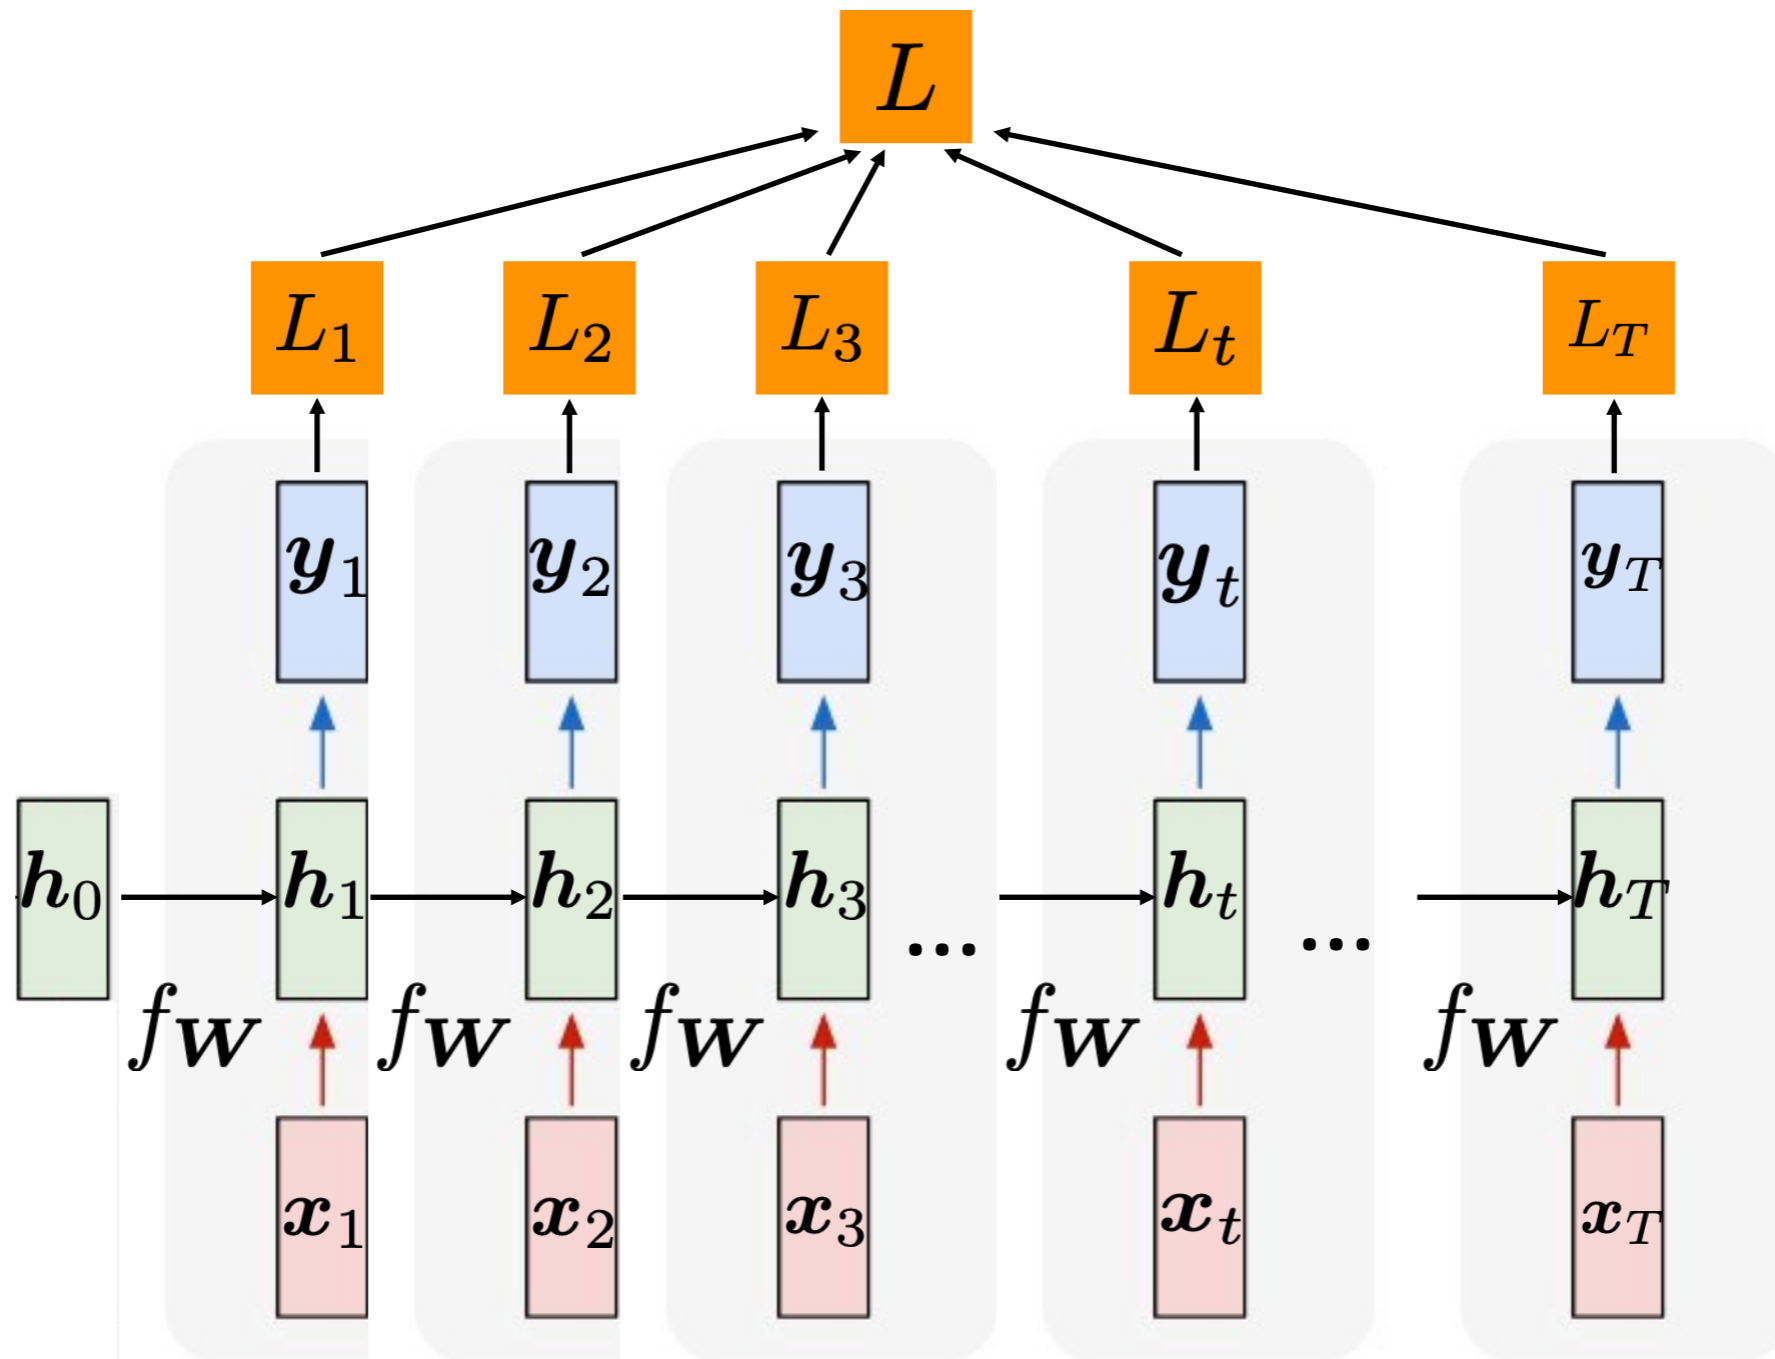
Recurrent Neural Networks

Computational Graph - Many to many




Recurrent Neural Networks

Computational Graph - Many to many




Recurrent Neural Networks

Computational Graph - Many to one



Recurrent Neural Networks

Computational Graph - One to many

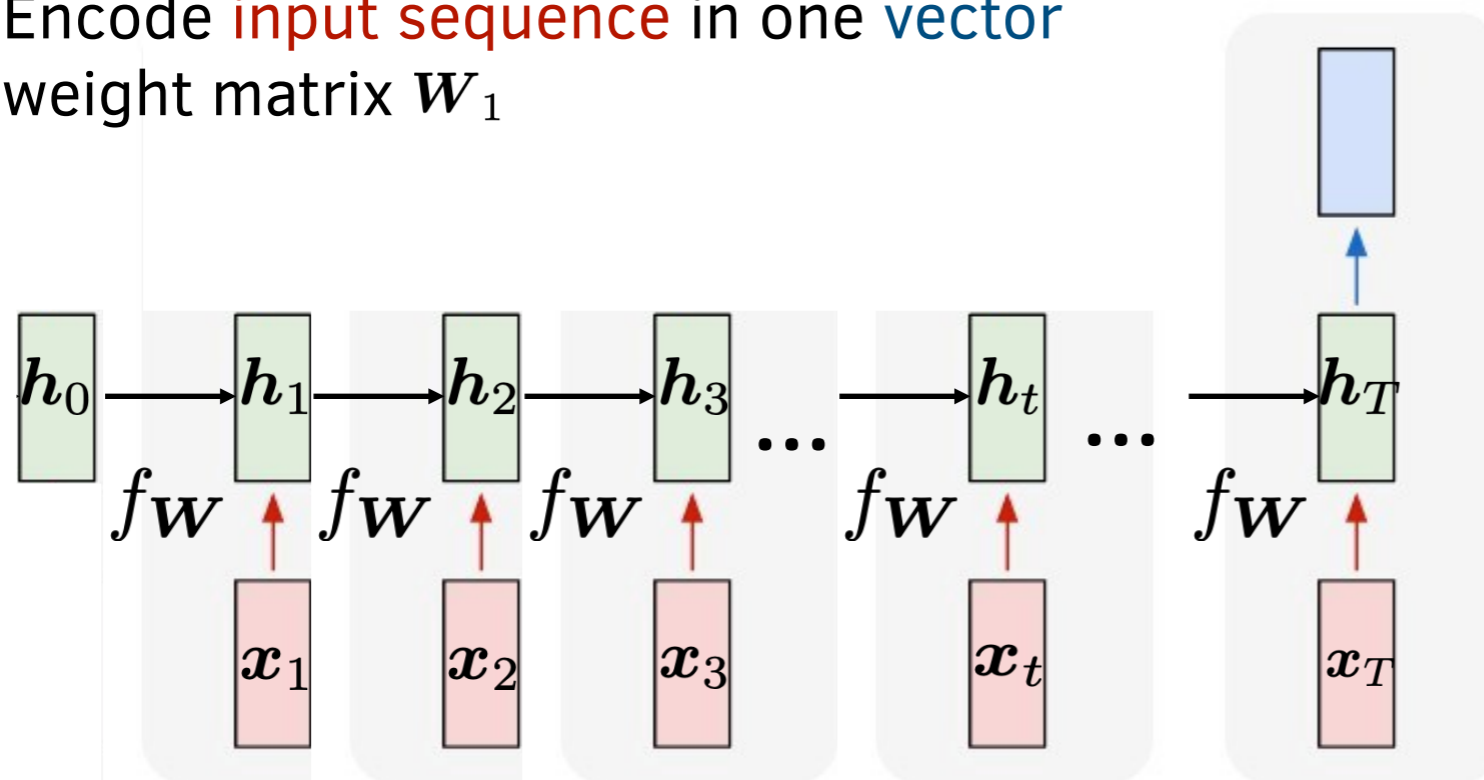


Recurrent Neural Networks

Computational Graph -
Sequence to sequence:
many to one + one to many

Many to one:

- Encode **input sequence** in one **vector**
- weight matrix W_1


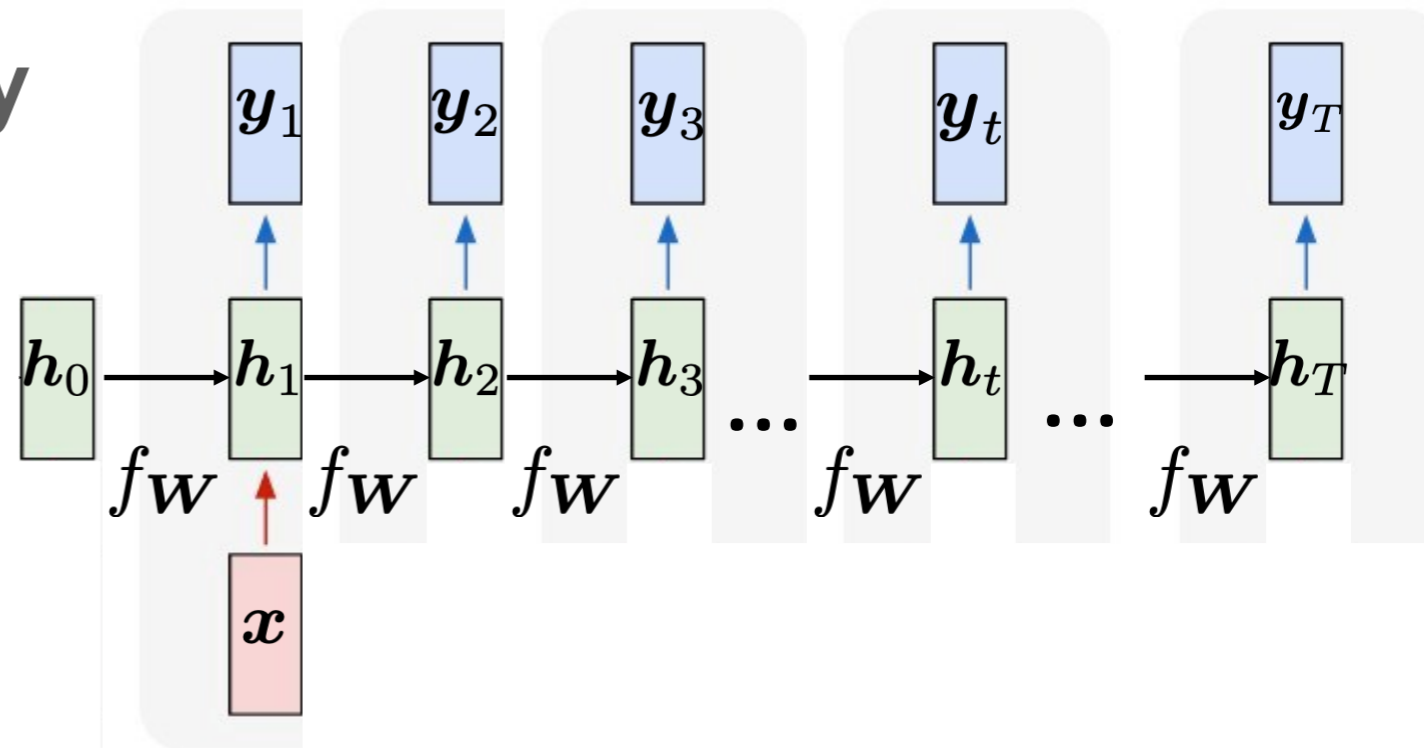


Recurrent Neural Networks

Computational Graph -
Sequence to sequence:
many to one + one to many

One to many:


- output sequence from input vector
- weight matrix W_2



Recurrent Neural Networks

Example: Character-level Language model

- Vocabulary:
[h,e,l,o]
- Example training:
“hello”



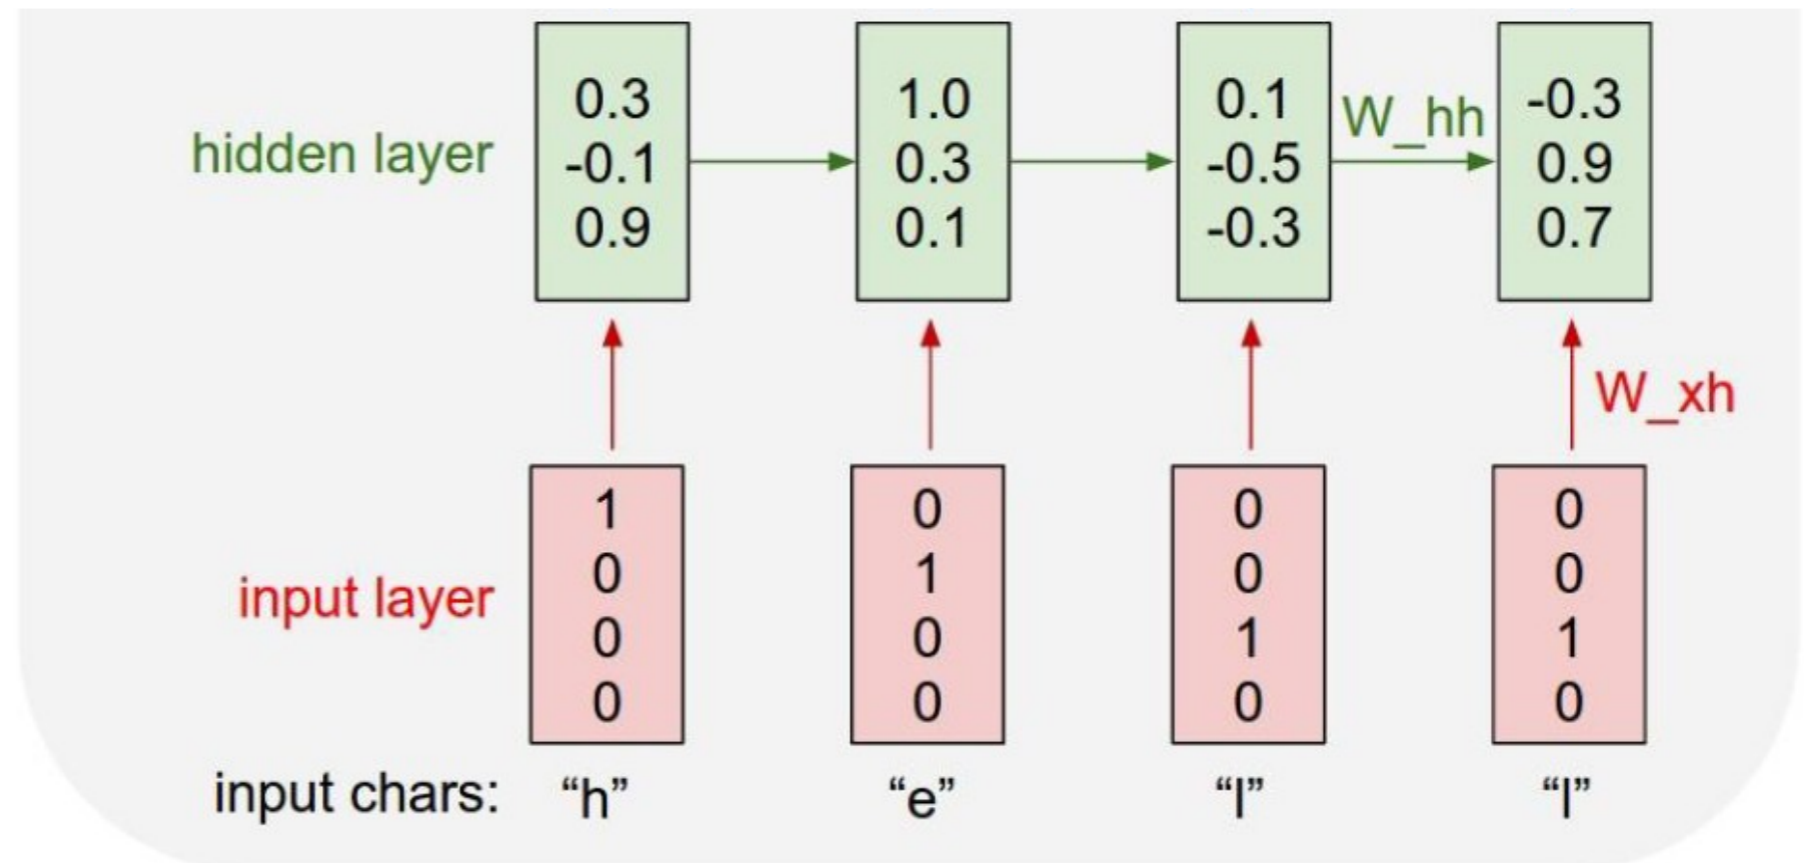
Recurrent Neural Networks

Example: Character-level Language model

- Vocabulary:
[h,e,l,o]

$$\mathbf{h}_t = \tanh(\mathbf{W}_{hh}\mathbf{h}_{t-1} + \mathbf{W}_{xh}\mathbf{x}_t)$$

- Example training:
“hello”

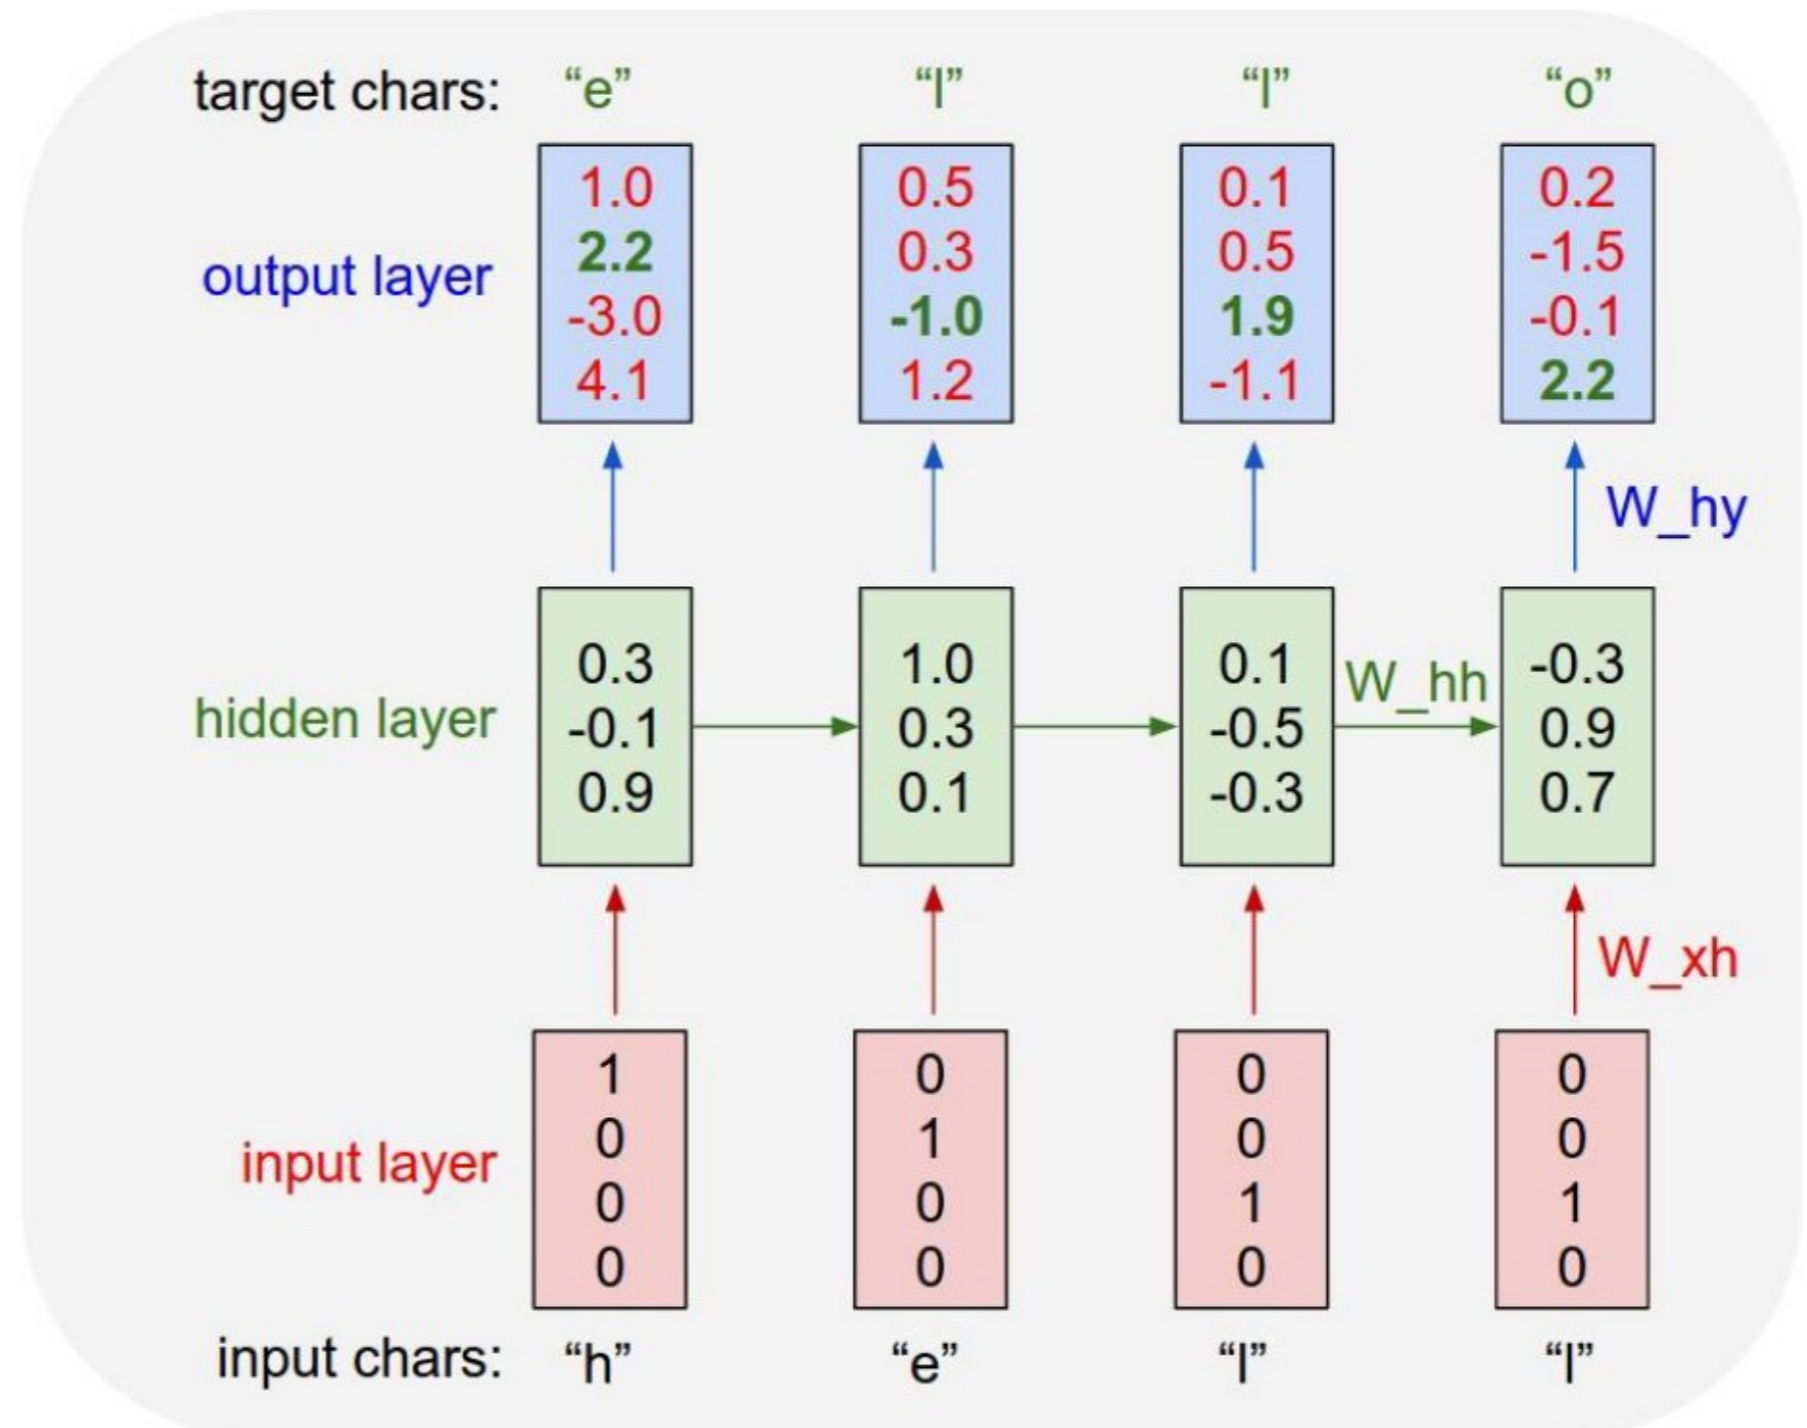


Recurrent Neural Networks

Example: Character-level Language model

- Vocabulary:
[h,e,l,o]
- Example training:
"hello"


$$y_t = W_{hy}h_t$$



Recurrent Neural Networks

Example: Character-level Language model


- Vocabulary:
[h,e,l,o]
- Test:
one character at a time



Recurrent Neural Networks

Example: Character-level Language model


- Vocabulary:
[h,e,l,o]
- Test:
one character at a time



Recurrent Neural Networks

Example: Character-level Language model

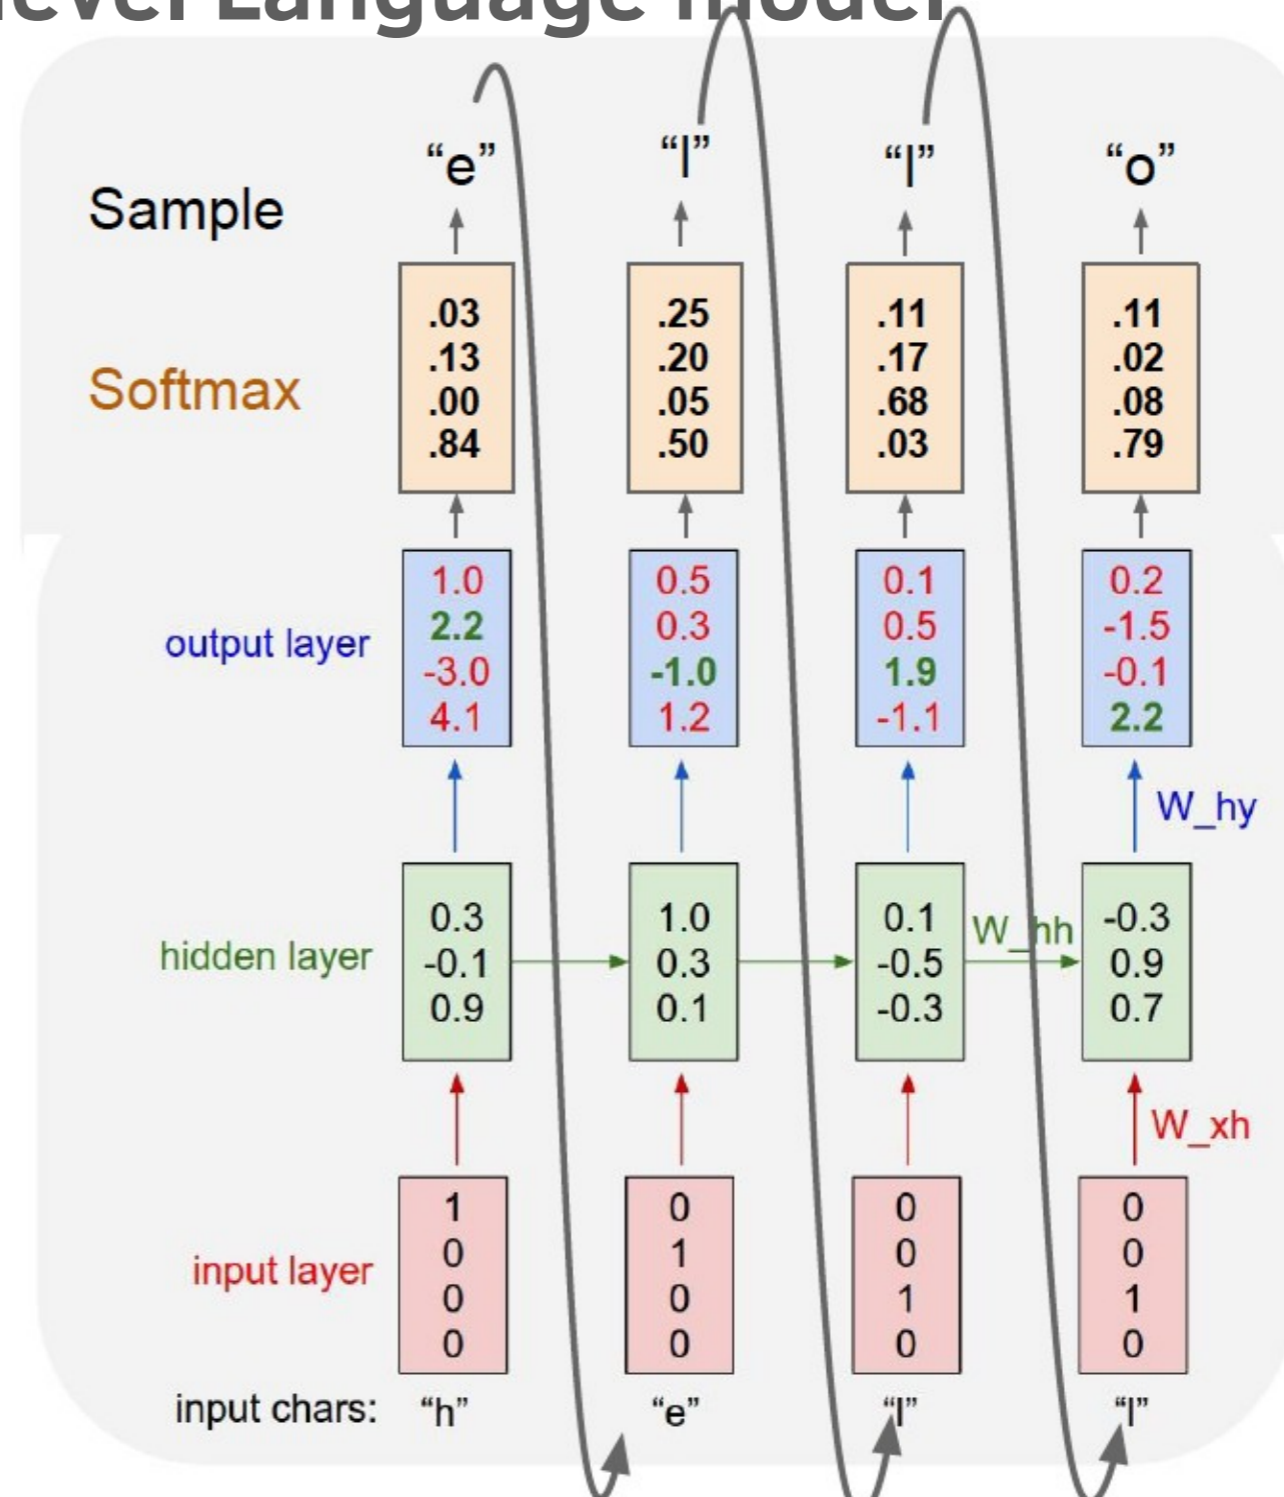
- Vocabulary:
[h,e,l,o]
- Test:
one character at a time



Recurrent Neural Networks

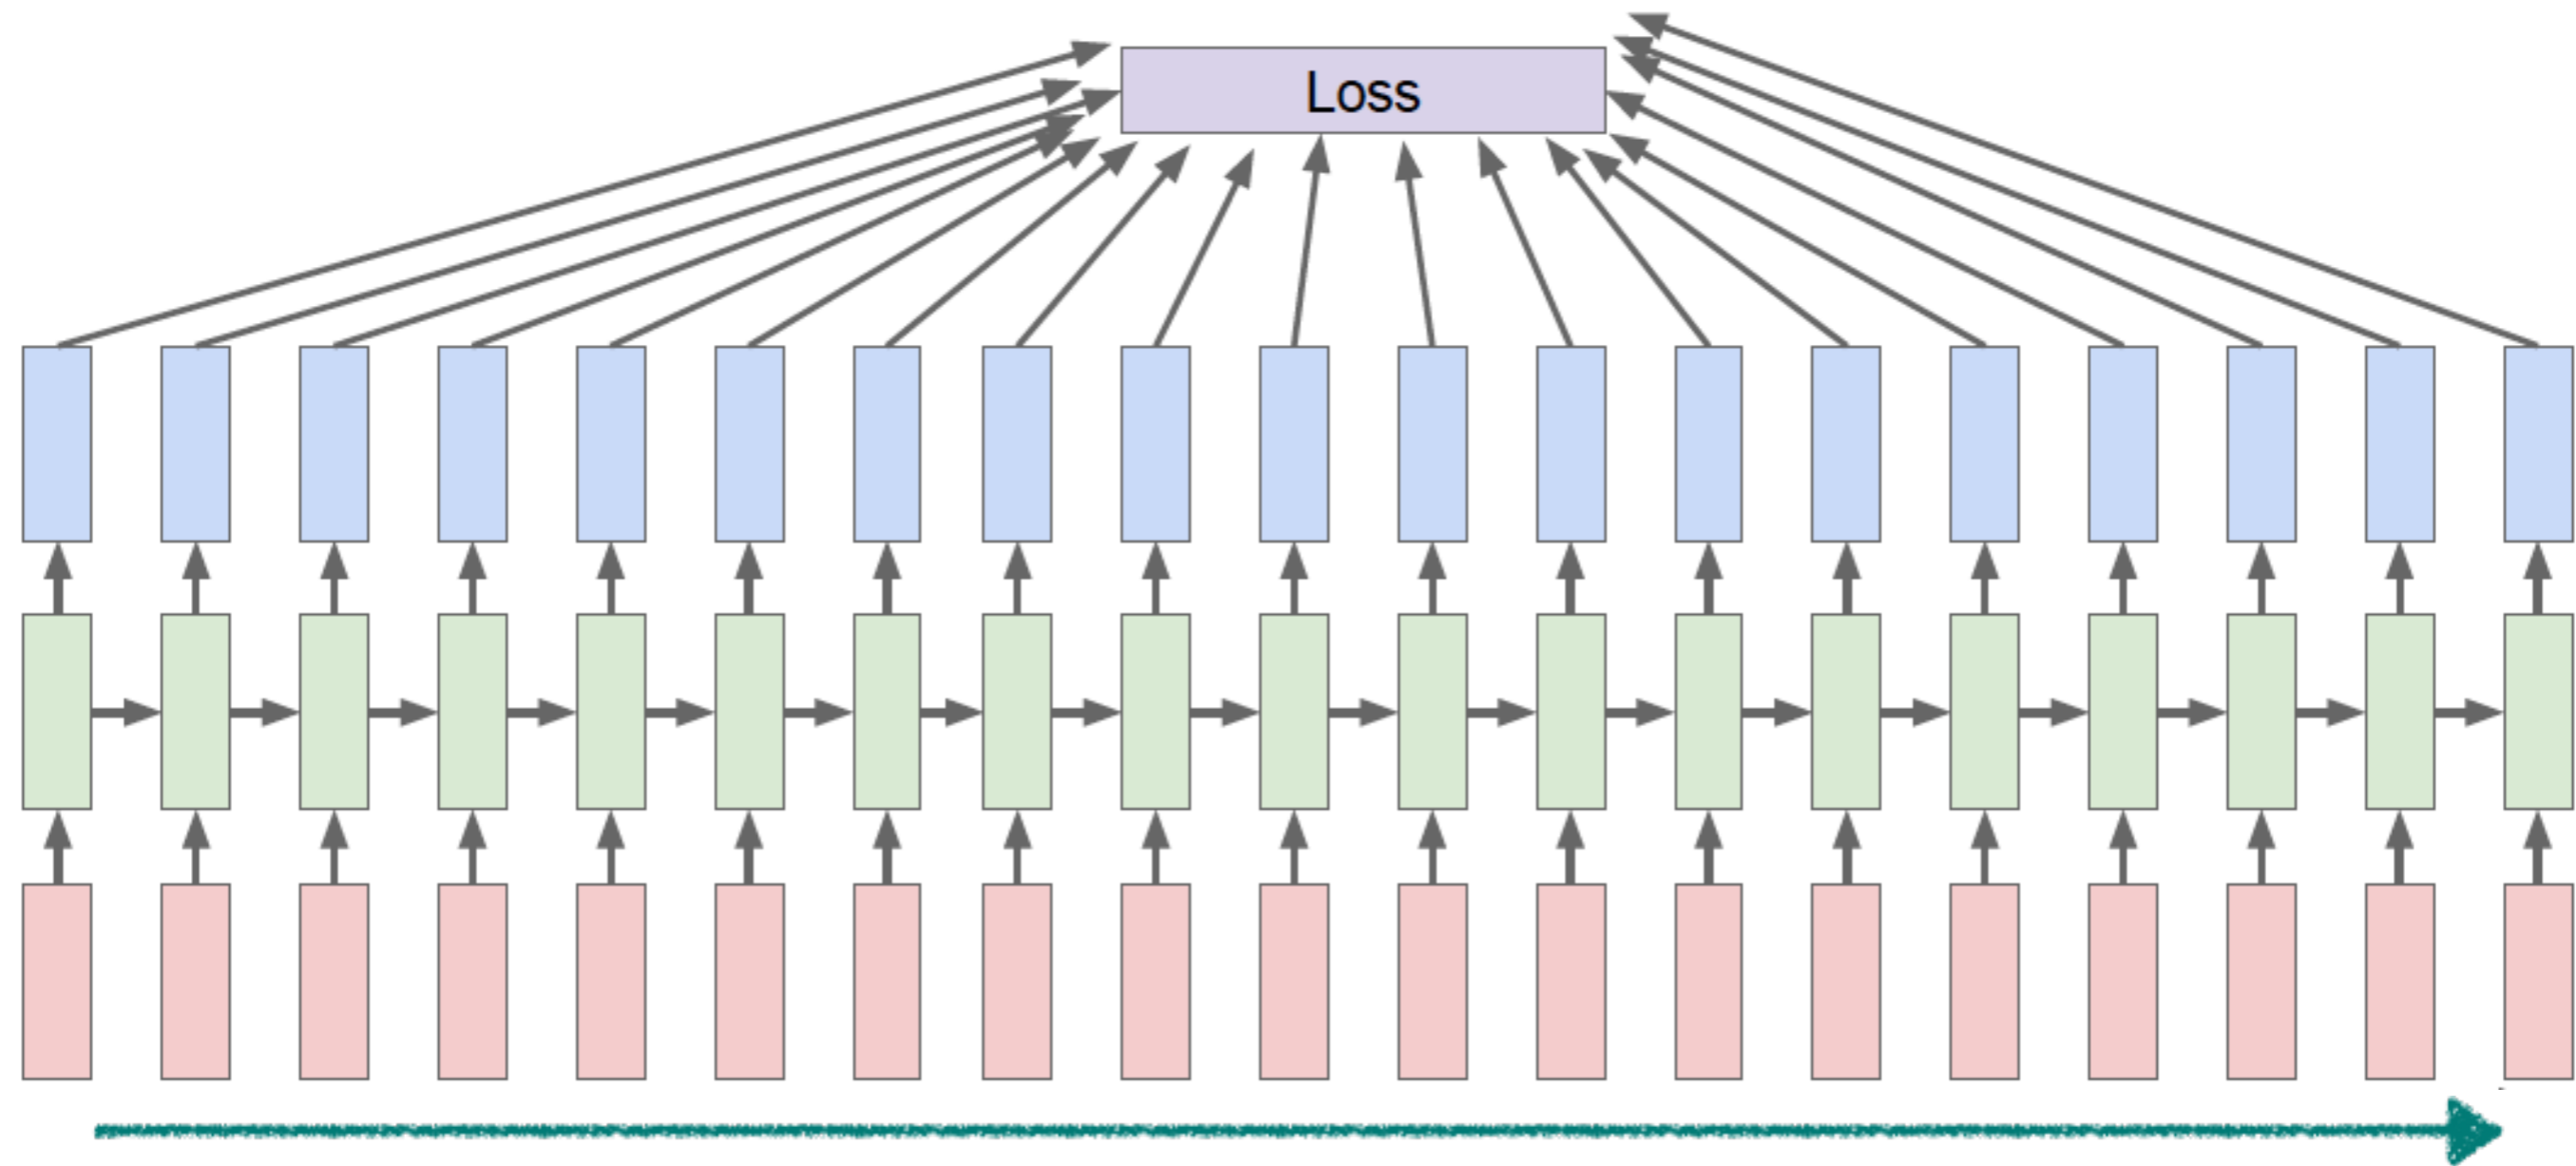
Example: Character-level Language model

- Vocabulary: [h,e,l,o]
- Test: one character at a time



Recurrent Neural Networks

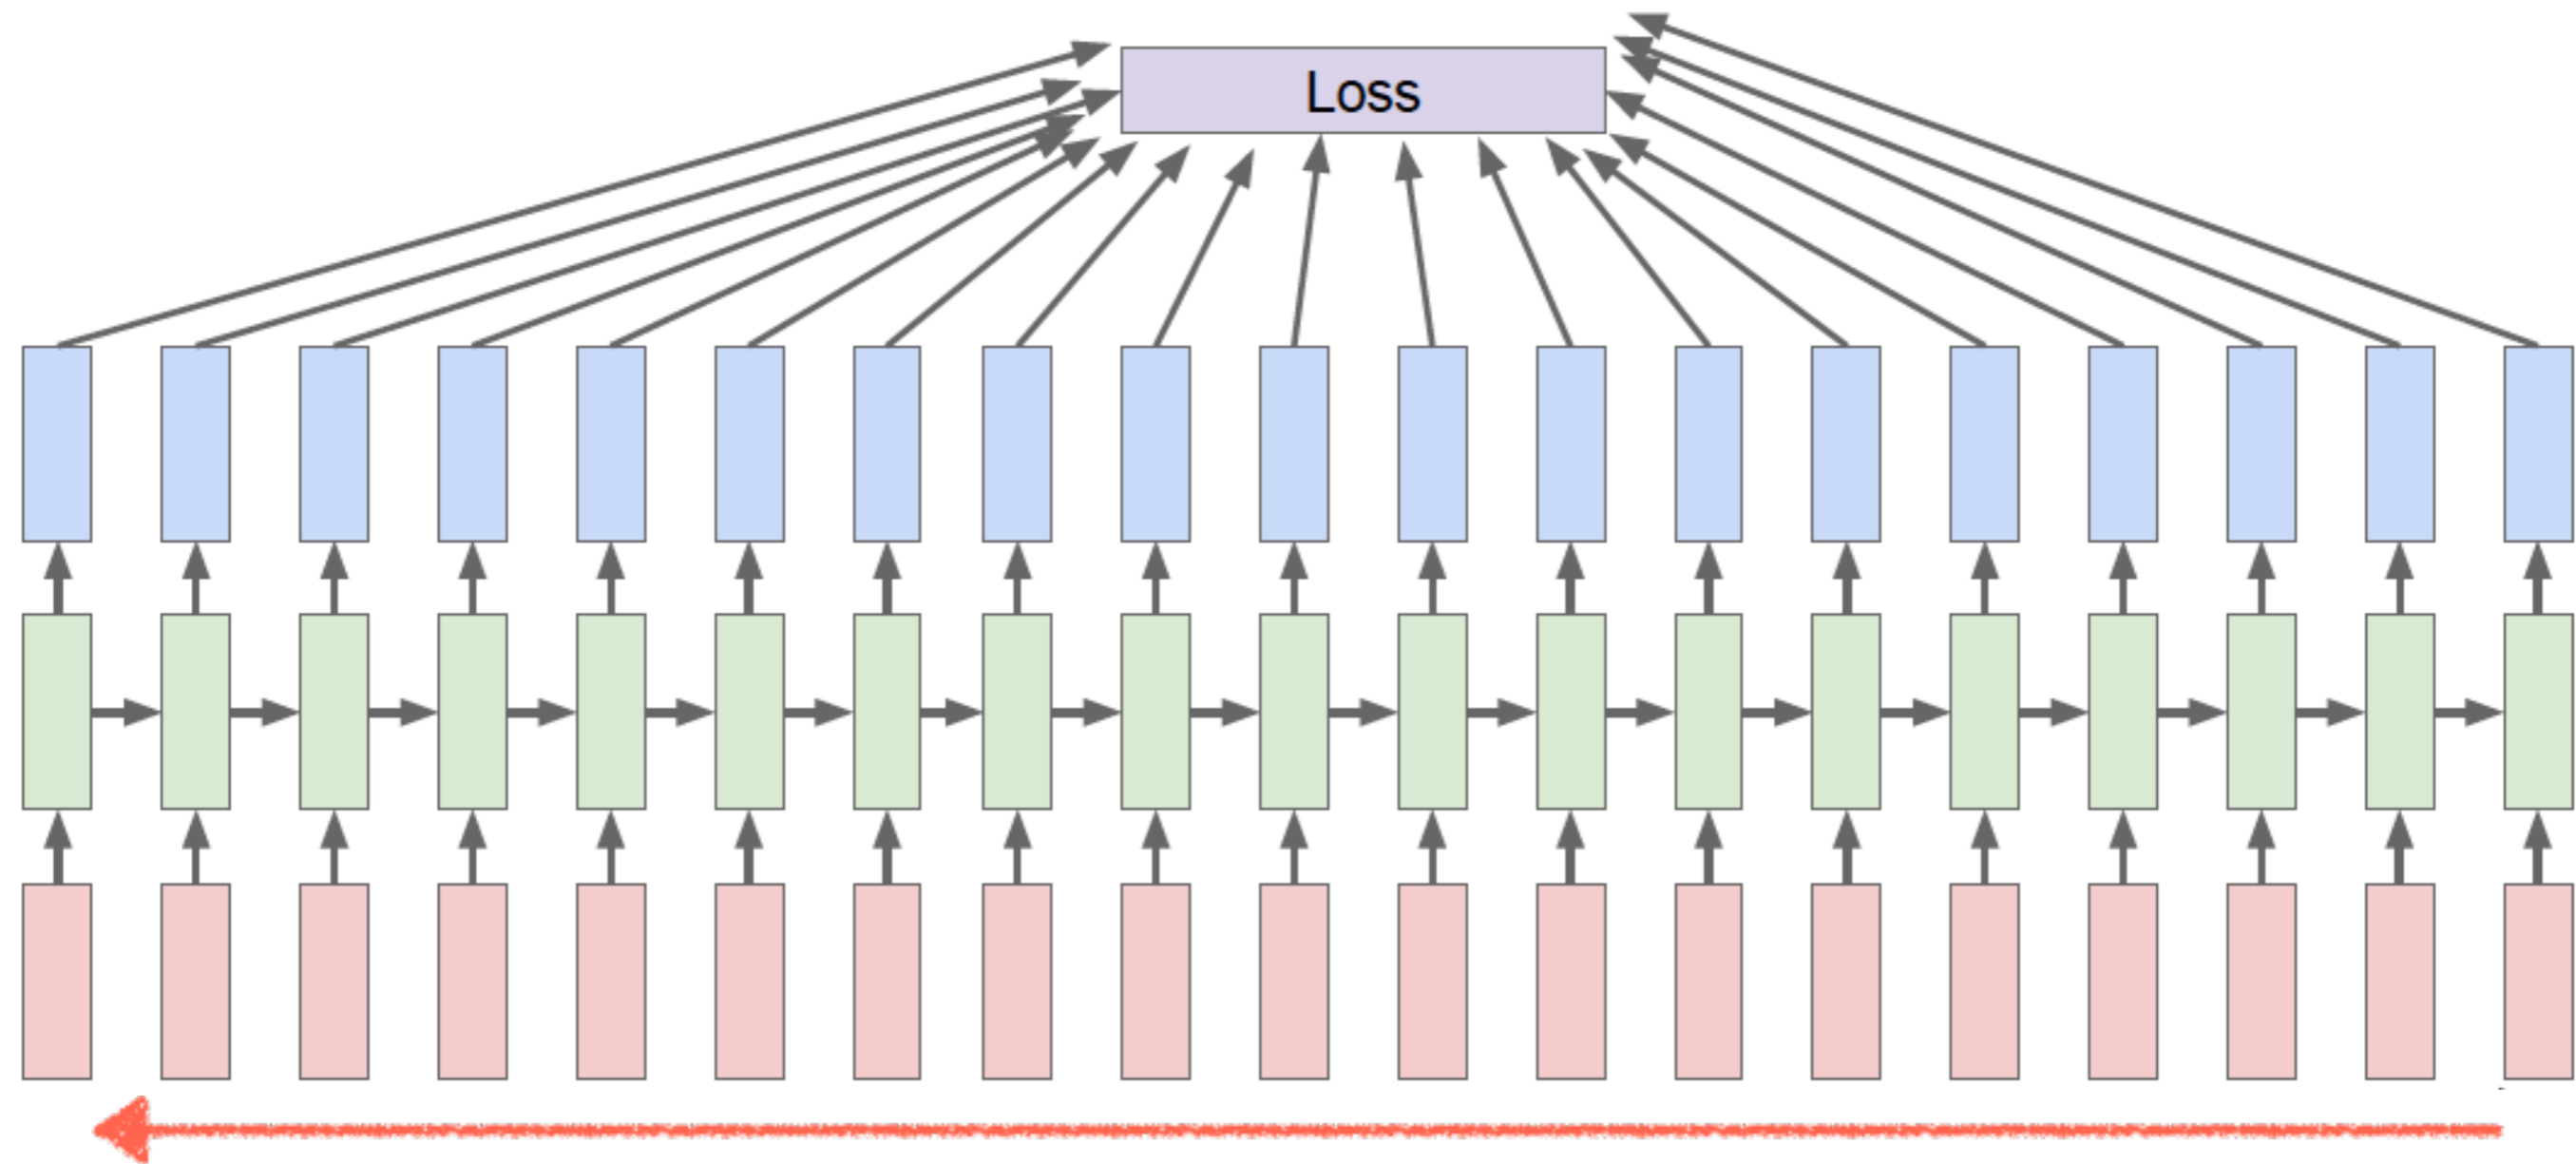
Learning of the weights - back-propagation through time



Forward through all the sequence to compute the loss

Recurrent Neural Networks


Learning of the weights - back-propagation through time



Backward through all the sequence to compute the gradient

Recurrent Neural Networks

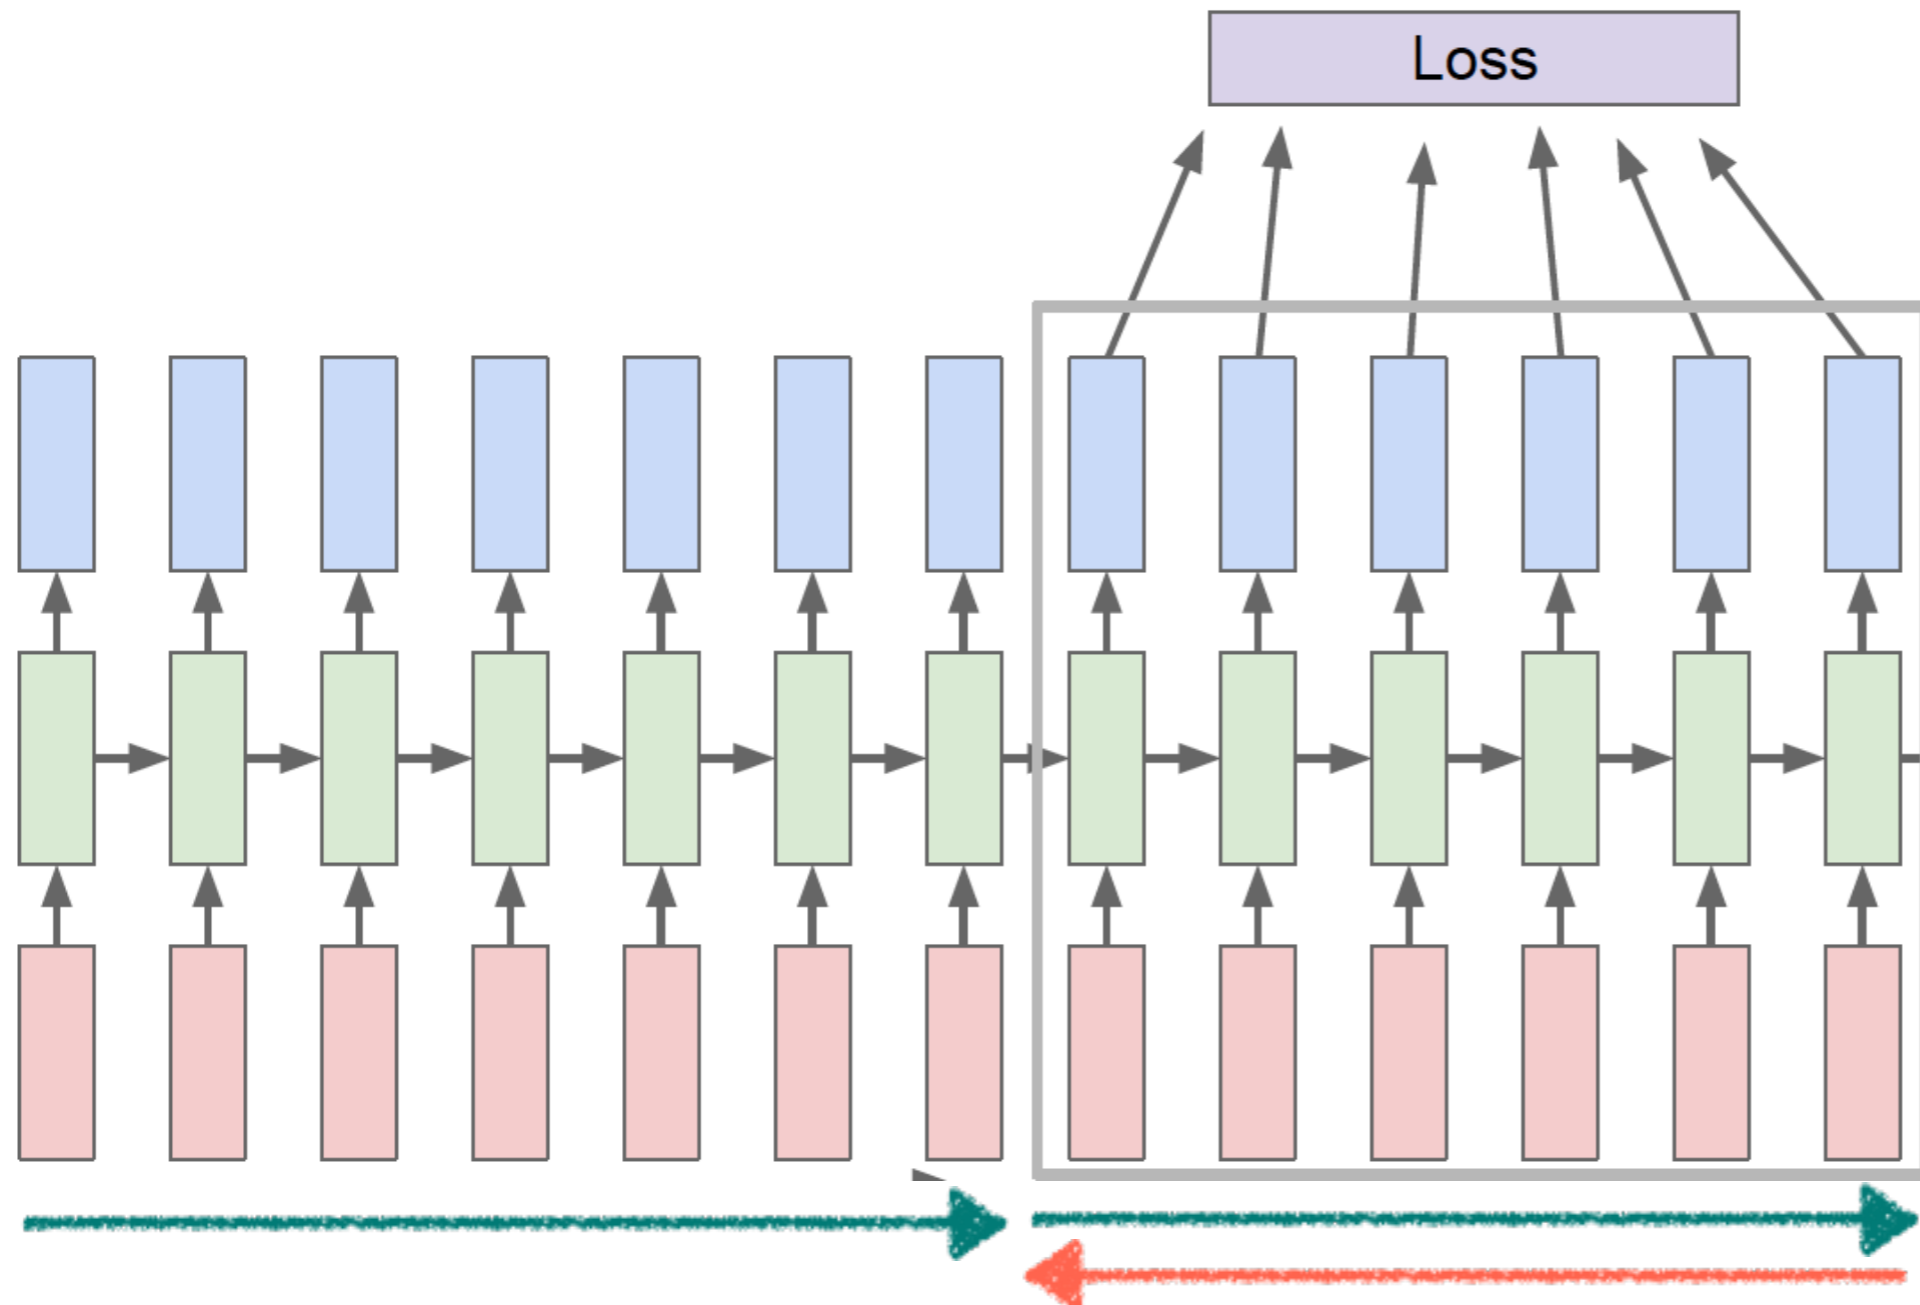
Truncated back-propagation through time



Run forward and backward through chunks of the sequence

Recurrent Neural Networks


Truncated back-propagation through time



Carry hidden states forward in time, but only back-propagate for some smaller number of steps

Recurrent Neural Networks

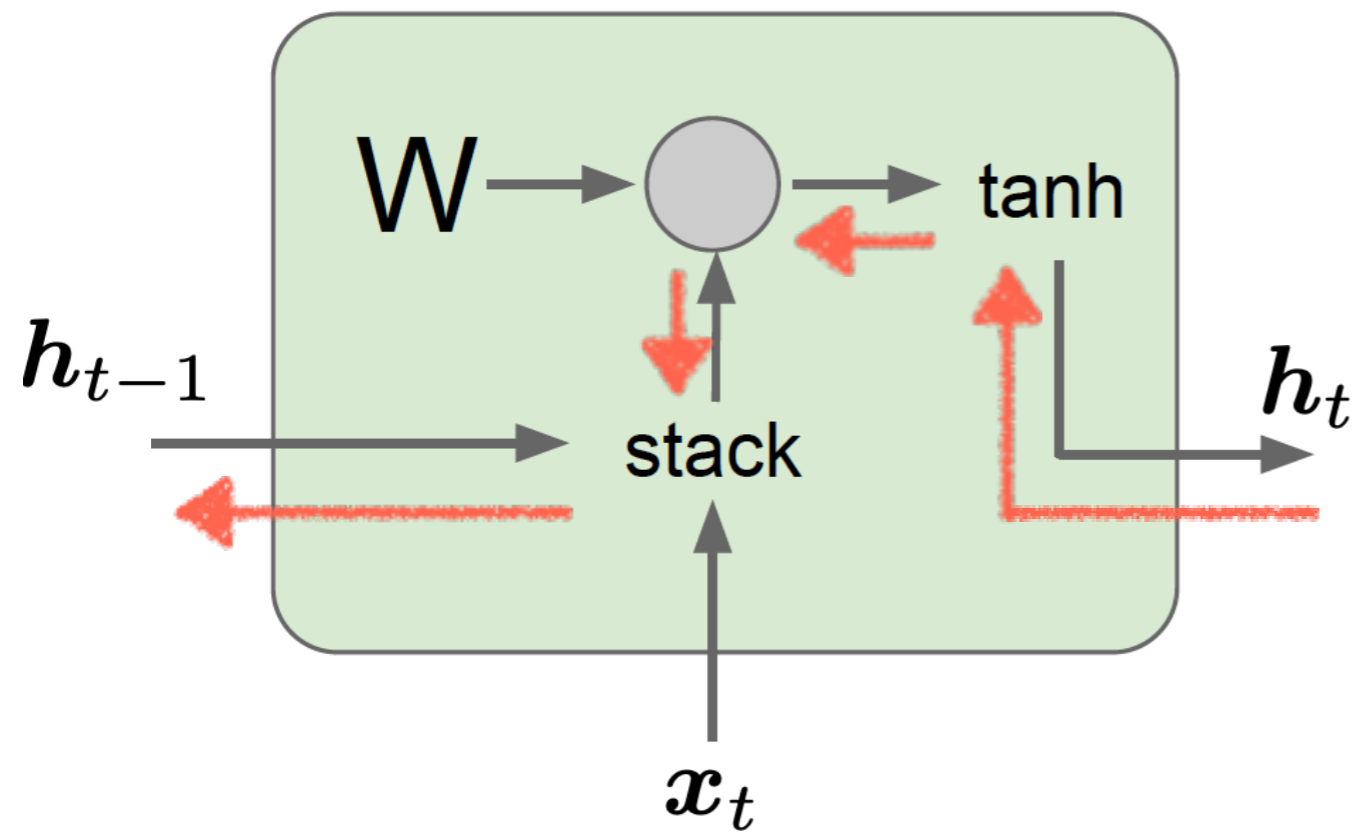
Truncated back-propagation through time



Carry hidden states forward in time, but only back-propagate for some smaller number of steps

Recurrent Neural Networks

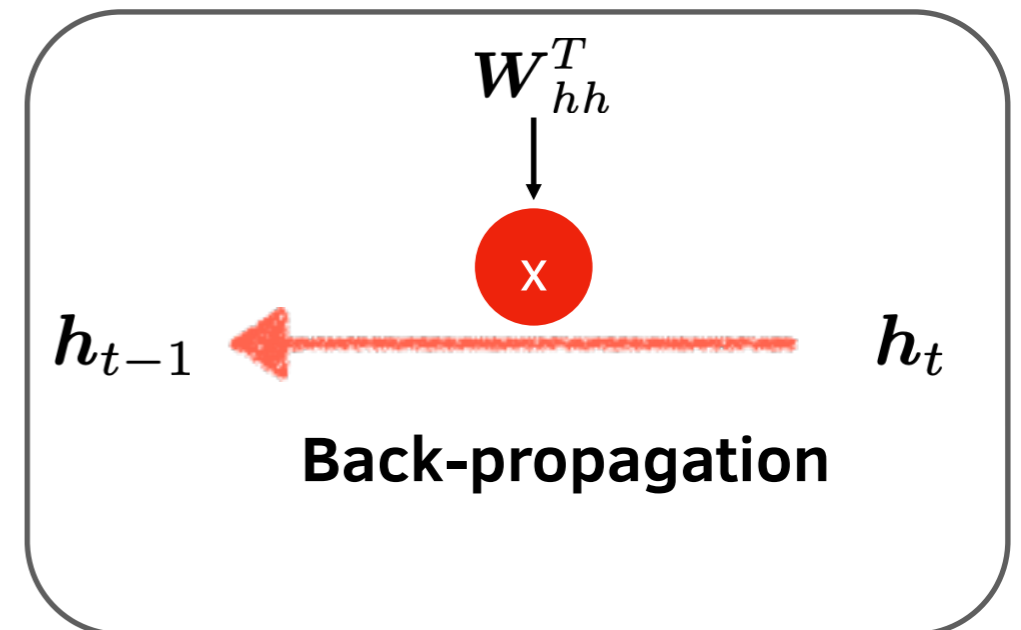
Backpropagation - Gradient flow



$$h_t = \tanh(\mathbf{W}_{hh}h_{t-1} + \mathbf{W}_{xh}x_t)$$

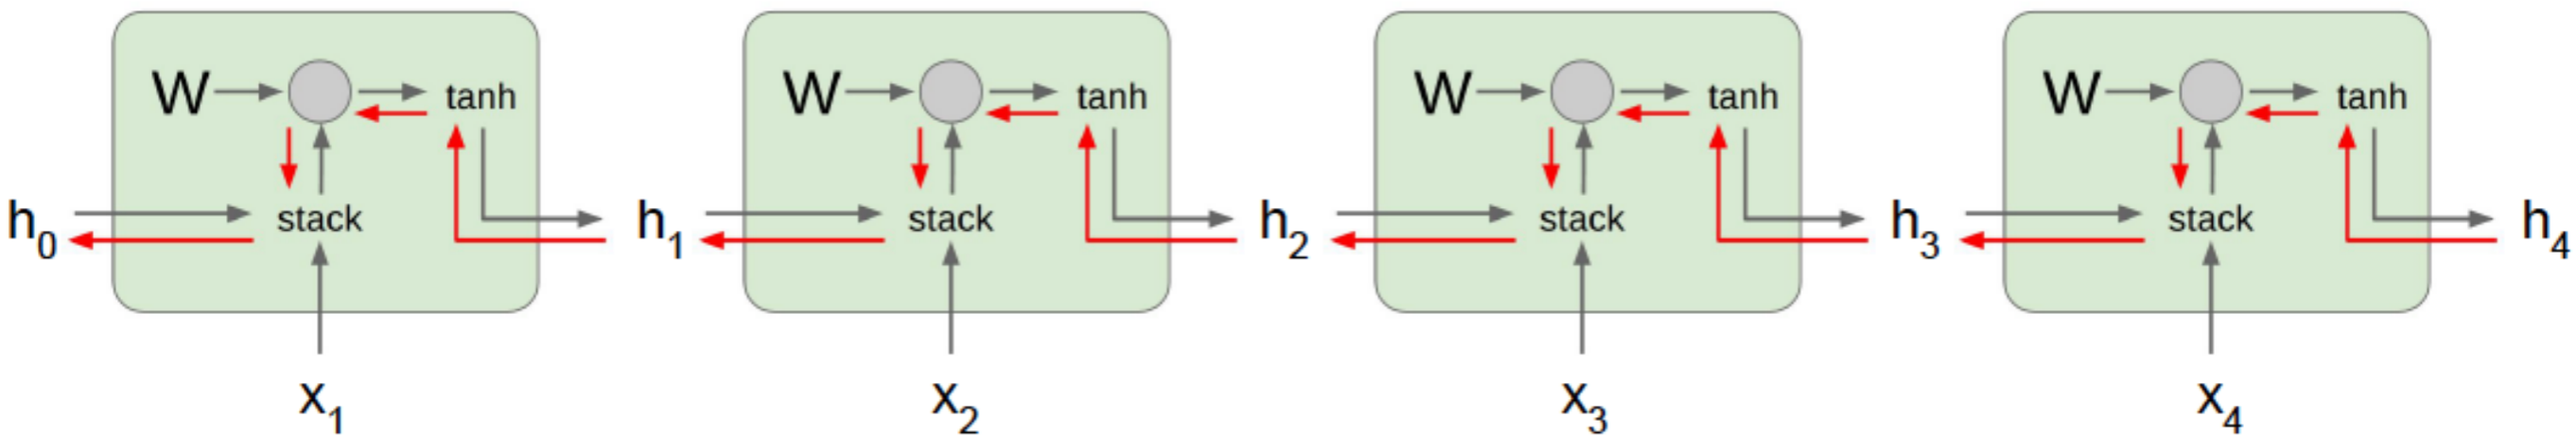
$$h_t = \tanh\left(\begin{bmatrix} \mathbf{W}_{hh} & \mathbf{W}_{hx} \end{bmatrix} \begin{bmatrix} h_{t-1} \\ x_t \end{bmatrix}\right)$$

$$h_t = \tanh\left(\mathbf{W} \begin{bmatrix} h_{t-1} \\ x_t \end{bmatrix}\right)$$



Recurrent Neural Networks

Backpropagation - Gradient flow



Computing gradient of h_0 involves many factors of W and repeated \tanh

- * **Exploding gradients** if largest singular value > 1
 - **SOLUTION:** gradient clipping with a threshold
- * **Vanishing gradient** if largest singular value < 1
 - **SOLUTION:** change RNN architecture

Recurrent Neural Networks

Example - Image Captioning

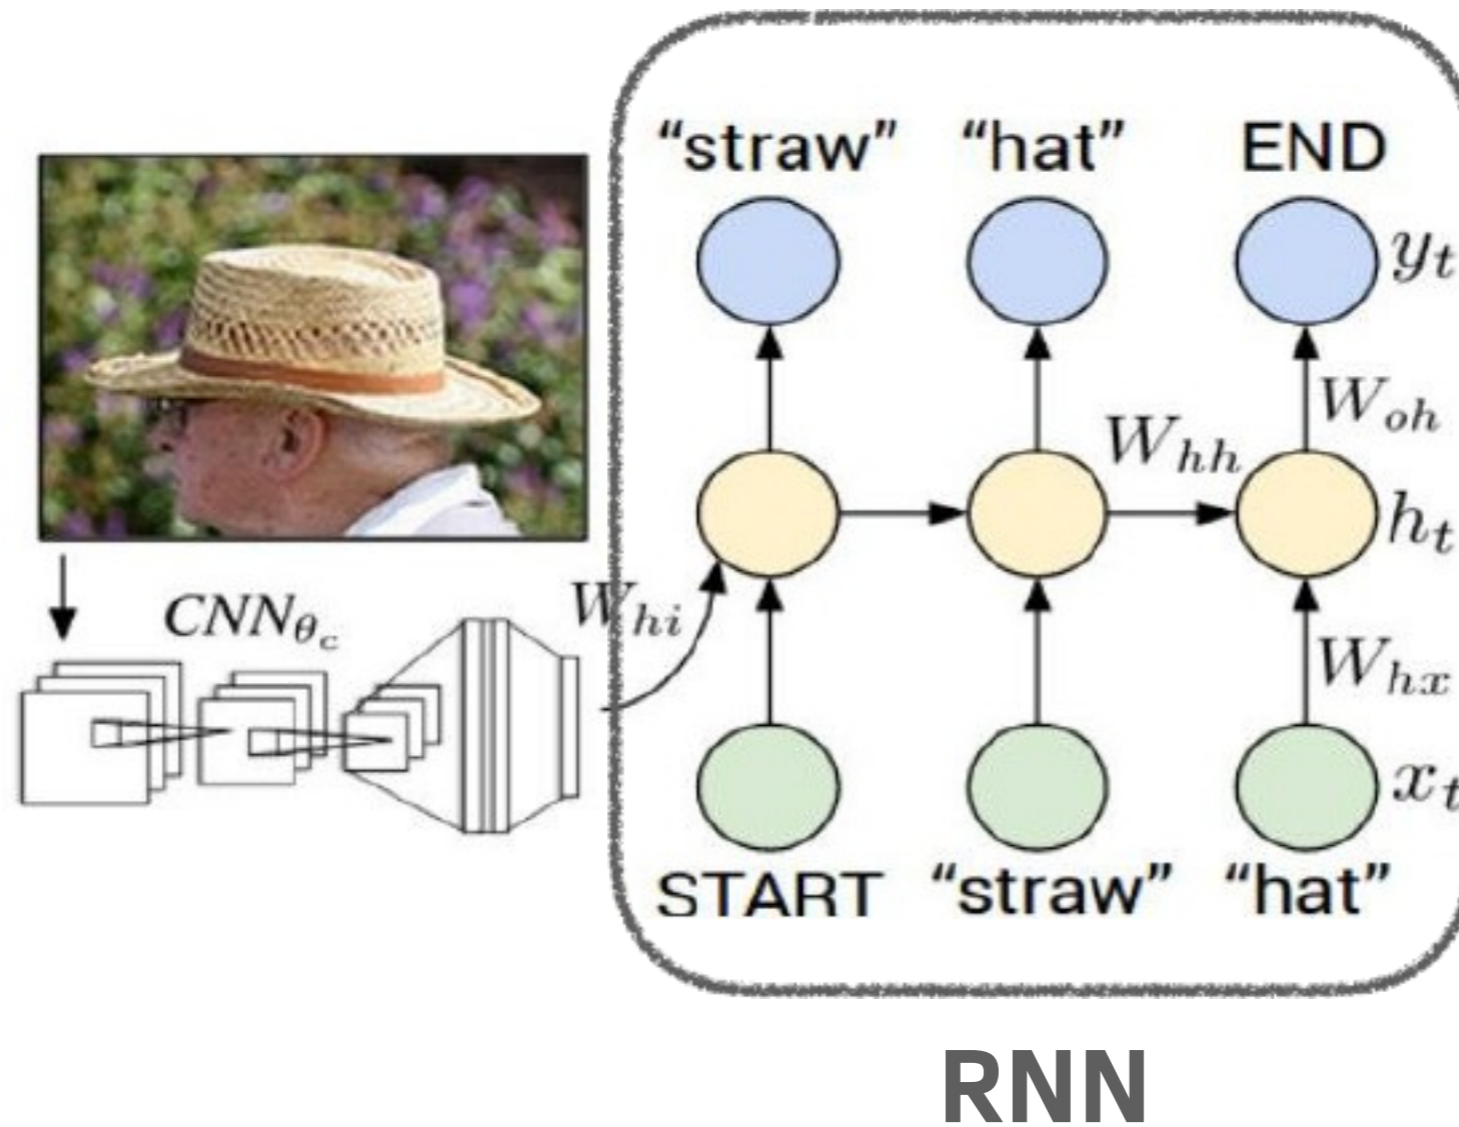
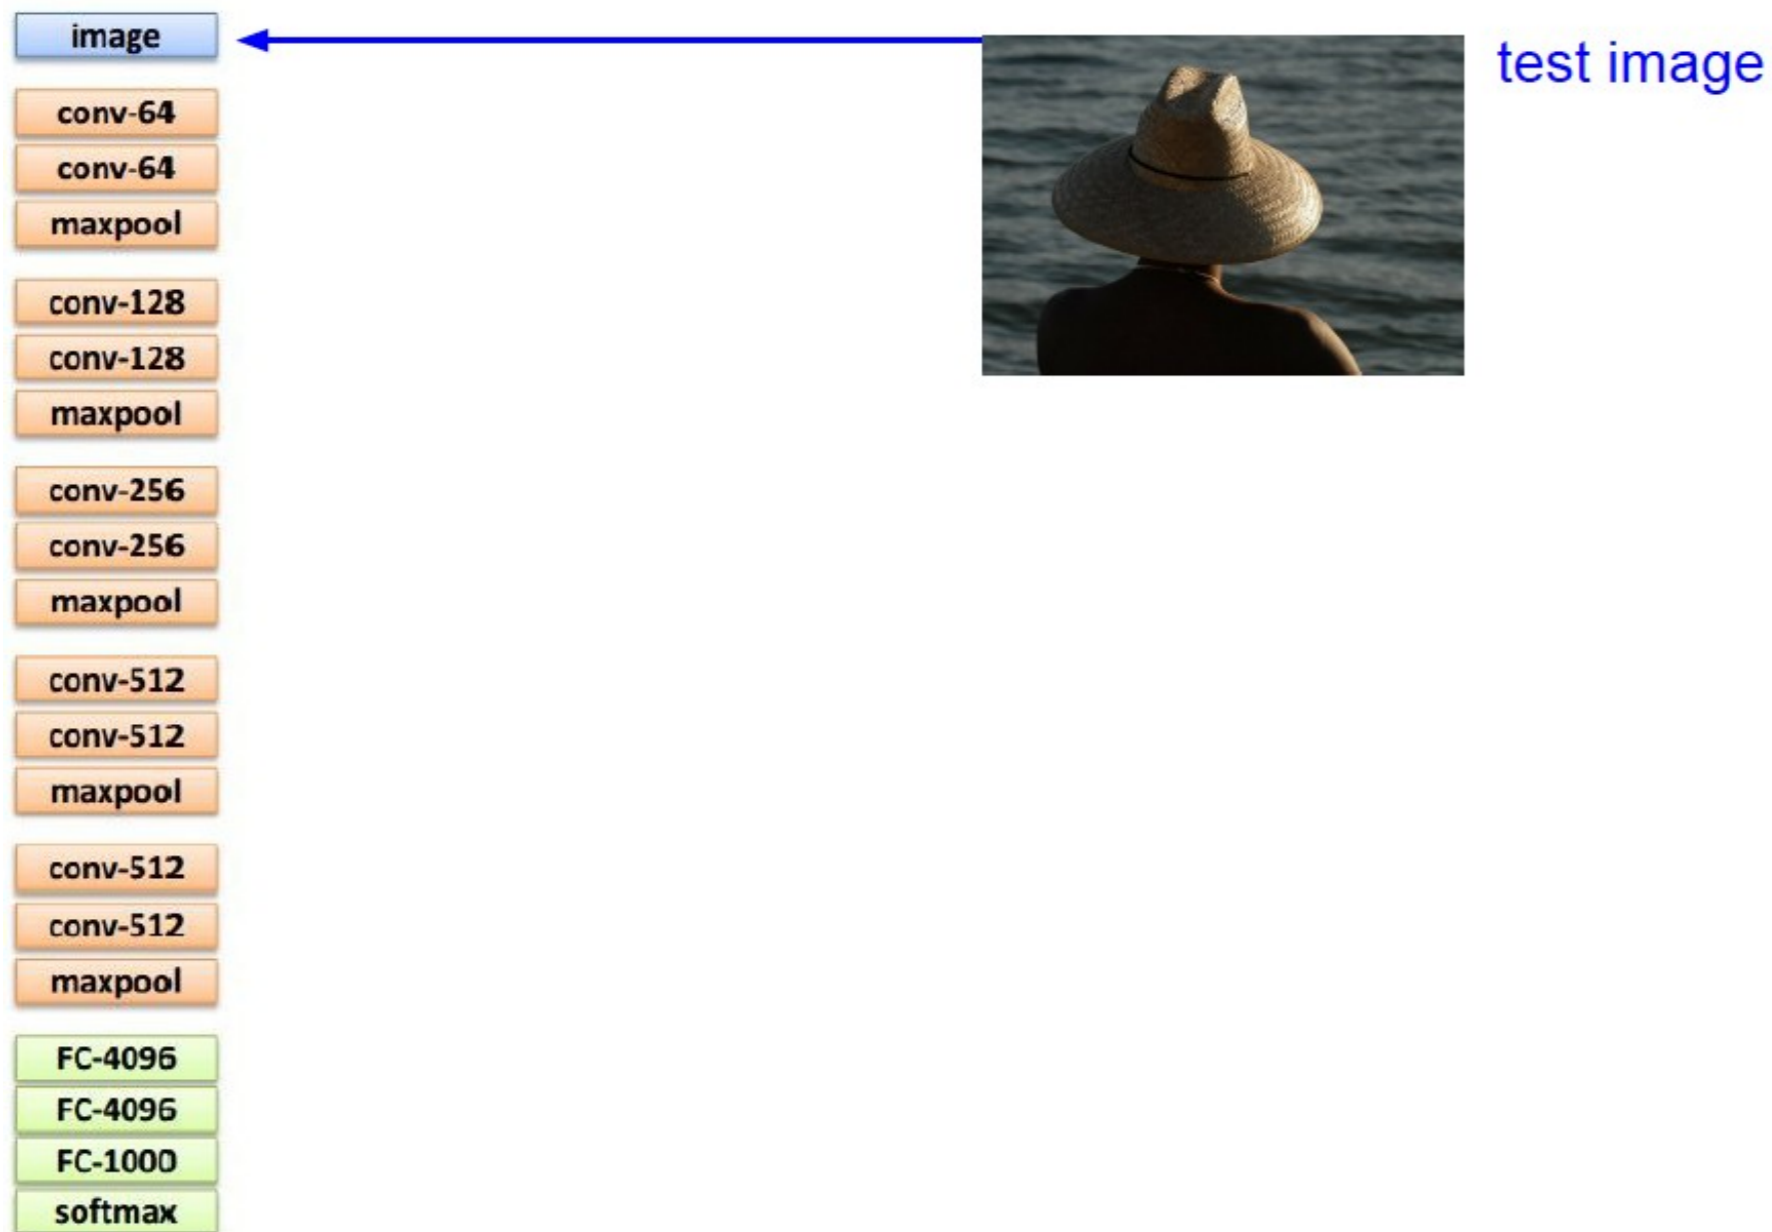


Figure: Karpathy et al.
Deep Visual Semantics Alignments for Generating Image Descriptions.
CVPR, 2015.

Recurrent Neural Networks

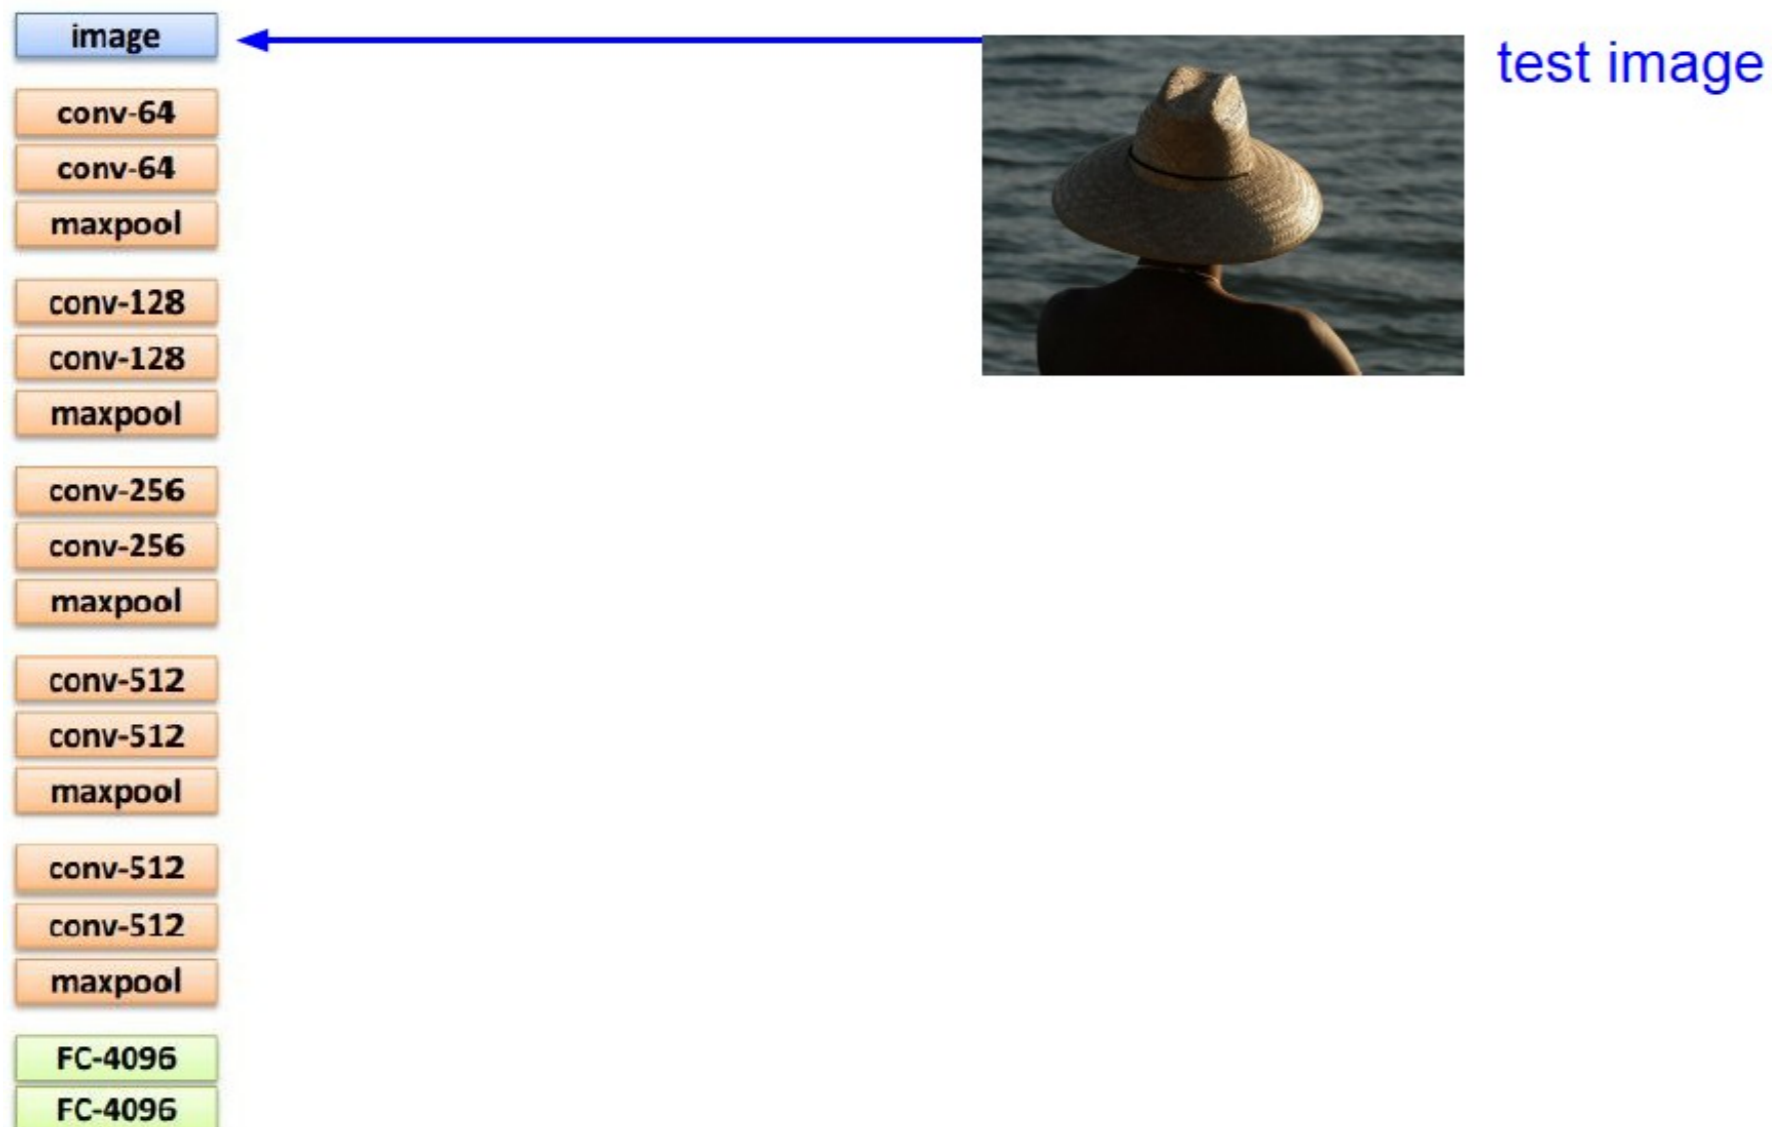
Example - Image Captioning



CNN trained on ImageNet

Recurrent Neural Networks

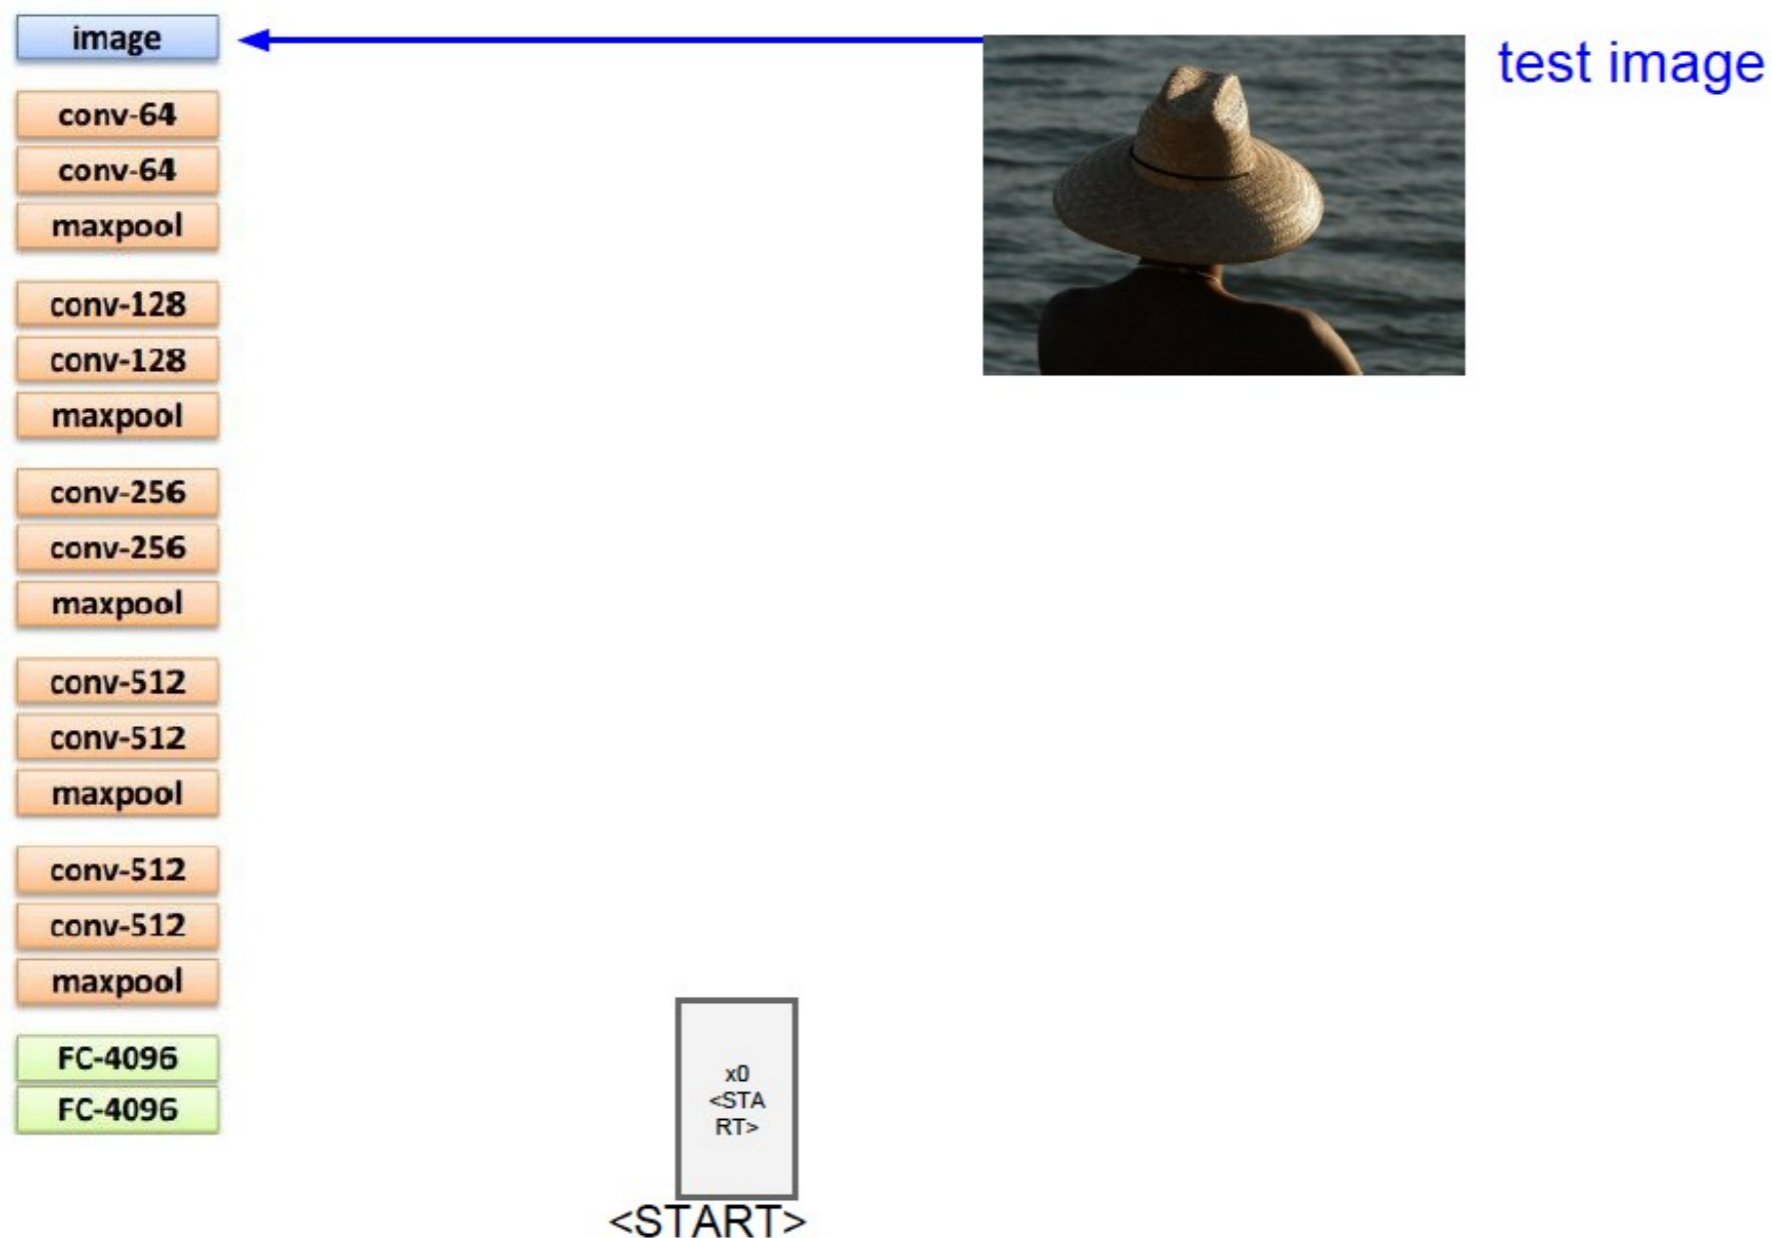
Example - Image Captioning



Take features before the last FC layer

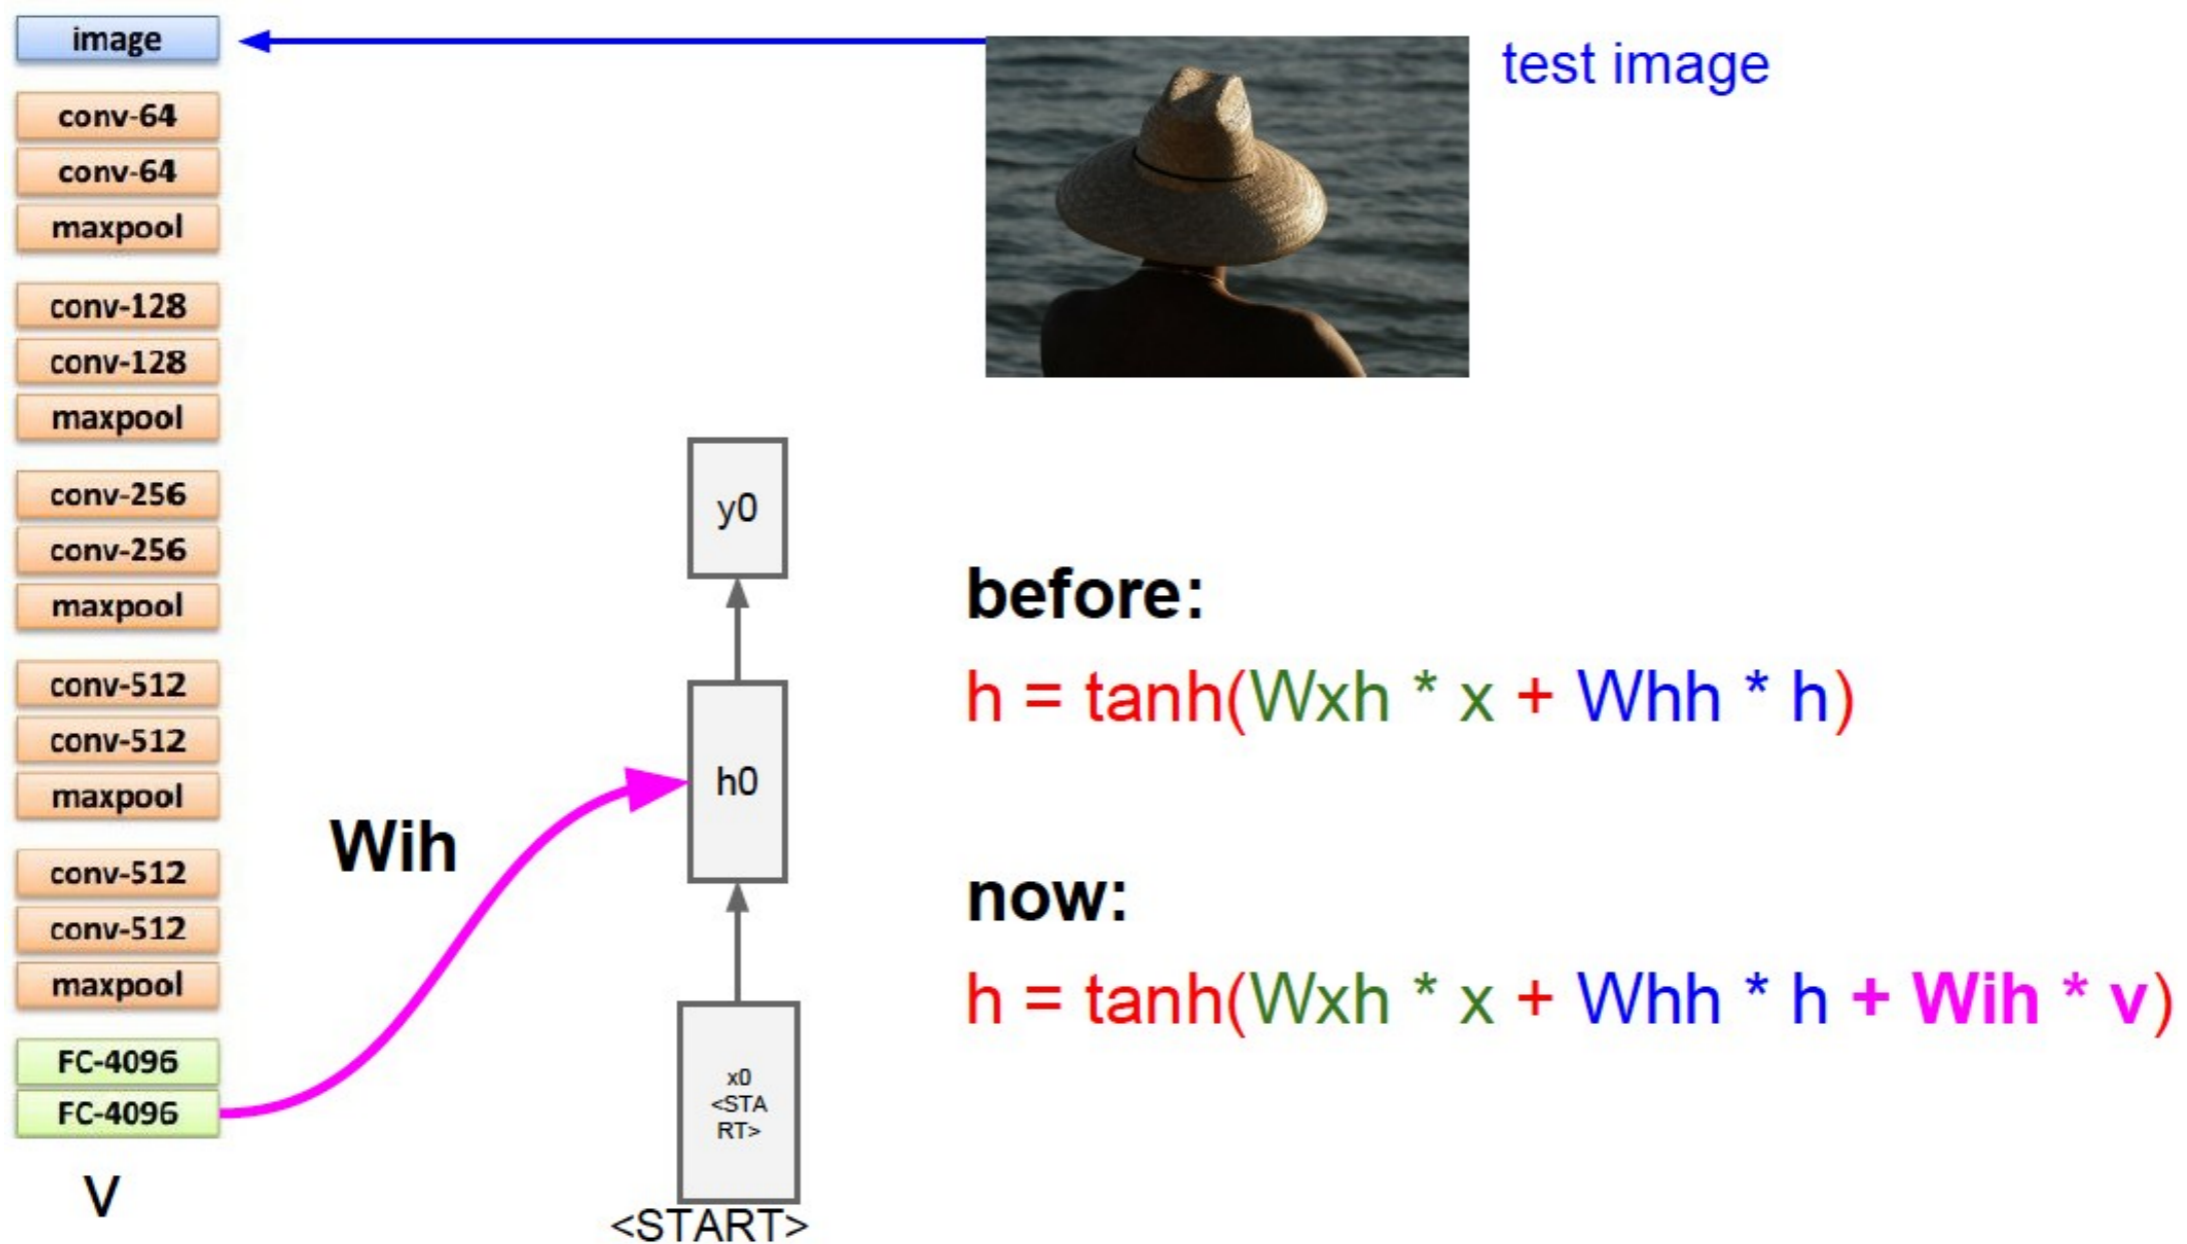
Recurrent Neural Networks

Example - Image Captioning



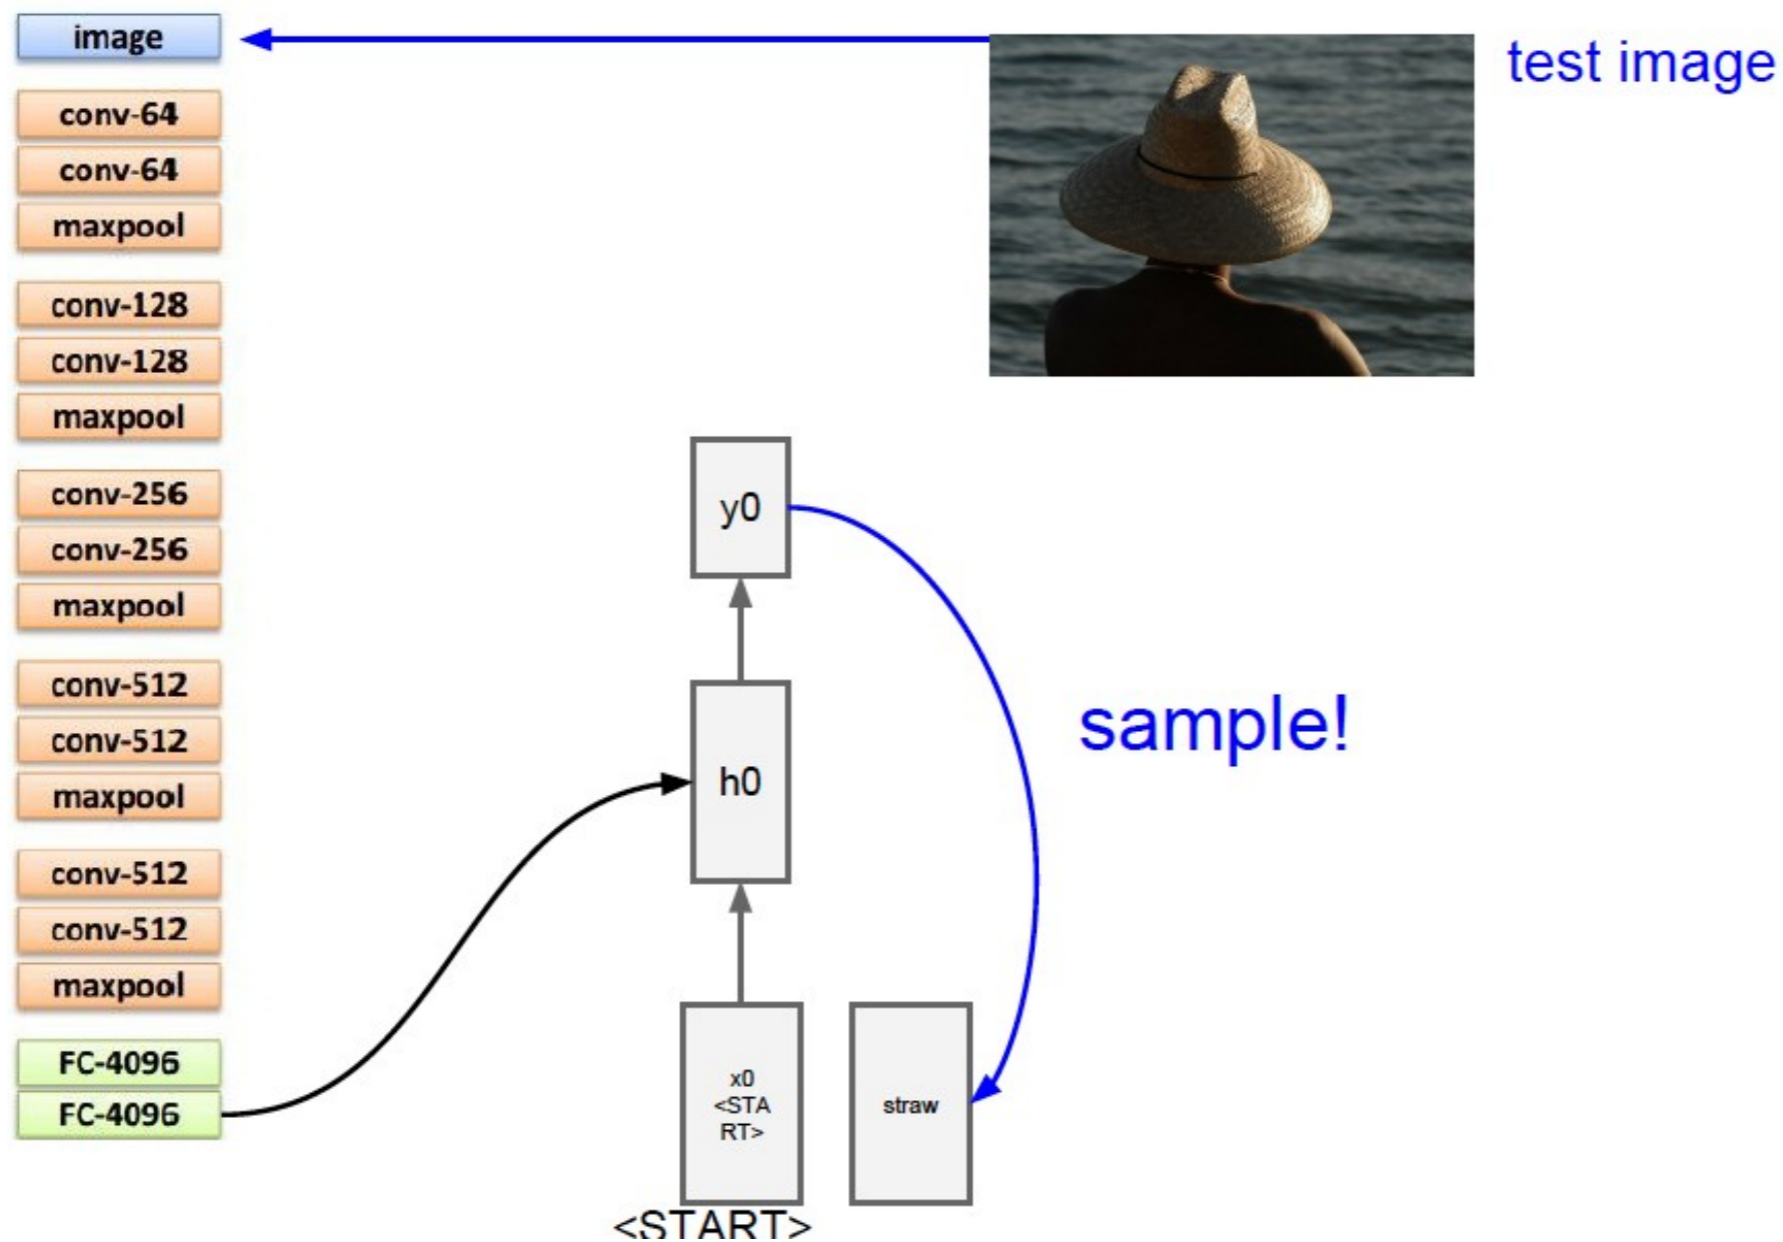
Recurrent Neural Networks

Example - Image Captioning



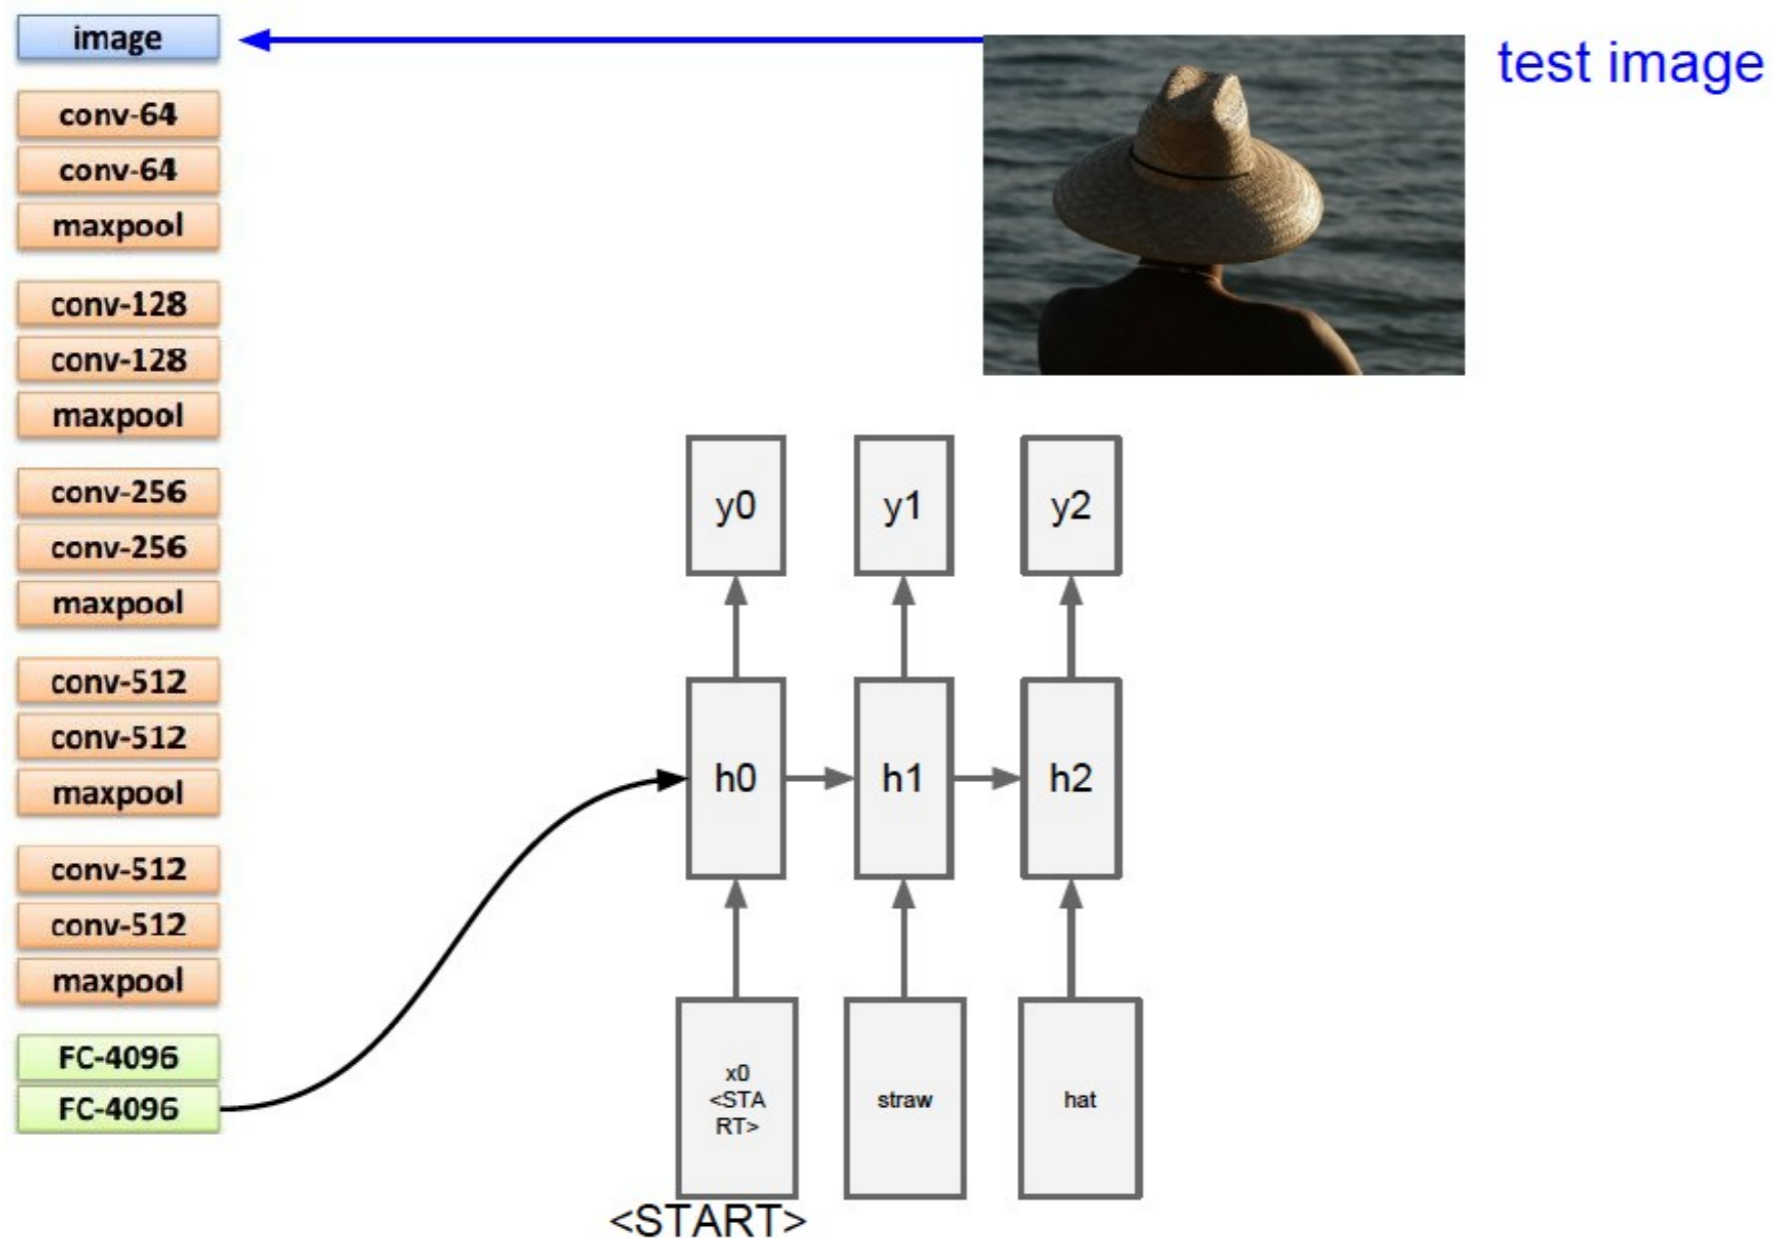
Recurrent Neural Networks

Example - Image Captioning



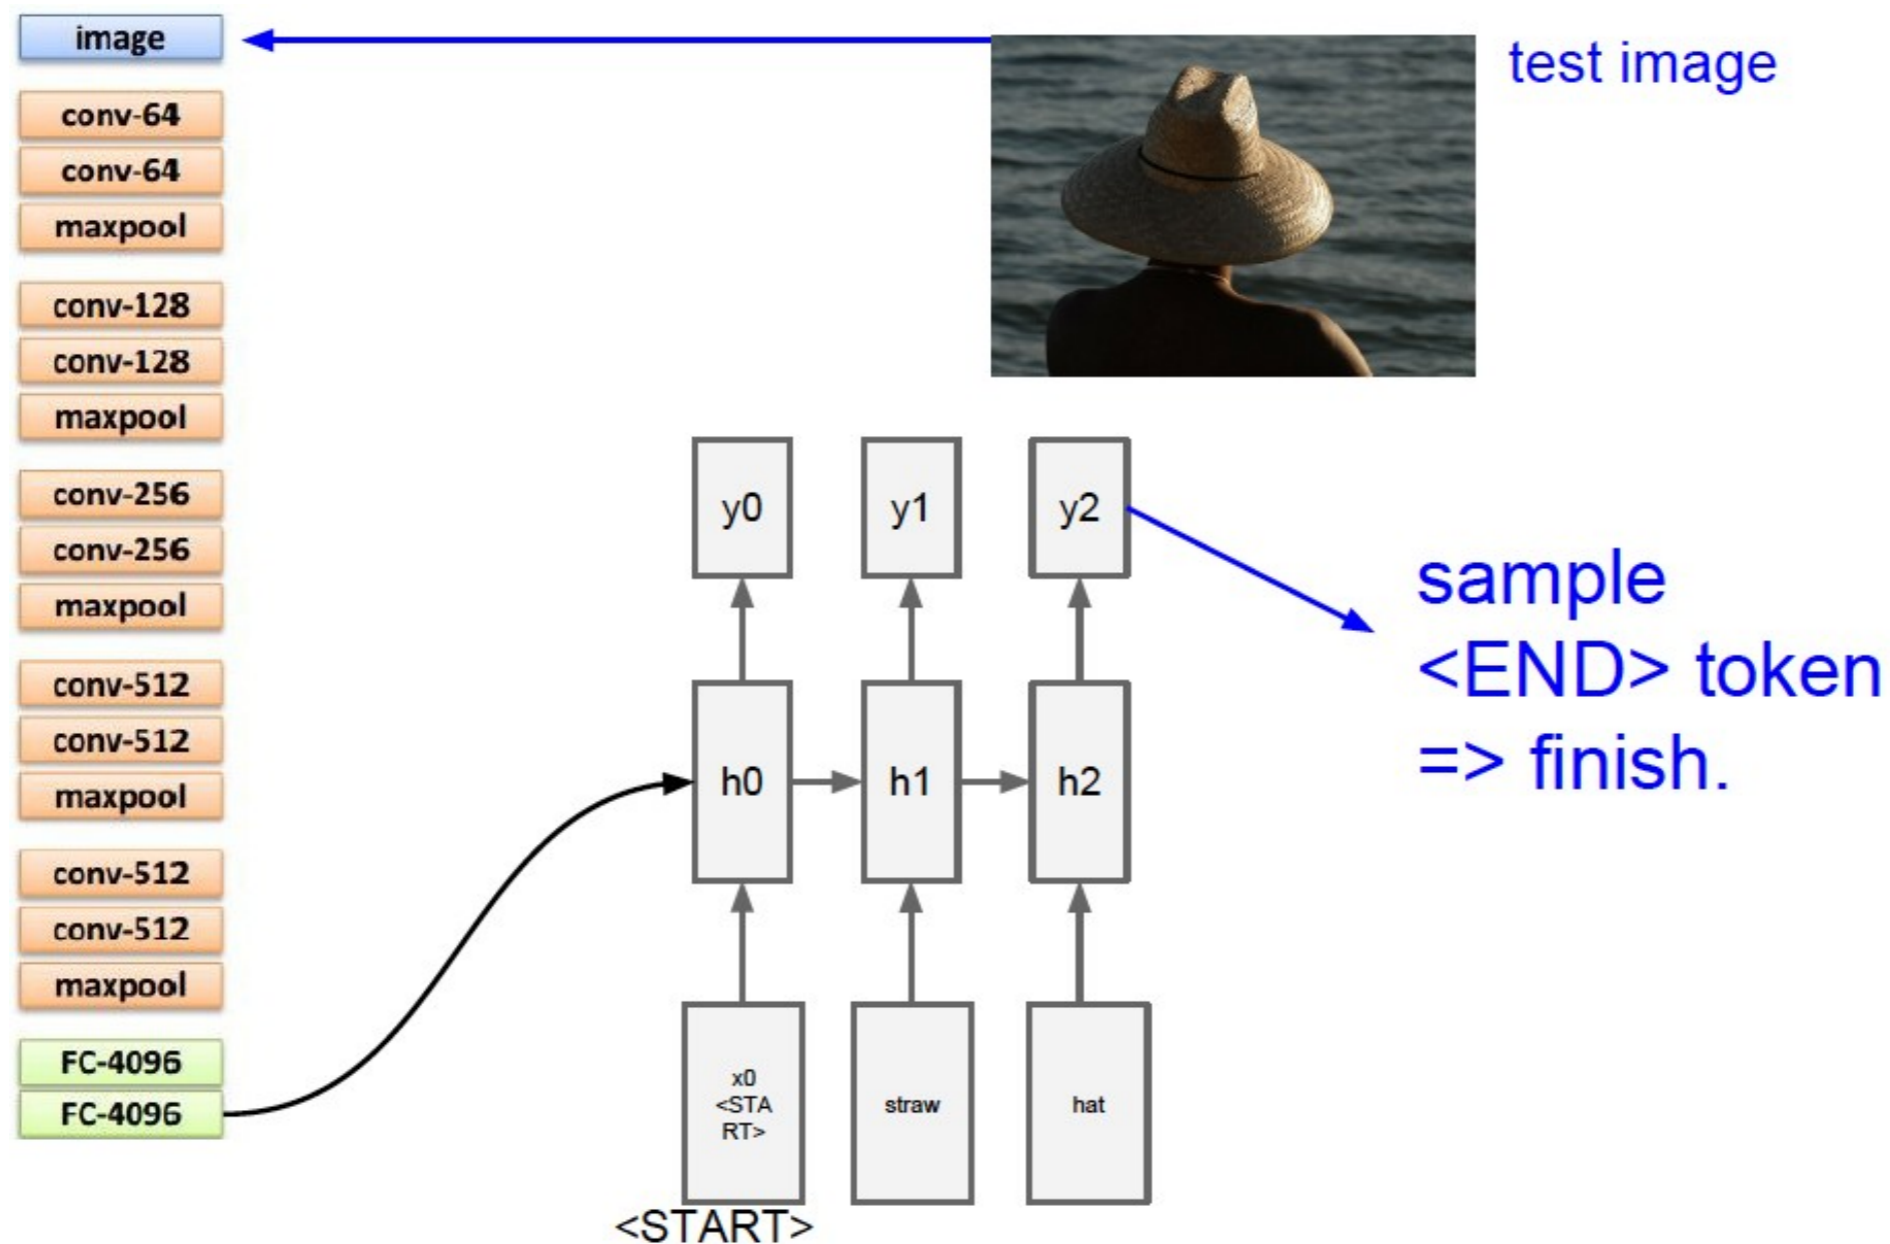
Recurrent Neural Networks

Example - Image Captioning



Recurrent Neural Networks

Example - Image Captioning



Recurrent Neural Networks

Example - Image Captioning



A cat sitting on a suitcase on the floor



A cat is sitting on a tree branch



A dog is running in the grass with a frisbee



A white teddy bear sitting in the grass



Two people walking on the beach with surfboards



A tennis player in action on the court



Two giraffes standing in a grassy field




A man riding a dirt bike on a dirt track

Recurrent Neural Networks

Example - Image Captioning (one to many)



A woman is holding a cat in her hand



A woman standing on a beach holding a surfboard



A bird is perched on a tree branch



A person holding a computer mouse on a desk



A man in a baseball uniform throwing a ball

Failure results

Recurrent Neural Networks

Deep RNN

- Multiple layer RNN


$$\mathbf{h}_t^l = \tanh \mathbf{W}^l [\mathbf{h}_t^{l-1} \ \mathbf{h}_{t-1}^l]^T$$

$$\mathbf{h} \in \mathcal{R}^n$$

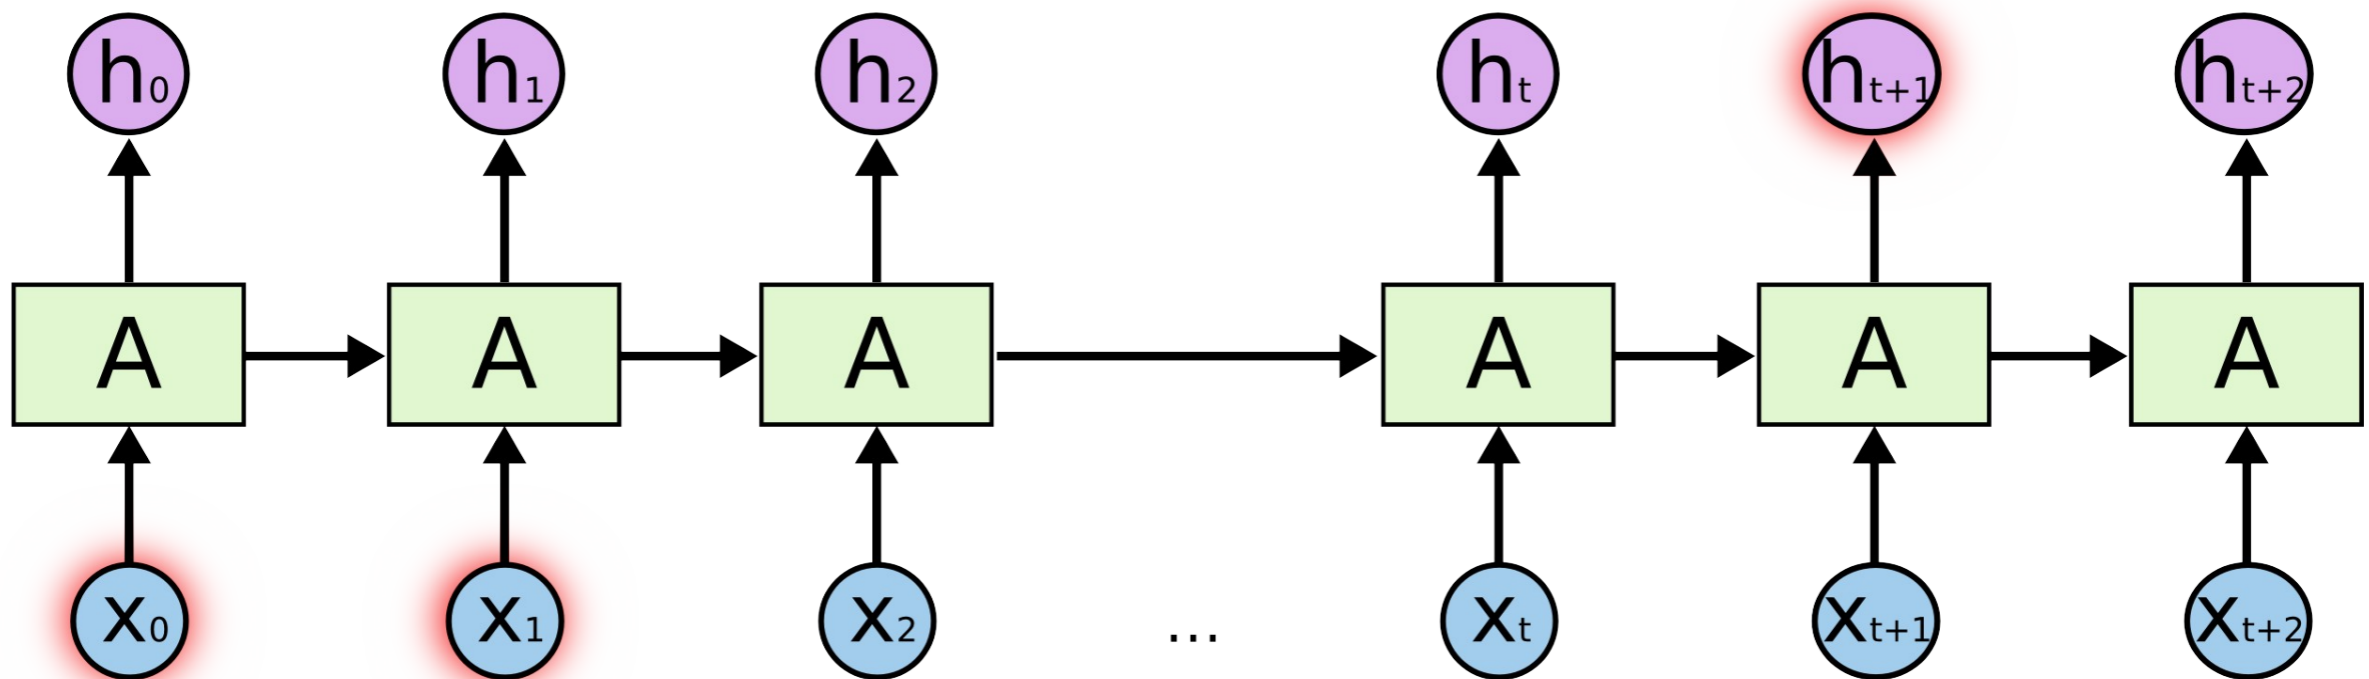
$$\mathbf{W}^l \in \mathcal{R}^{n \times 2n}$$

- Recall for one layer RNN:

$$\mathbf{h}_t = \tanh(\mathbf{W}_{hh} \mathbf{h}_{t-1} + \mathbf{W}_{xh} \mathbf{x}_t)$$




Recurrent Neural Networks



A -> RNN
h -> state
x -> input sequence

Long-term dependencies are hard to model!

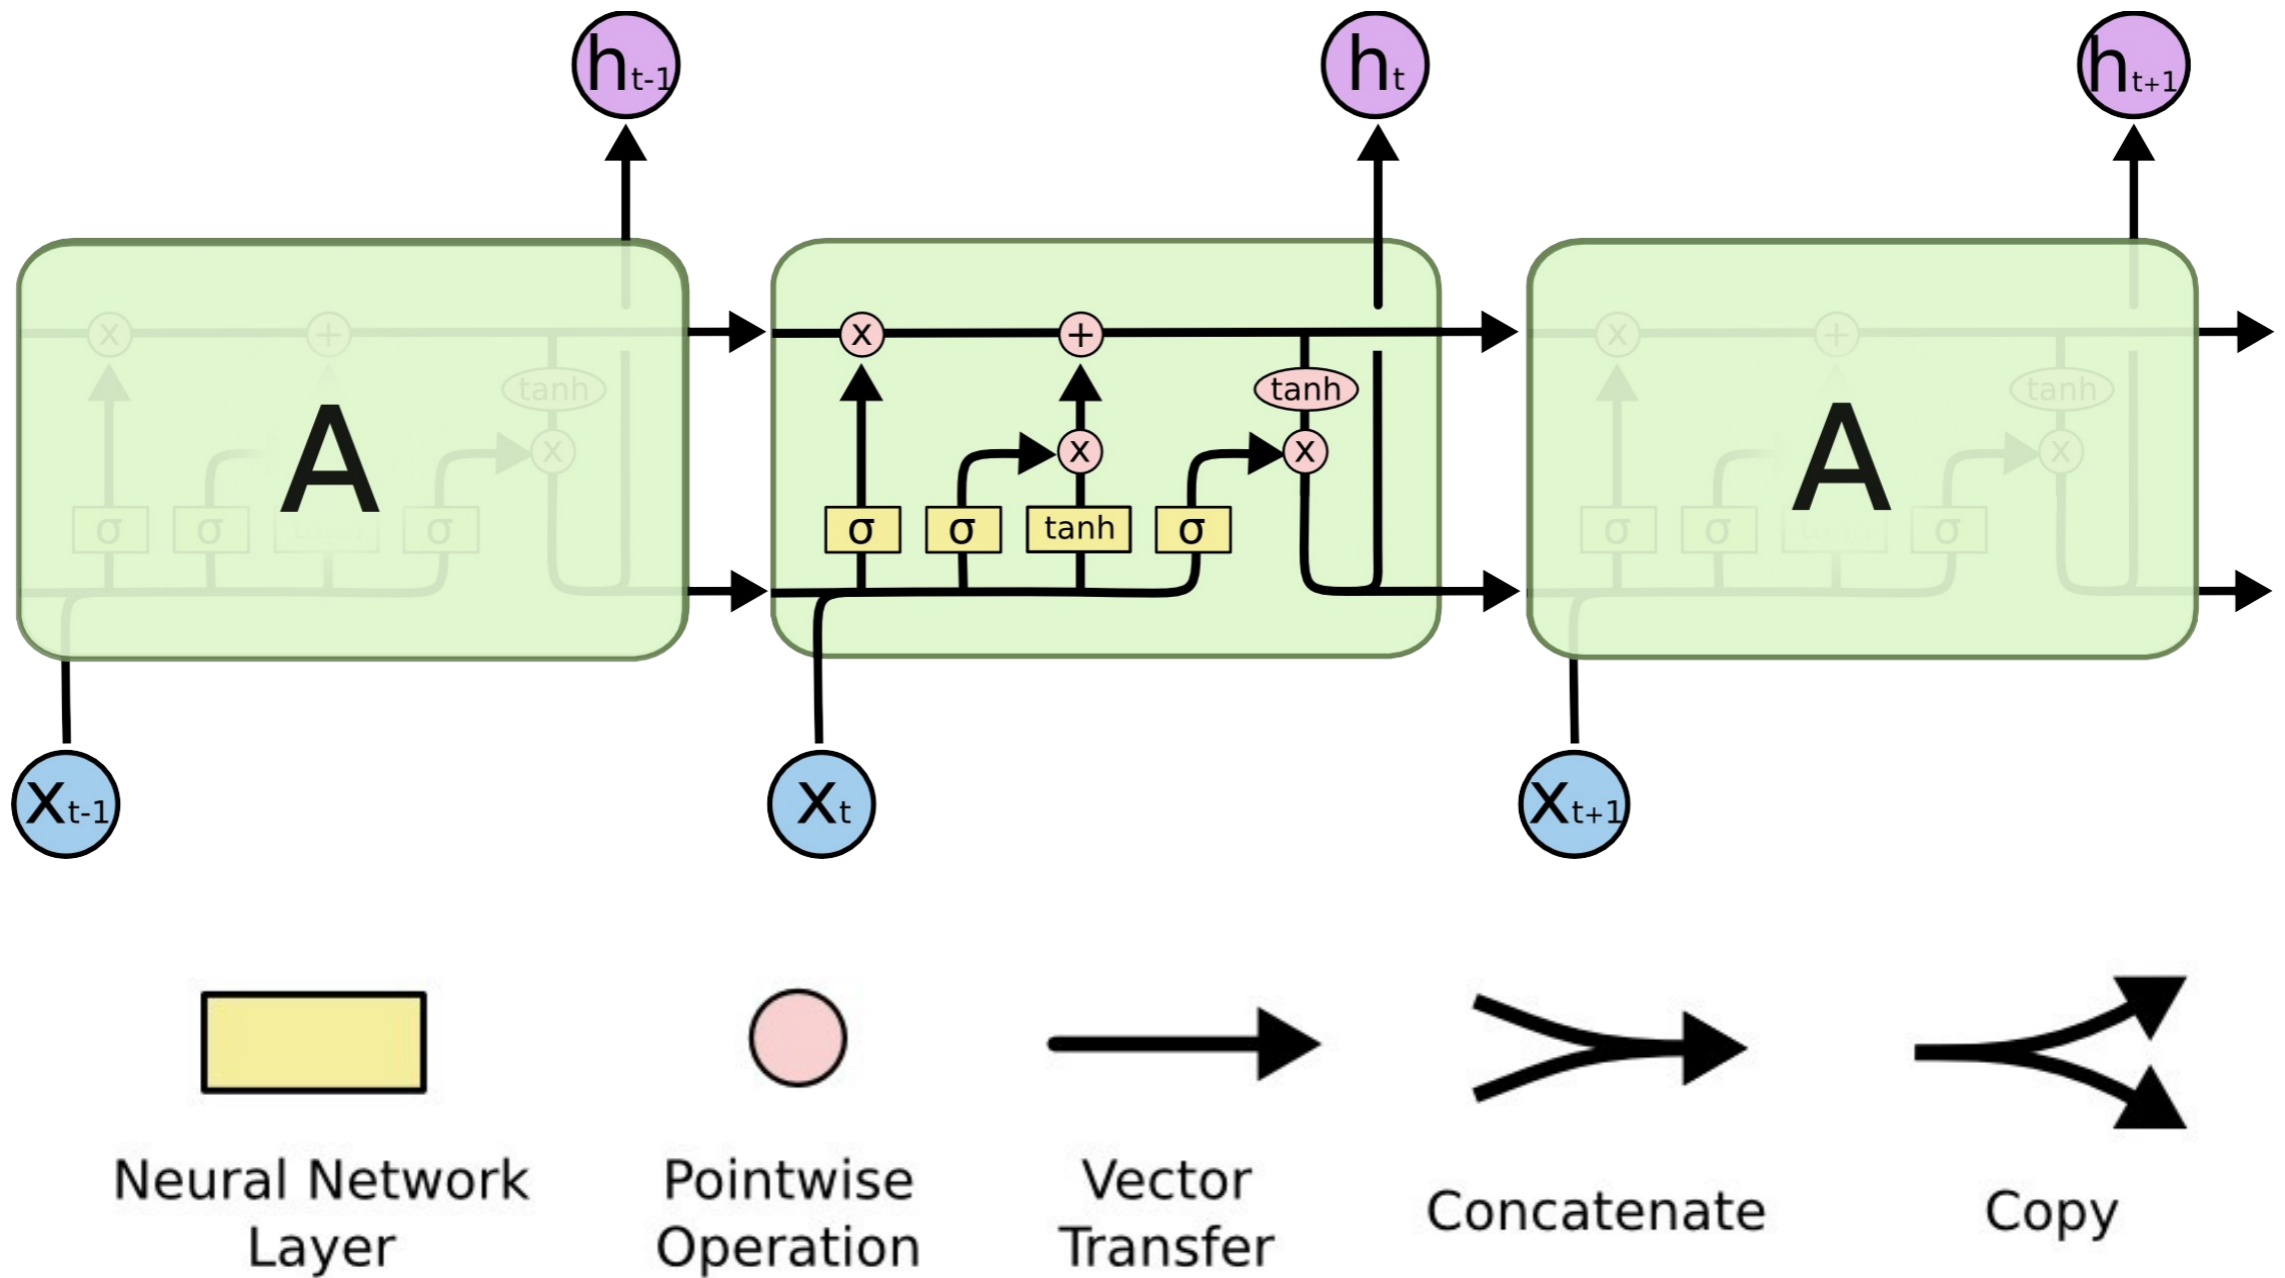
Recurrent Neural Networks



A -> RNN
h -> state
x -> input sequence

Plain RNN - what we have seen so far


From plain RNN to LSTM



LSTM: Long Short Term Memory Networks

LSTM

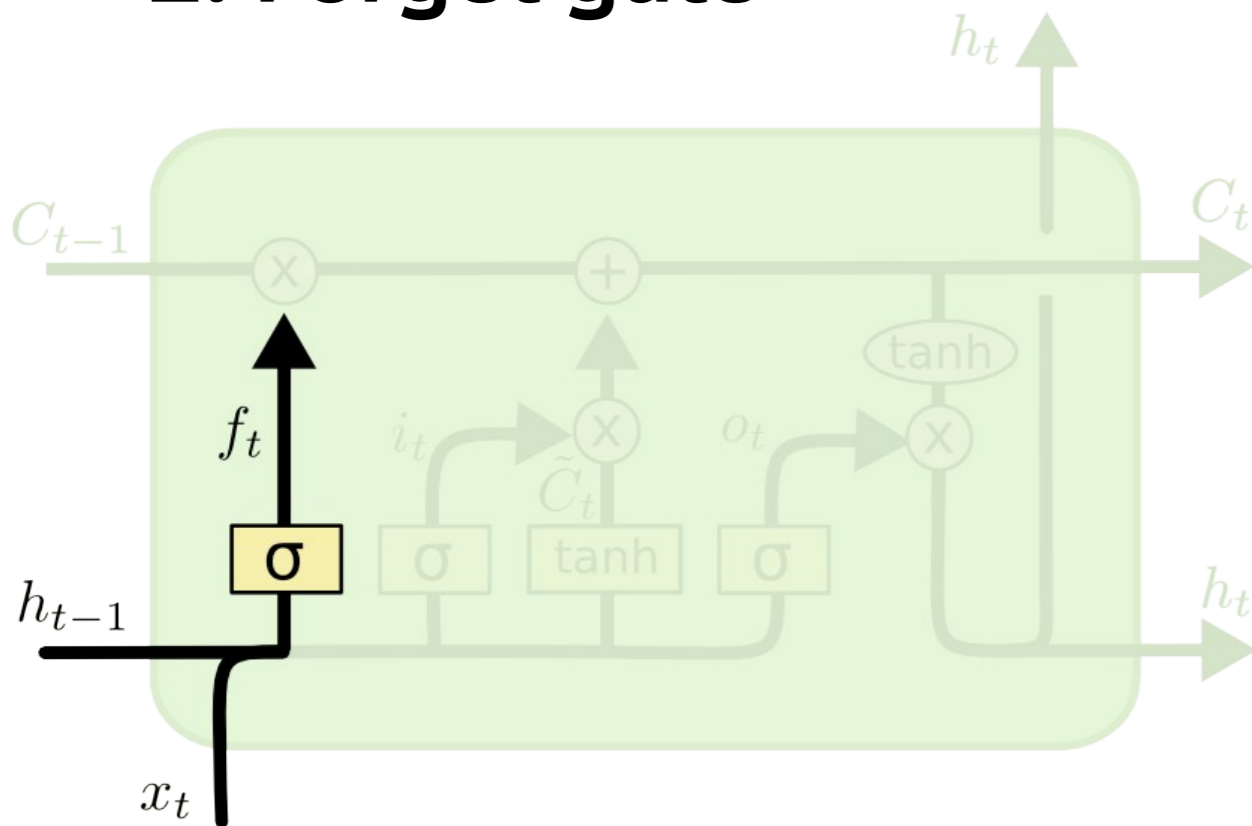
1. Cell State / Memory



The LSTM have the ability to remove or add information to the cell state, carefully regulated by structures called gates

LSTM

2. Forget gate




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

Should we continue to remember this “bit” of information or not?

The first step in our LSTM is to decide what information we’re going to throw away from the cell state. This decision is made by a sigmoid layer called the “forget gate layer.”

LSTM

3. Input gate



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


Should we update this “bit” of information or not? If so, with what?

The next step is to decide what new information we’re going to store in the cell state. This has two parts. First, a sigmoid layer called the “input gate layer” decides which values we’ll update. Next, a

tanh layer creates a vector of new candidate values, \tilde{C}_t , that could be added to the state.

LSTM

4. Memory Update




Forget that Memorize this

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

Decide what will be kept in the cell state/memory

LSTM

5. Output gate



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

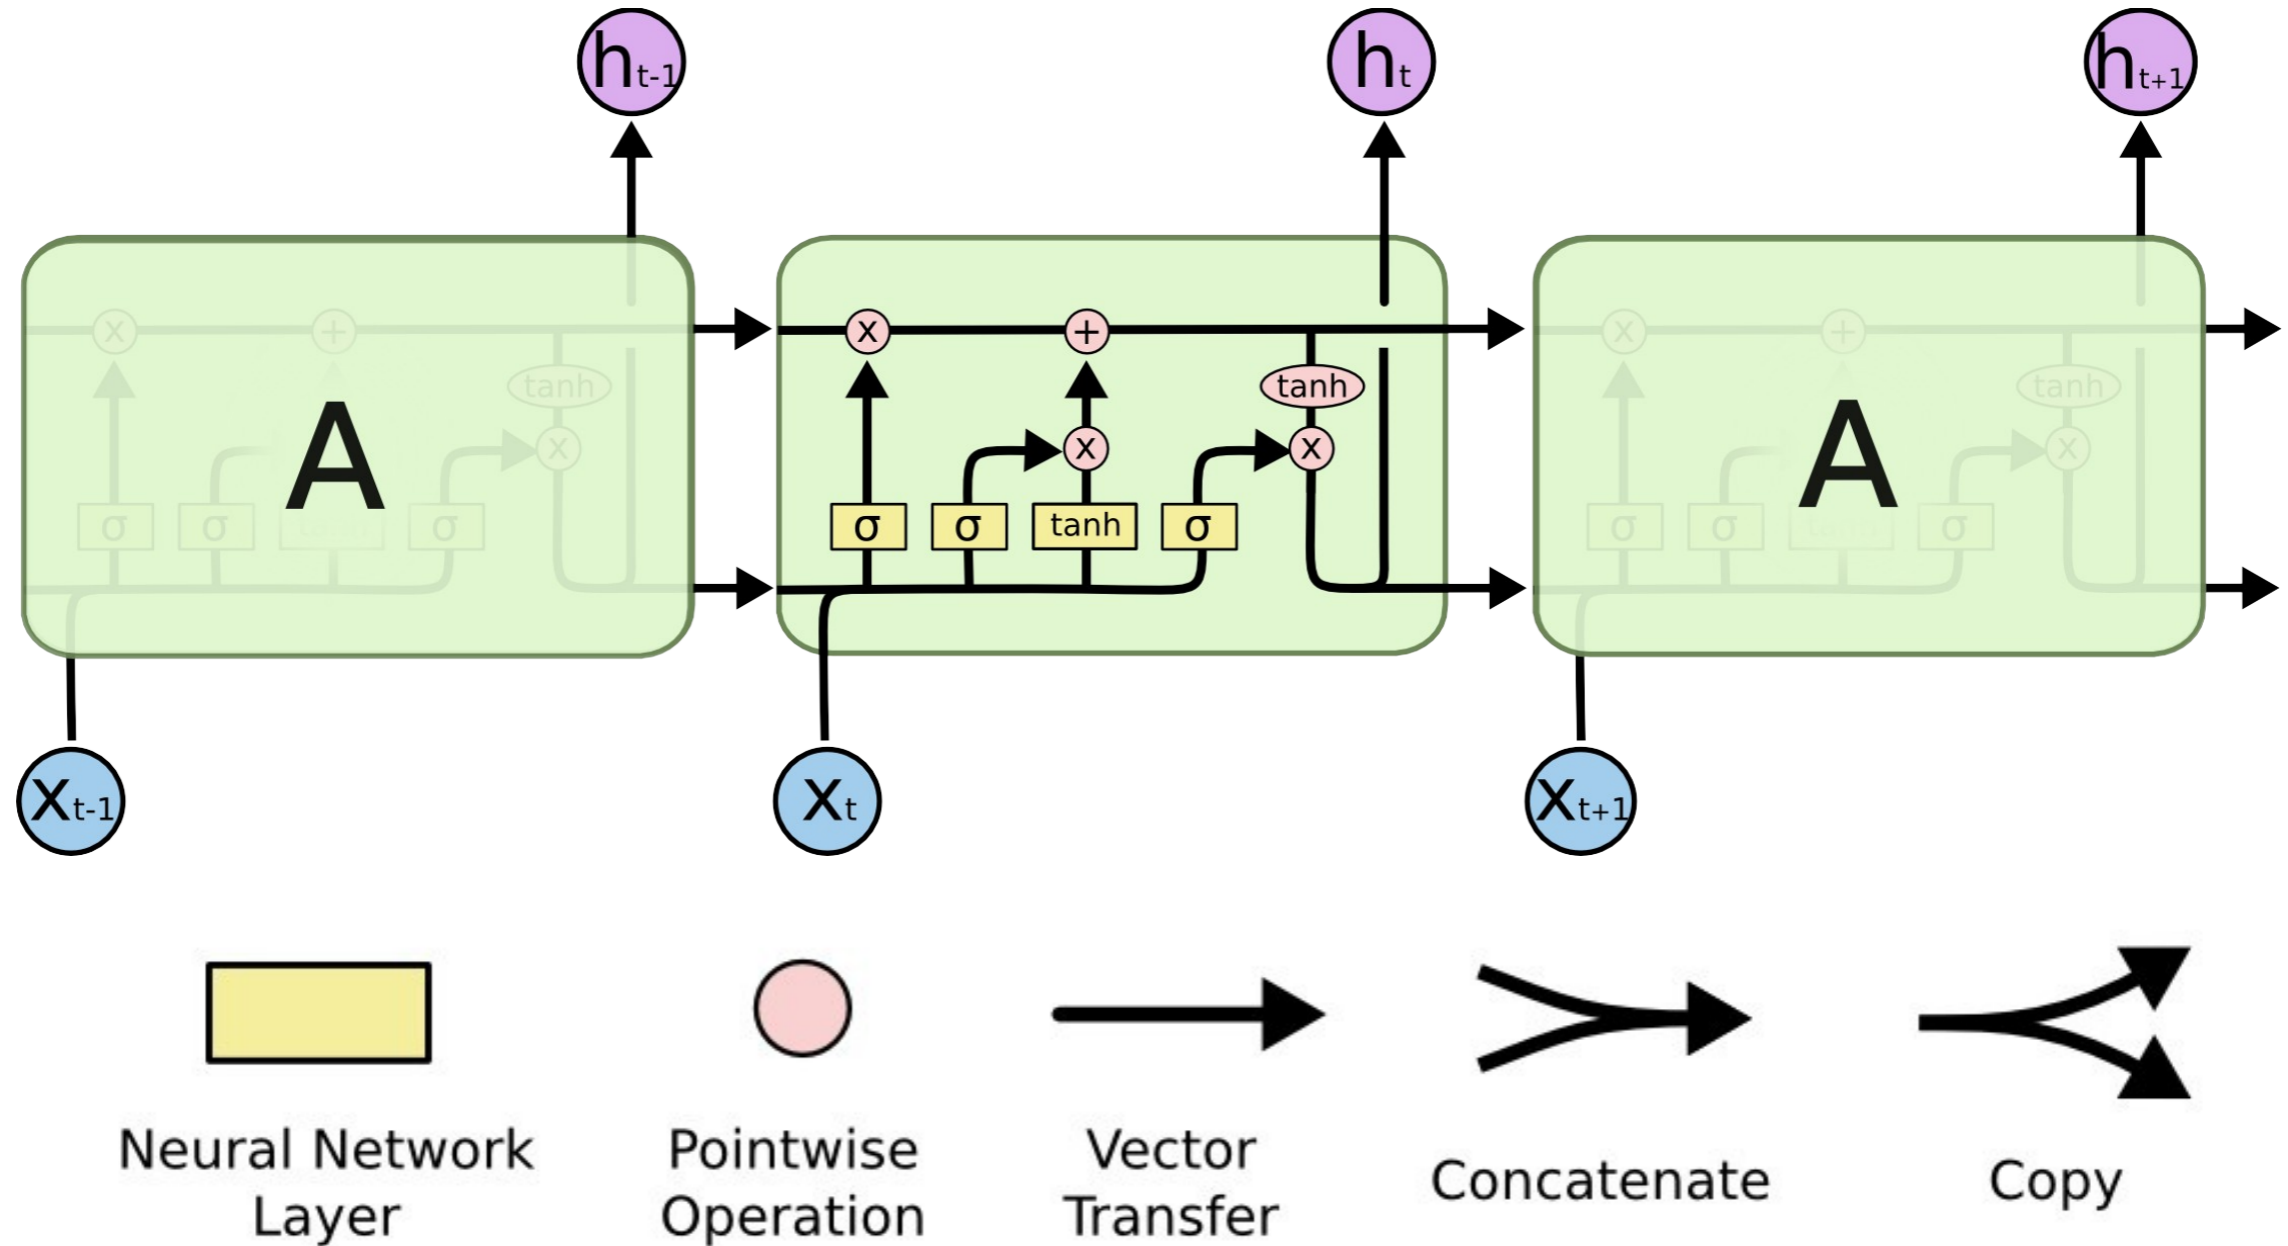
$$h_t = o_t * \tanh (C_t)$$

Should we output this “bit” of information?

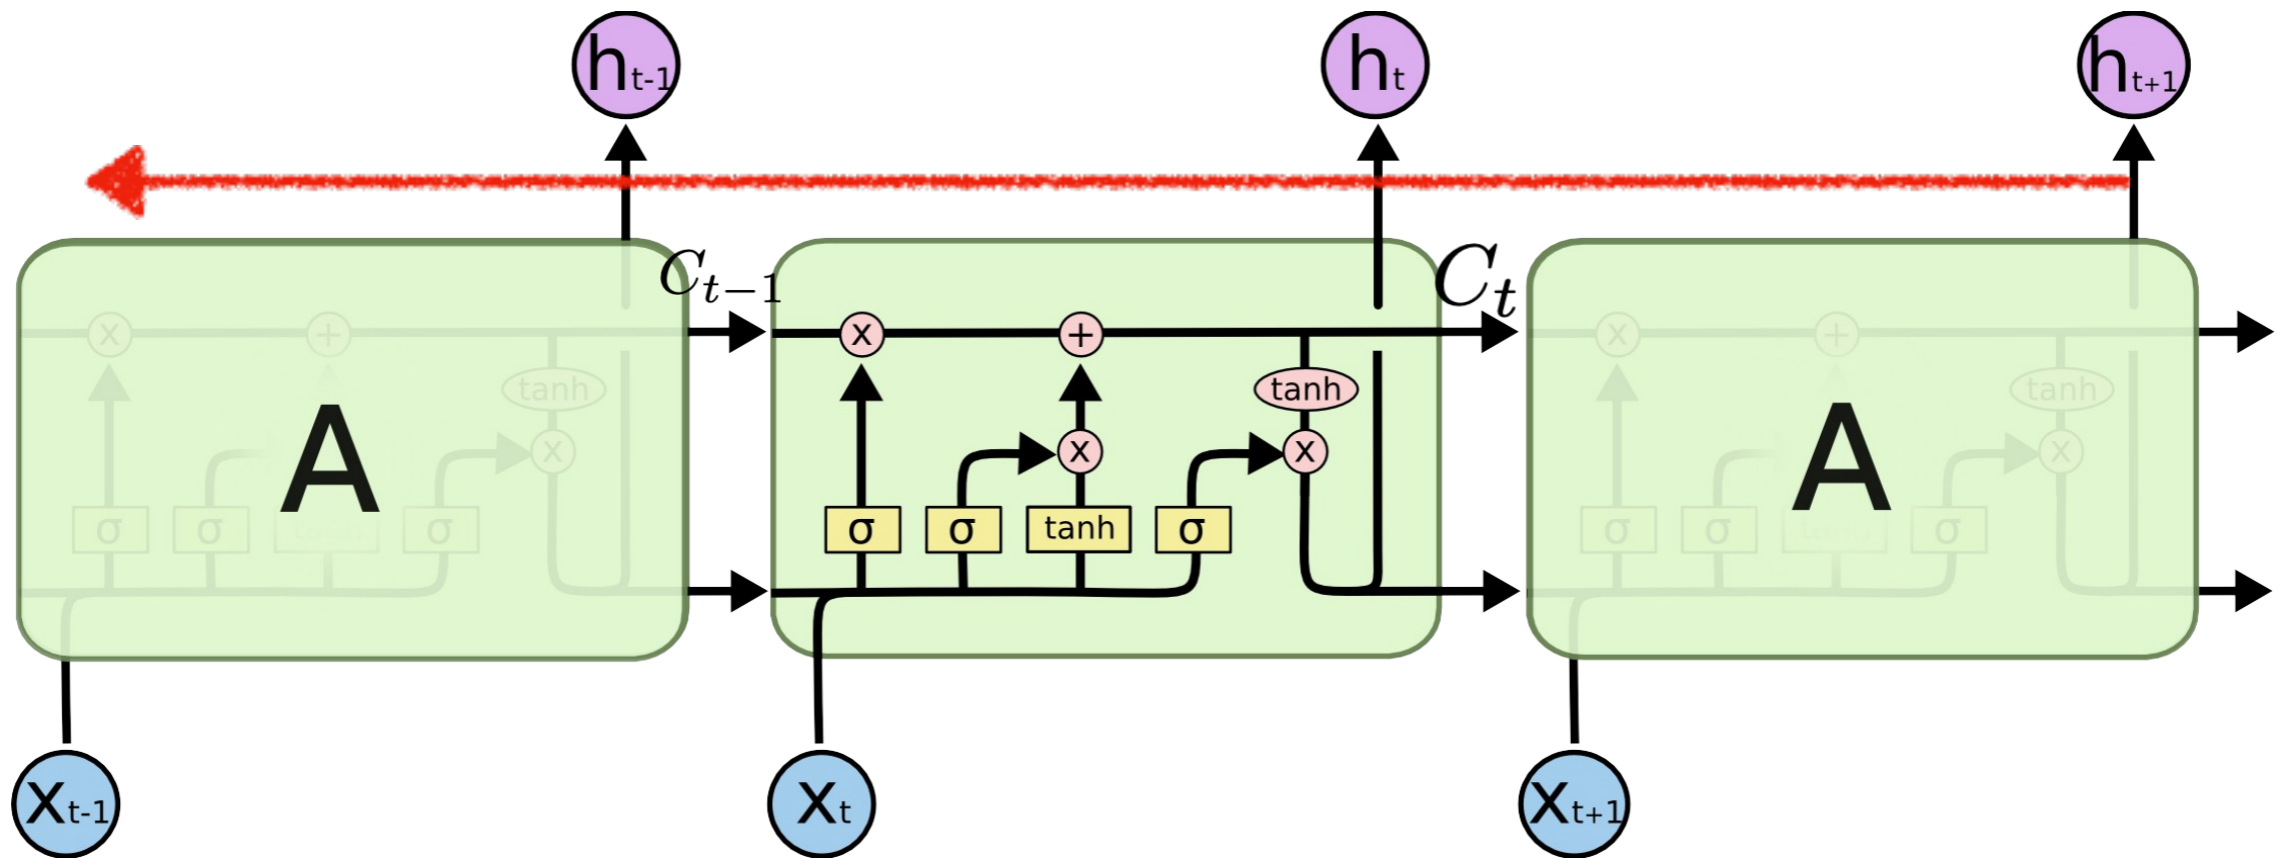
This output will be based on our cell state, but will be a filtered version. First, we run a sigmoid layer which decides what parts of the cell state we’re going to output. Then, we put the cell state through tanh (to push the values to be between -1 and 1) and multiply it by the output of the sigmoid gate, so that we only output the parts we decided to.

LSTM

Complete LSTM - A pretty sophisticated cell




LSTM



Back-propagation from C_t to C_{t-1} only has element wise multiplication by f_t and no by weight matrix
 There is an uninterrupted gradient flow

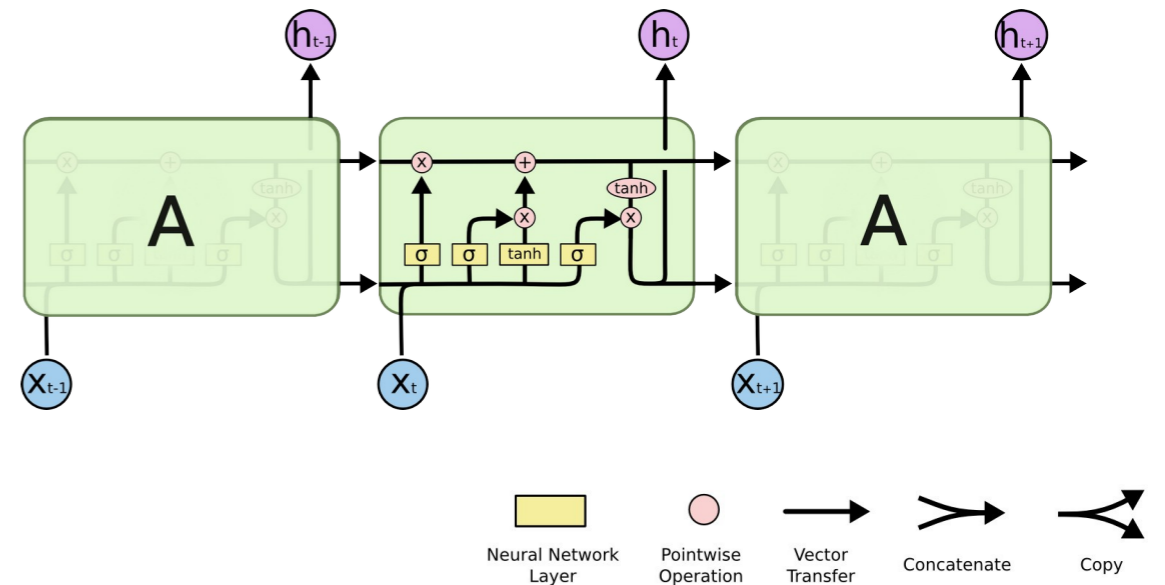
Plain RNN vs LSTM

Plain RNN



$$h_t = \tanh(\mathbf{W} \begin{bmatrix} h_{t-1} \\ x_t \end{bmatrix})$$

LSTM




$$\begin{bmatrix} i_t \\ f_t \\ o_t \\ \tilde{C}_t \end{bmatrix} = \begin{bmatrix} \sigma \\ \sigma \\ \sigma \\ \tanh \end{bmatrix} \mathbf{W} \begin{bmatrix} h_{t-1} \\ x_t \end{bmatrix} + b$$

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

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

Plain RNN vs LSTM

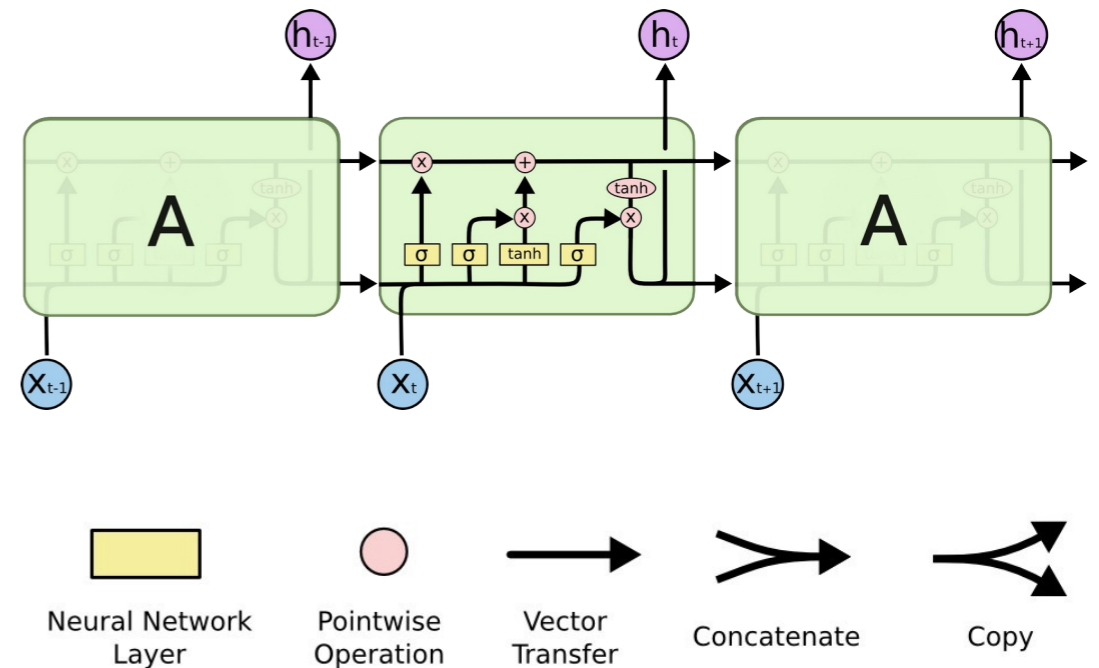
Plain RNN



Backward flow of gradient in RNN can:

- explode -> gradient clipping
- vanish -> use LSTM


LSTM



Backward flow of gradient in LSTM:

- their additive interactions improve the gradient flow

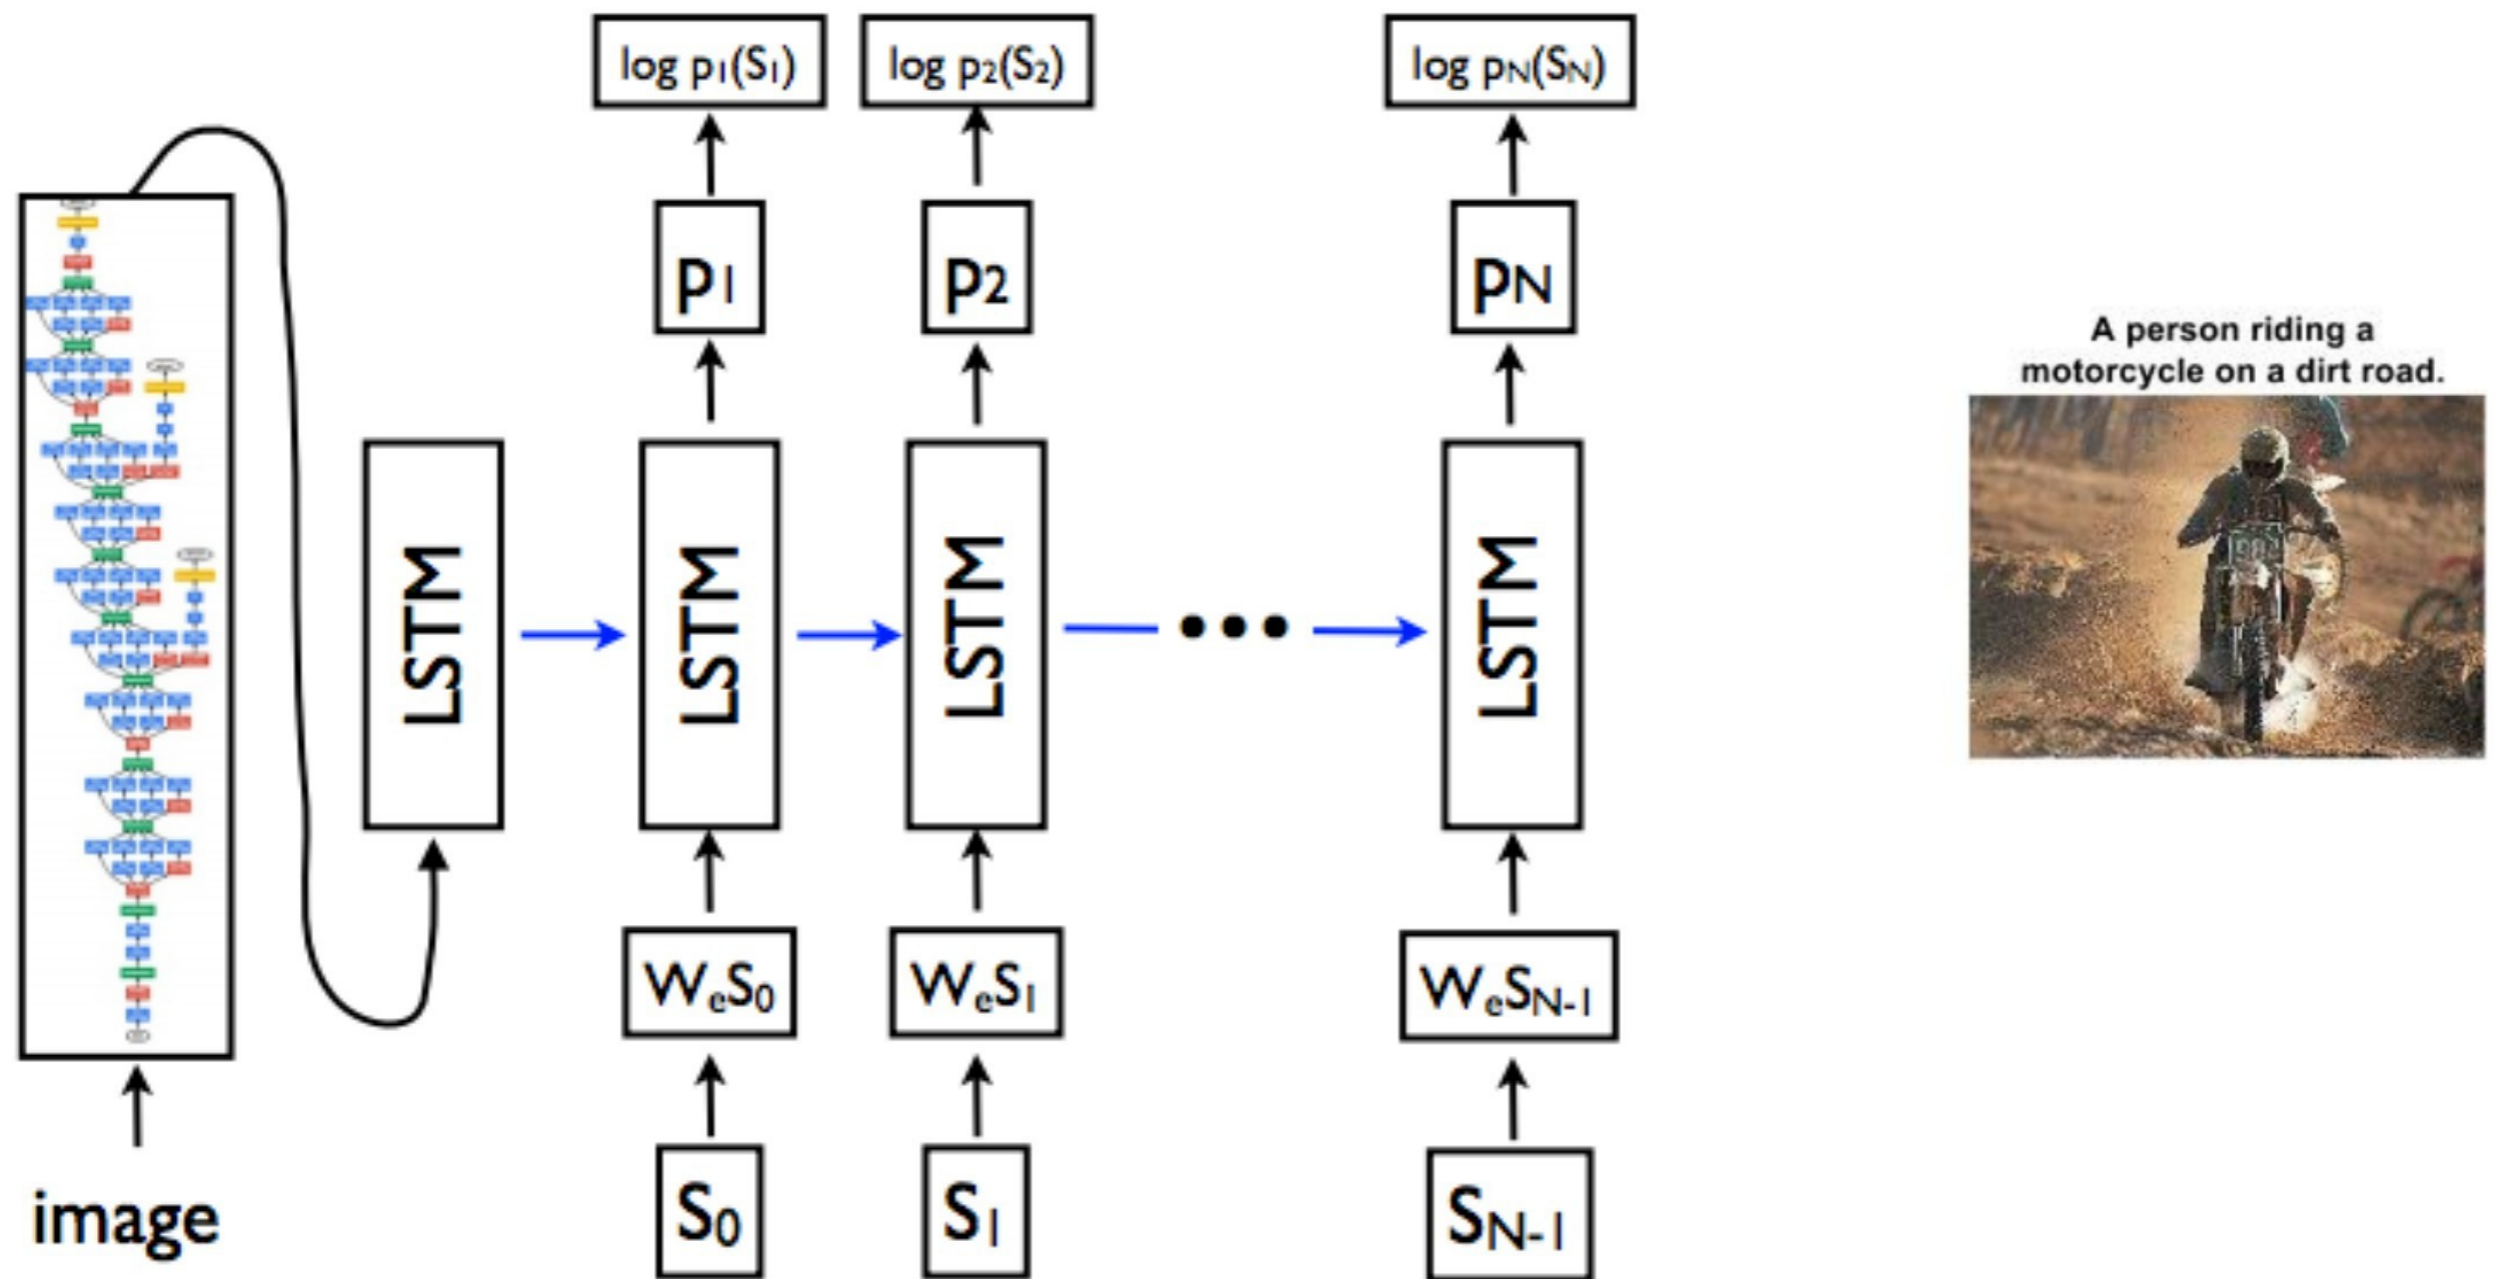
Show and Tell: A neural Image Caption Generator



A group of people shopping at an outdoor market.

There are many vegetables at the fruit stand.

Show and Tell: A neural Image Caption Generator



Show and Tell: A neural Image Caption Generator

A person riding a motorcycle on a dirt road.




Two dogs play in the grass.



A skateboarder does a trick on a ramp.



A dog is jumping to catch a frisbee.



A group of young people playing a game of frisbee.



Two hockey players are fighting over the puck.



A little girl in a pink hat is blowing bubbles.



A refrigerator filled with lots of food and drinks.



A herd of elephants walking across a dry grass field.



A close up of a cat laying on a couch.



A red motorcycle parked on the side of the road.



A yellow school bus parked in a parking lot.



Describes without errors

Describes with minor errors

Somewhat related to the image

Unrelated to the image

Show and Tell: A neural Image Caption Generator

A person riding a motorcycle on a dirt road.




Two dogs play in the grass.



A skateboarder does a trick on a ramp.



A dog is jumping to catch a frisbee.




A group of young people playing a game of frisbee.



Two hockey players are fighting over the puck.



A little girl in a pink hat is blowing bubbles.



A refrigerator filled with lots of food and drinks.



A herd of elephants walking across a dry grass field.



A close up of a cat laying on a couch.



A red motorcycle parked on the side of the road.



A yellow school bus parked in a parking lot.



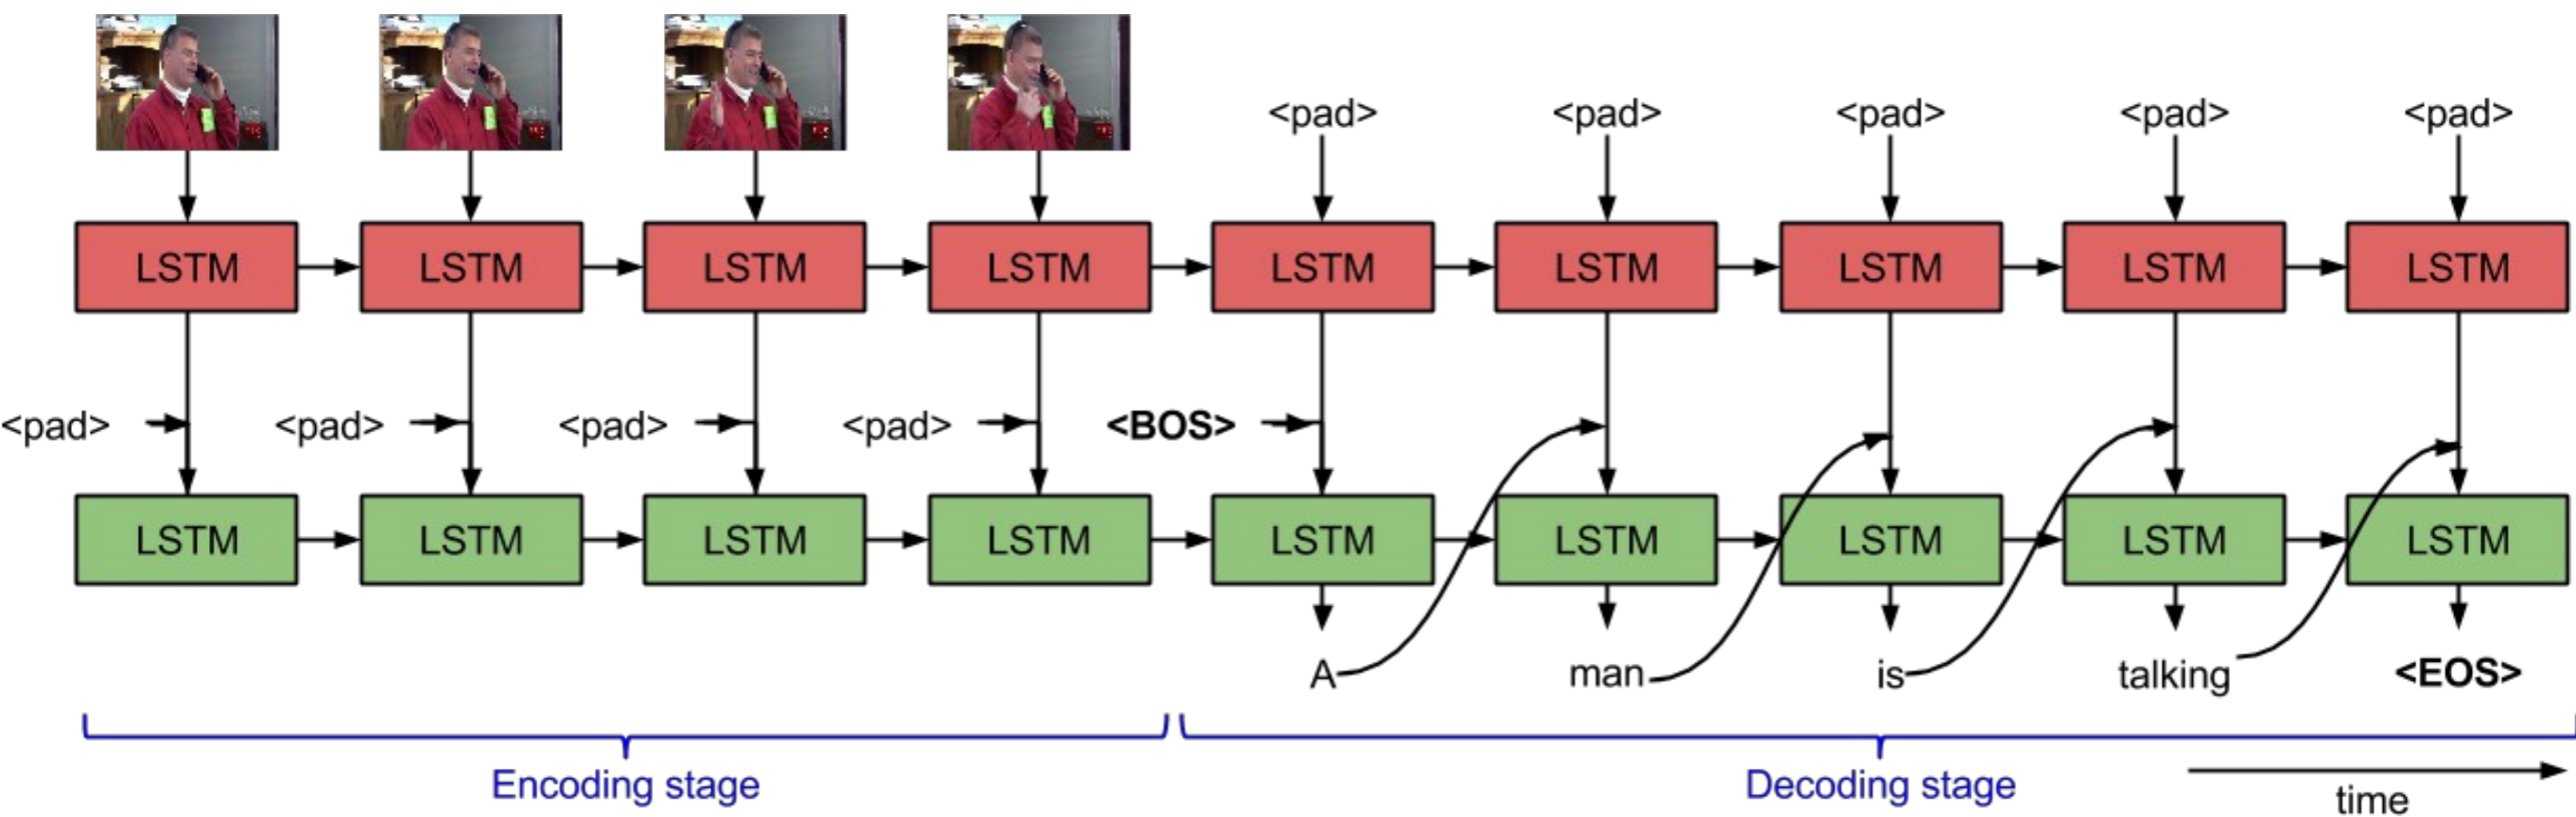
Describes without errors

Describes with minor errors

Somewhat related to the image

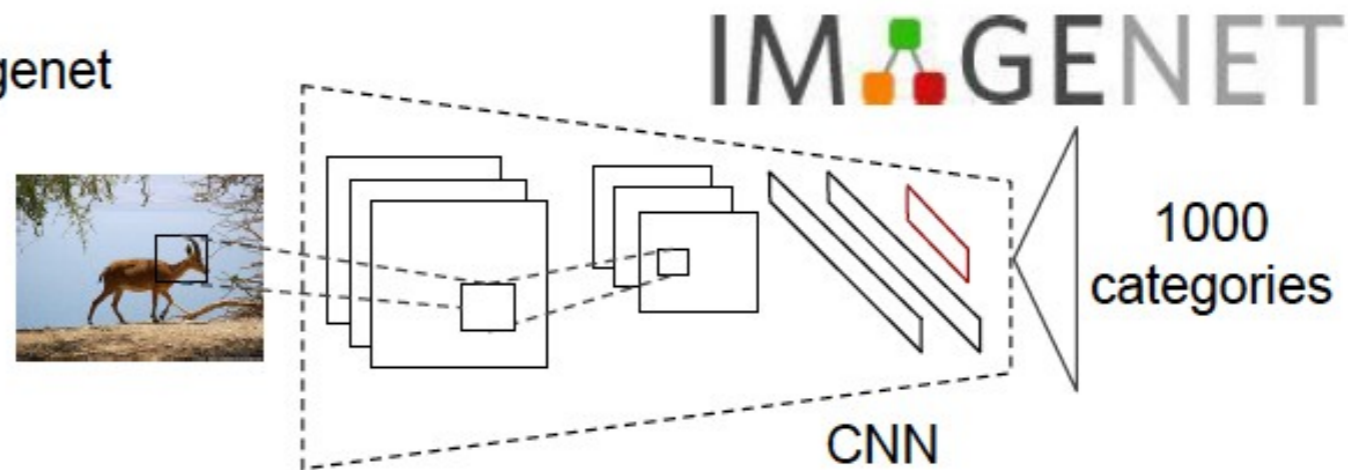
Unrelated to the image

Sequence to Sequence Video to Text

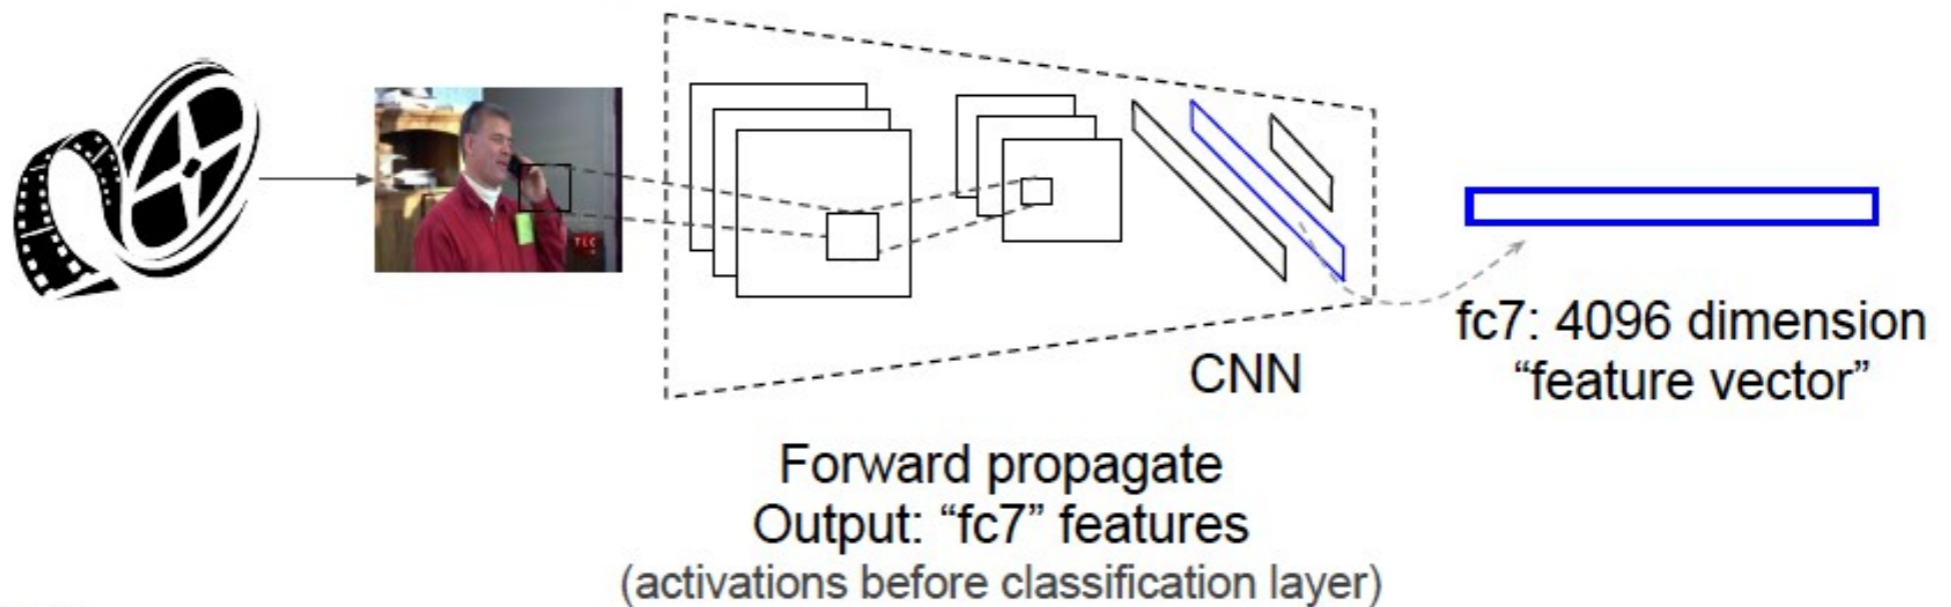


Sequence to Sequence Video to Text

1. Train on Imagenet




2. Take activations from layer before classification



Frames: RGB

Sequence to Sequence Video to Text

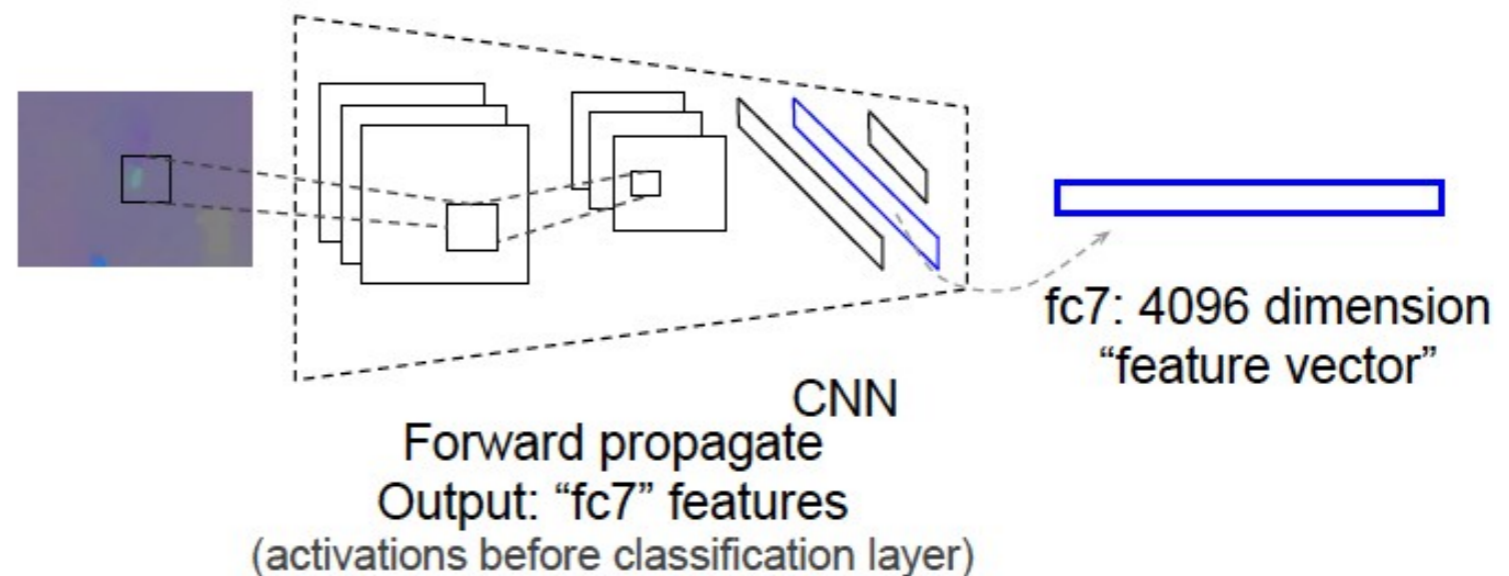
1. Train CNN on Activity classes



2. Use optical flow to extract flow images.



3. Take activations from layer before classification



Frames: Flow

Sequence to Sequence Video to Text

Movie Corpus - DVS



CC: Queen: "Which estate?"

DVS: Looking troubled, the Queen descends the stairs.

The Queen rushes into the courtyard. She then puts a head scarf on ...

...and gets into the driver's side of a nearby Land Rover.

The Land Rover pulls away.

Three bodyguards quickly jump into a nearby car and follow her.

Processed:
Looking troubled, someone descends the stairs.

Someone rushes into the courtyard. She then puts a head scarf on ...

Sequence to Sequence Video to Text

Summary and examples of results:

<https://youtu.be/-xNI7e7YgDk>

Summary of Today's class

RN





N

one to many

many to one

many to many

many to many



Input: No sequence

Input: Sequence

Input: Sequence

Output: Sequence

Output: No sequence

Output: Sequence

Example:

- Image captioning (image -> words)

Example:

- sentence classification
- sentiment classification (words seq.->sentiment)

Example:


- machine translation, (words seq-> words seq)

Example:

- video captioning

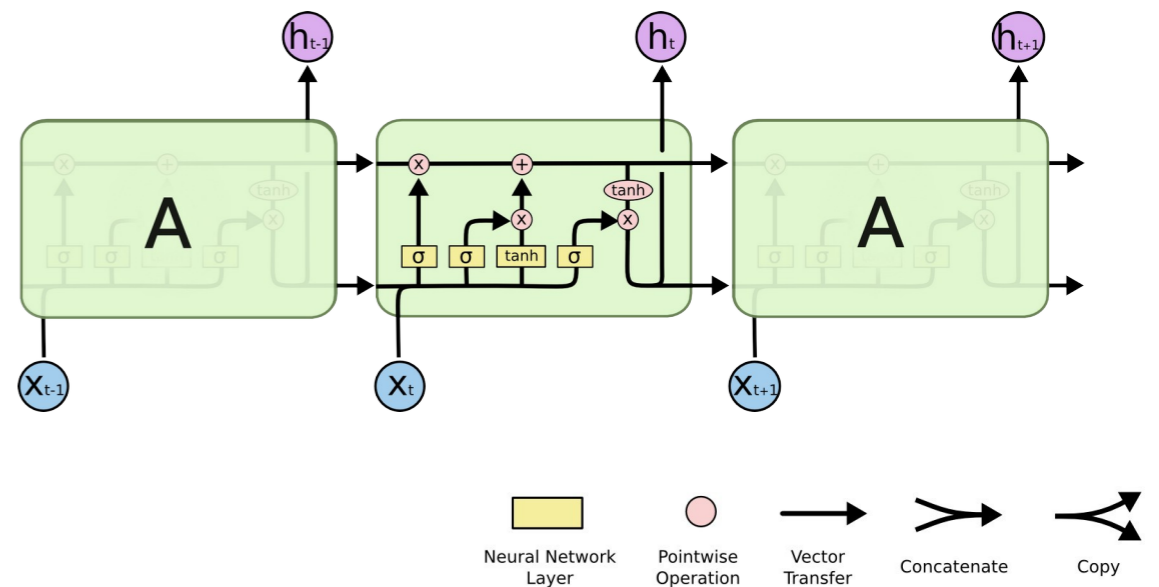
RNN vs LSTM

Plain RNN



$$h_t = \tanh(\mathbf{W} \begin{bmatrix} h_{t-1} \\ x_t \end{bmatrix})$$

LSTM



$$\begin{bmatrix} i_t \\ f_t \\ o_t \\ \tilde{C}_t \end{bmatrix} = \begin{bmatrix} \sigma \\ \sigma \\ \sigma \\ \tanh \end{bmatrix} \mathbf{W} \begin{bmatrix} h_{t-1} \\ x_t \end{bmatrix} + b$$

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

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