مرکز آموزش‌های الکترونیکی دانشگاه شهید بهشتی

شبکه‌های عصبی مصنوعی

Artificial Neural Networks

Ali Katanforoush*

** Department of Data and Computer Sciences, Shahid Beheshti University*

Recurrent Neural Network (RNN)

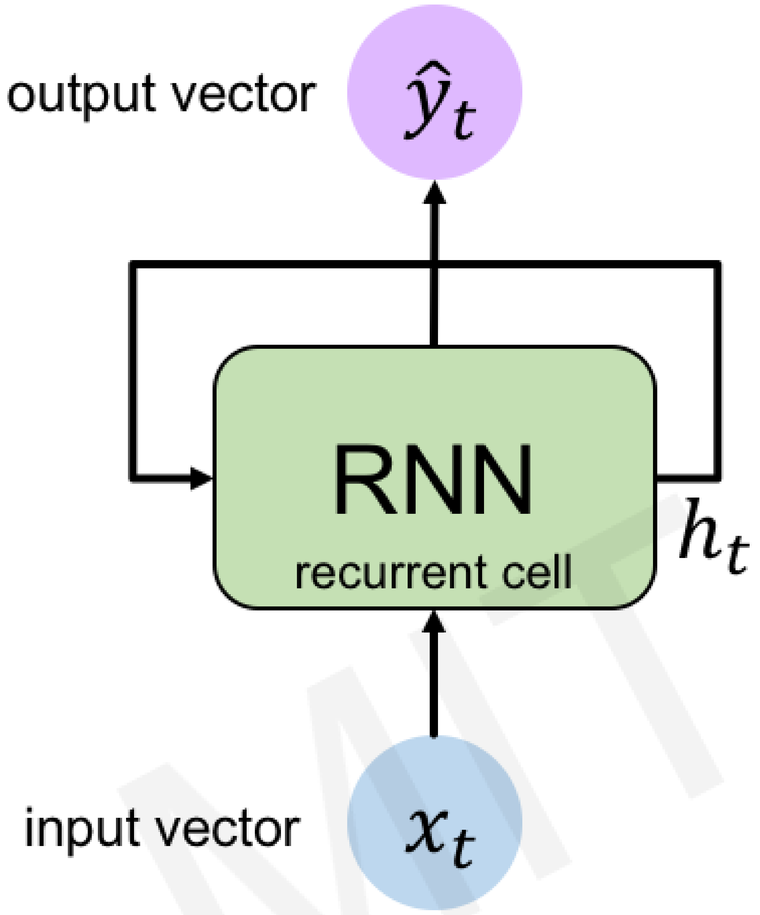

Apply a **recurrence relation** at every time step to process a sequence:

$$\boxed{h_t} = \boxed{f_W}\left(\boxed{h_{t-1}}, \boxed{x_t}\right)$$

cell state

function parameterized by W

old state

input vector at time step *t*

Note: the same function and set of parameters are used at every time step

RNN State Update and Output

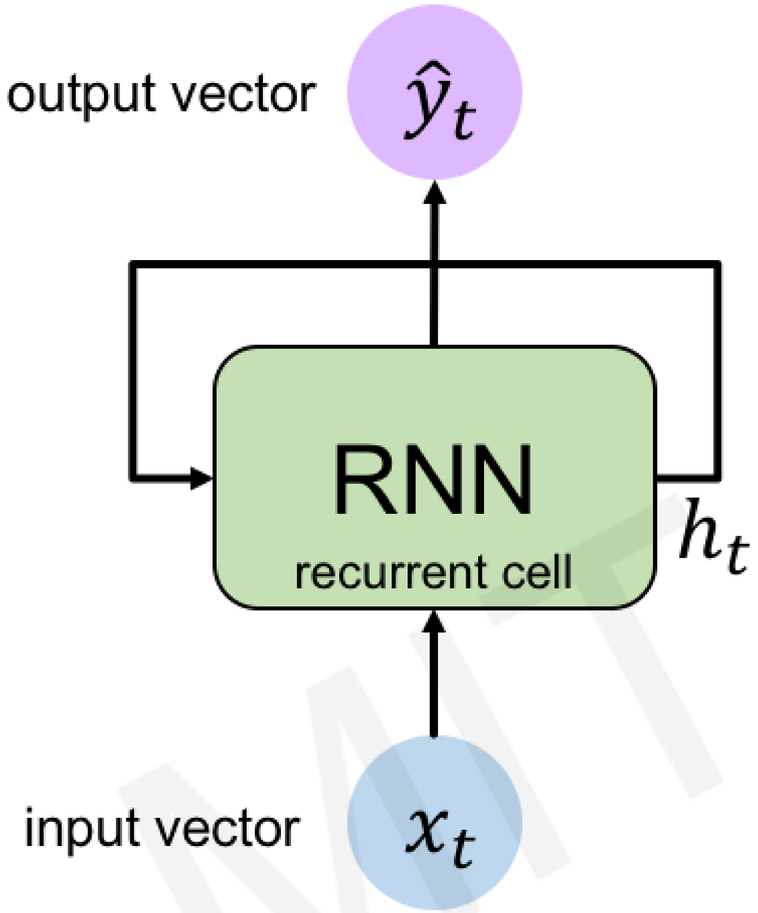

Output Vector

$$\hat{y}_t = \boldsymbol{W}_{\boldsymbol{hy}}^{\boldsymbol{T}} h_t$$

Update Hidden State

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

Input Vector

x_t

RNNs: Computational Graph Across Time

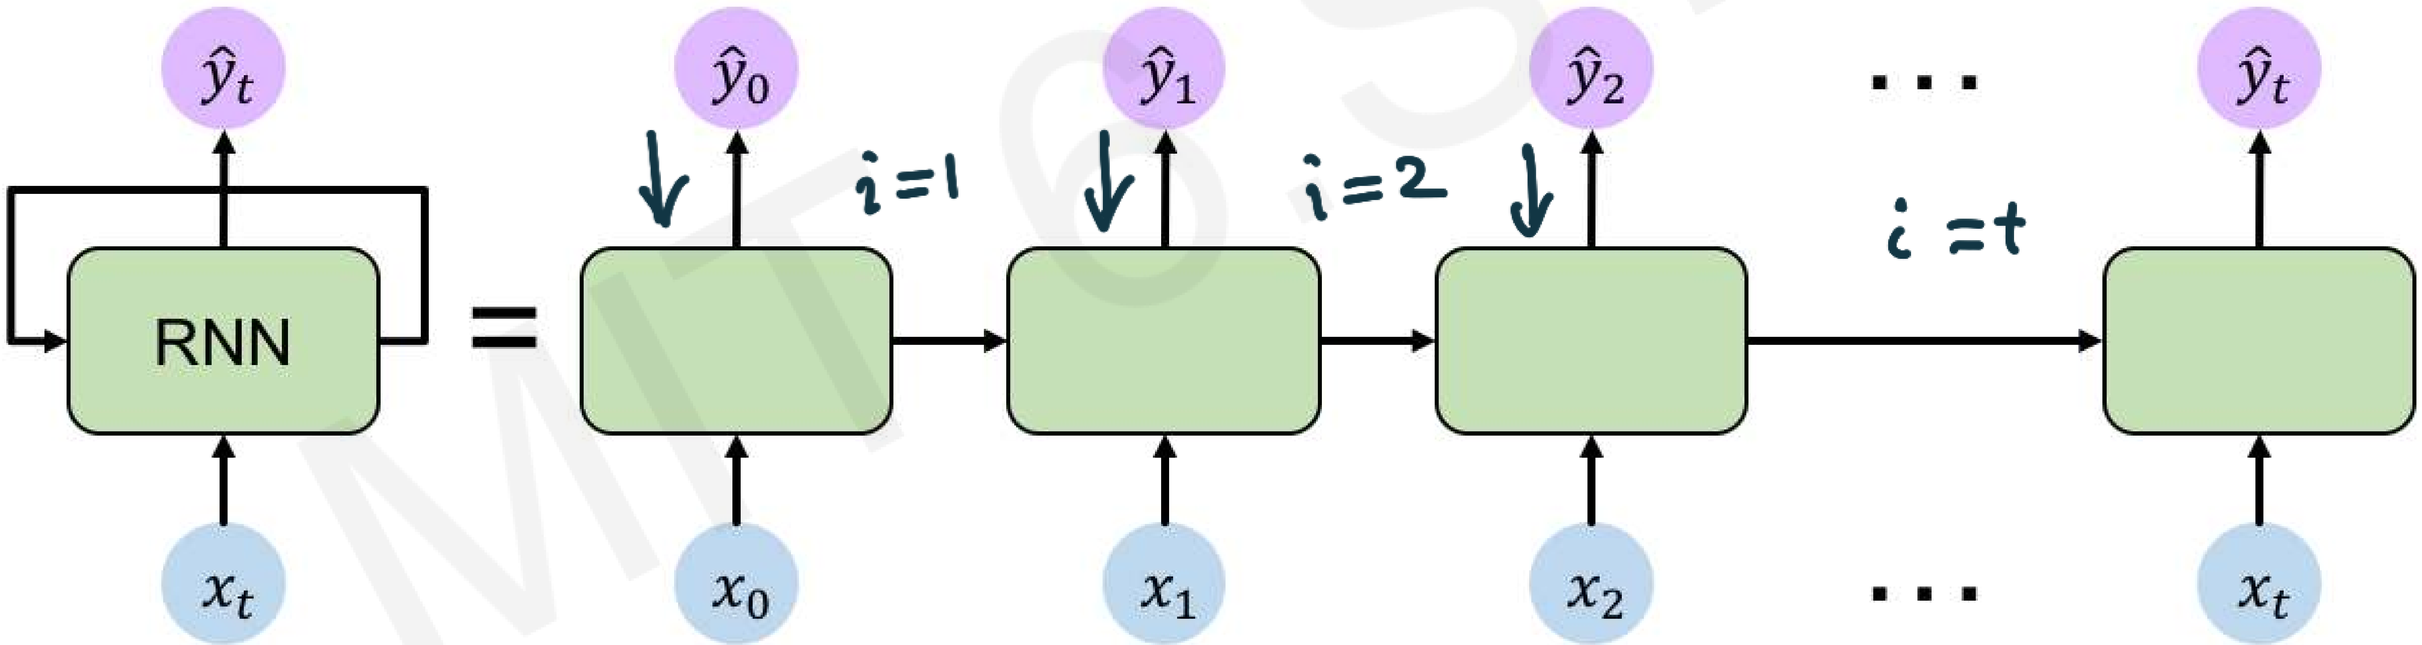

Backpropagation Through Time (BPTT)

Recall: Backpropagation in Feed Forward Models

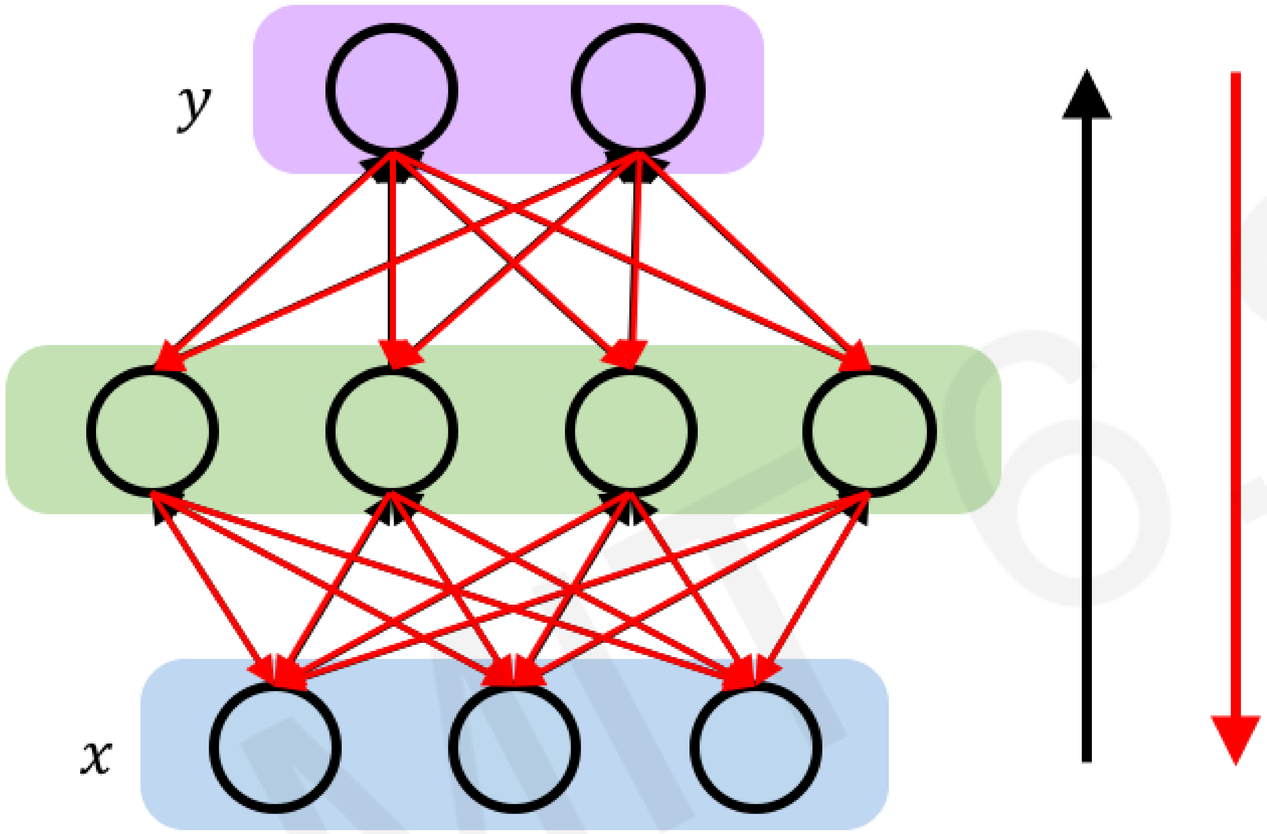

Backpropagation algorithm:

1. Take the derivative (gradient) of the loss with respect to each parameter
2. Shift parameters in order to minimize loss

RNNs: Backpropagation Through Time

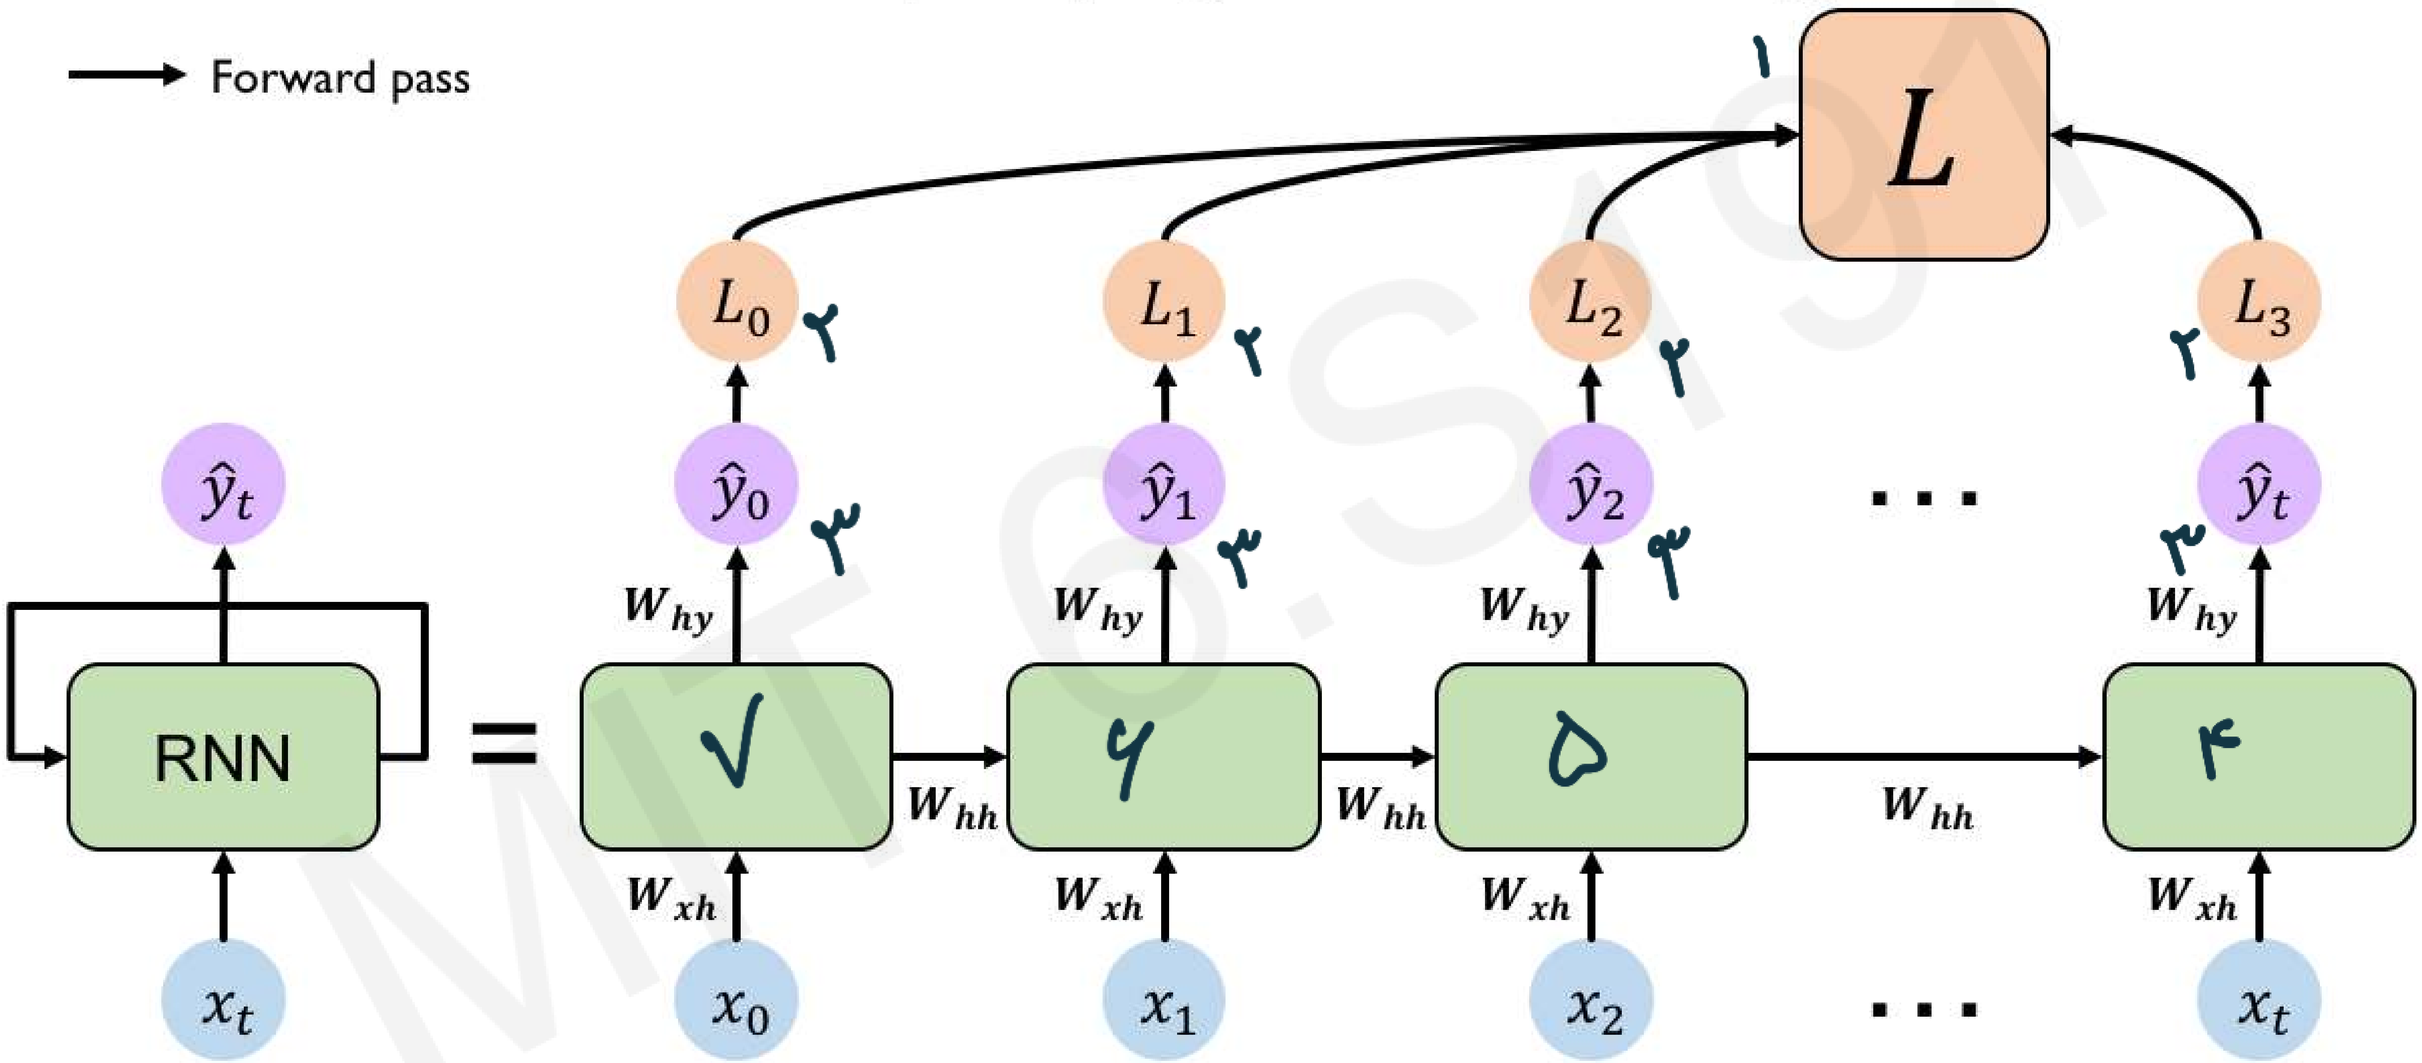

RNNs: Backpropagation Through Time

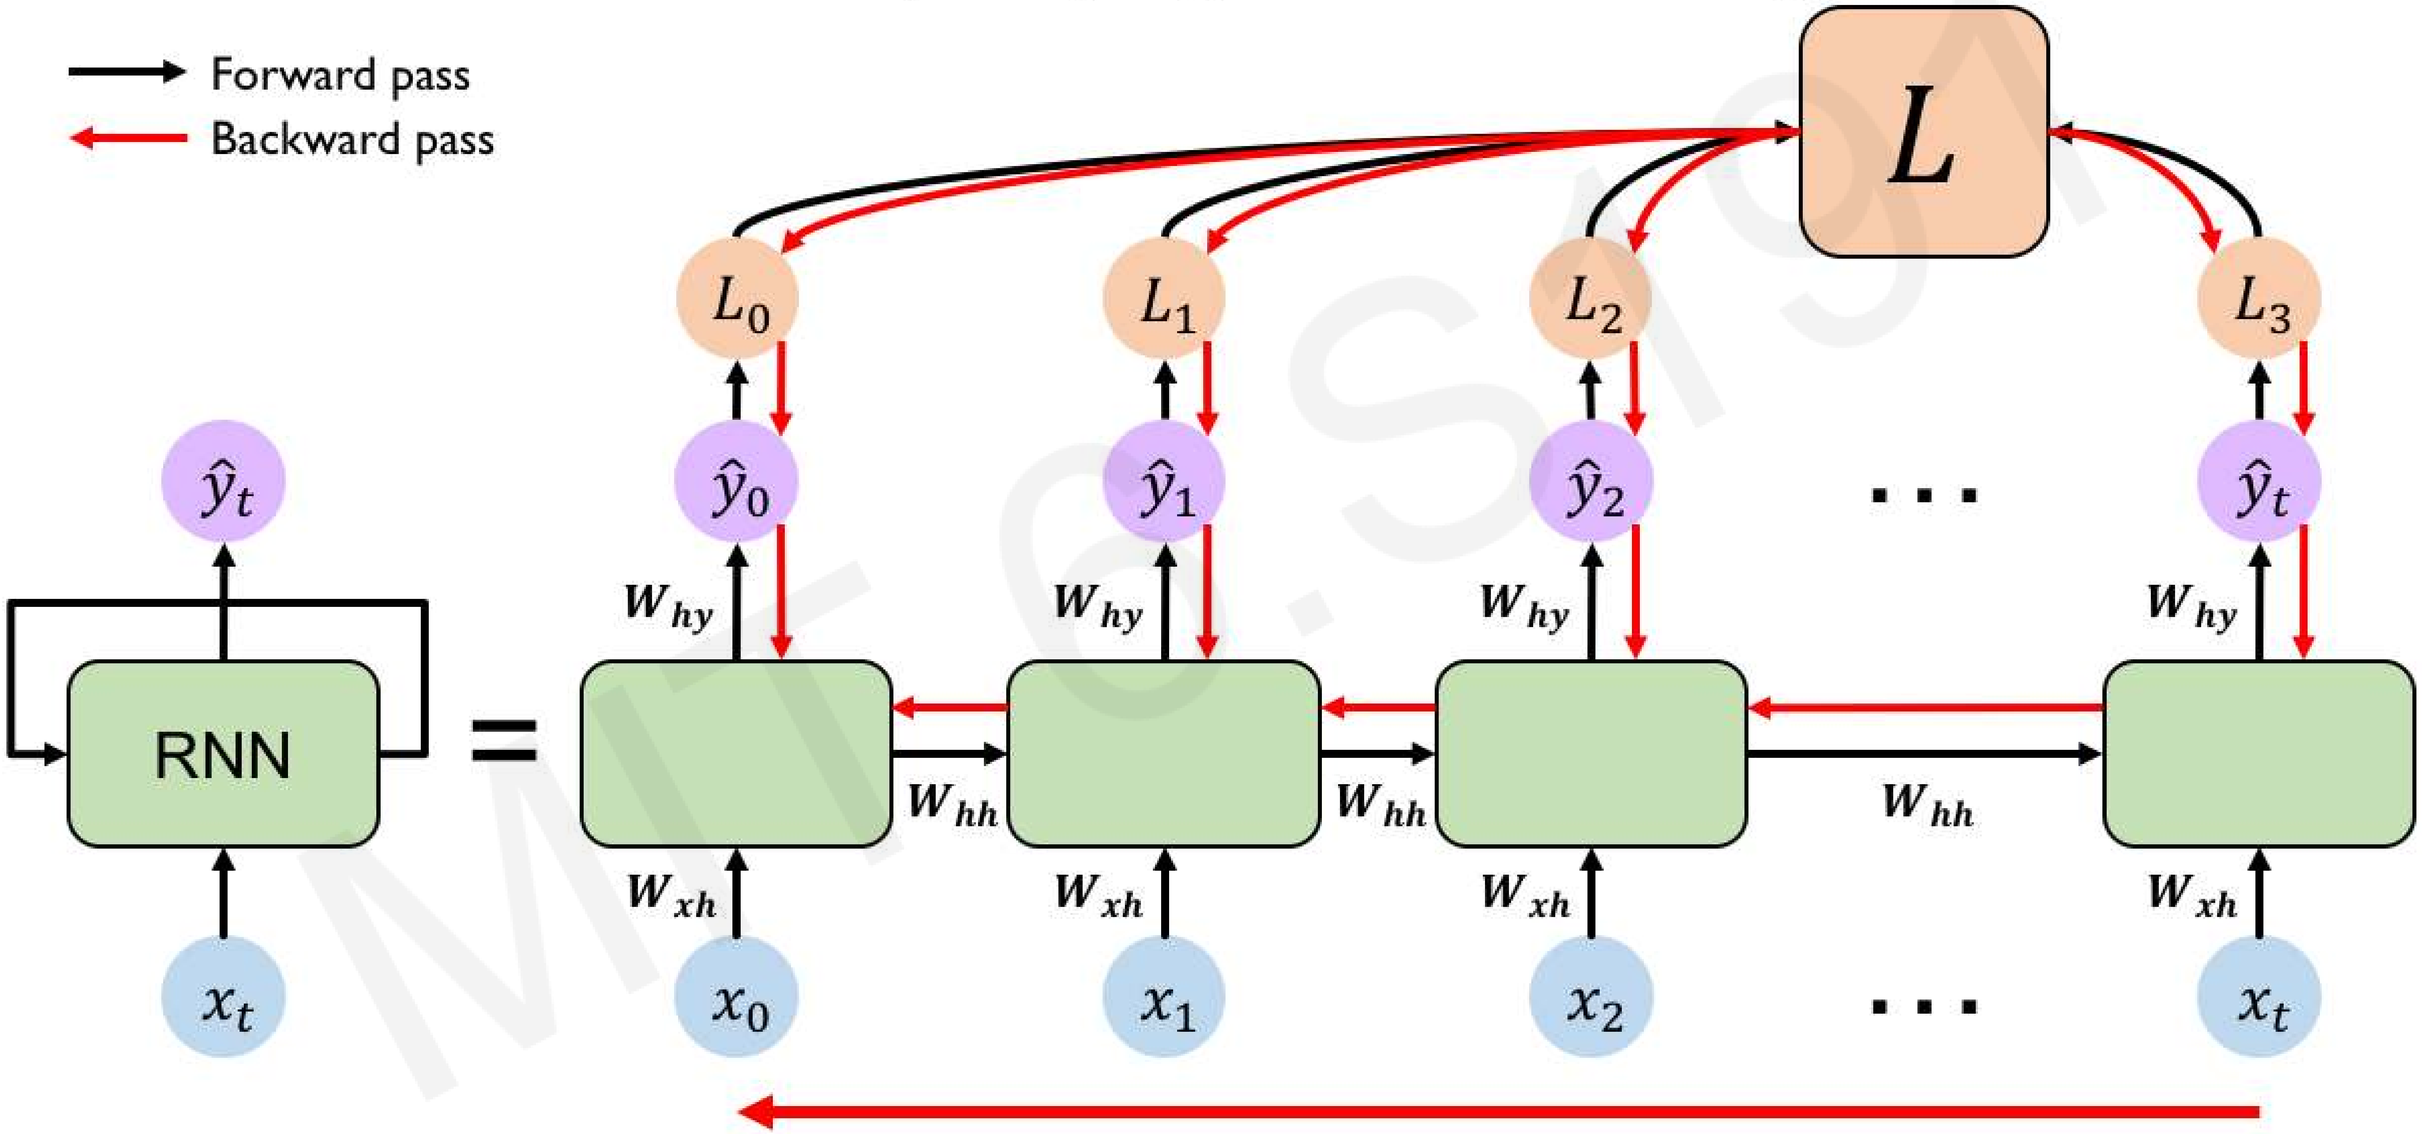

Standard RNN Gradient Flow

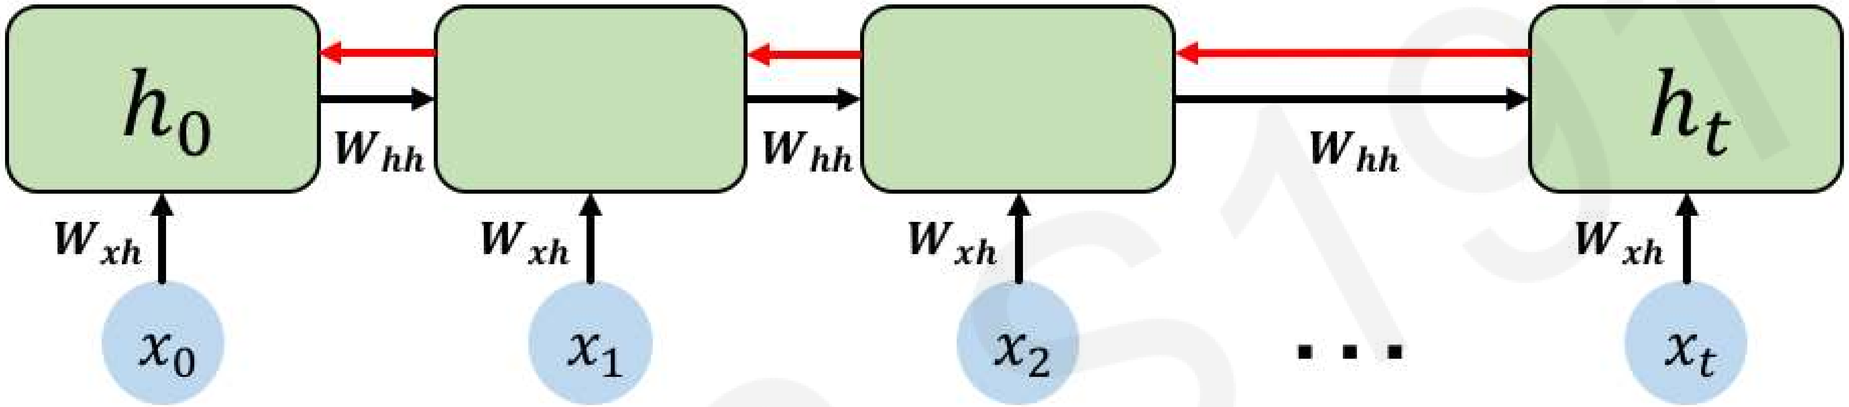

Standard RNN Gradient Flow

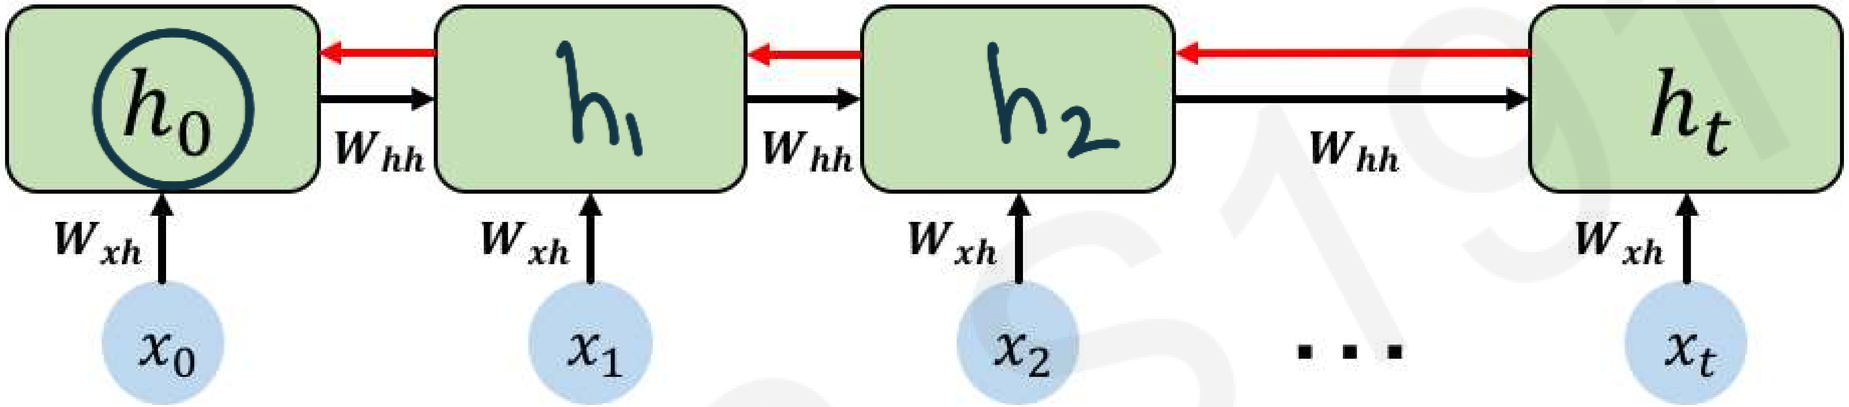

Computing the gradient wrt h_0 involves **many factors of W_{hh} + repeated gradient computation!**

$$W_{hh}^{(new)} = W_{hh}^{(old)} - \eta \sum_{i=1}^{t} \nabla \text{loss}_{W_{hh}^{(i)}}$$

BPTT

Standard RNN Gradient Flow: Exploding Gradients

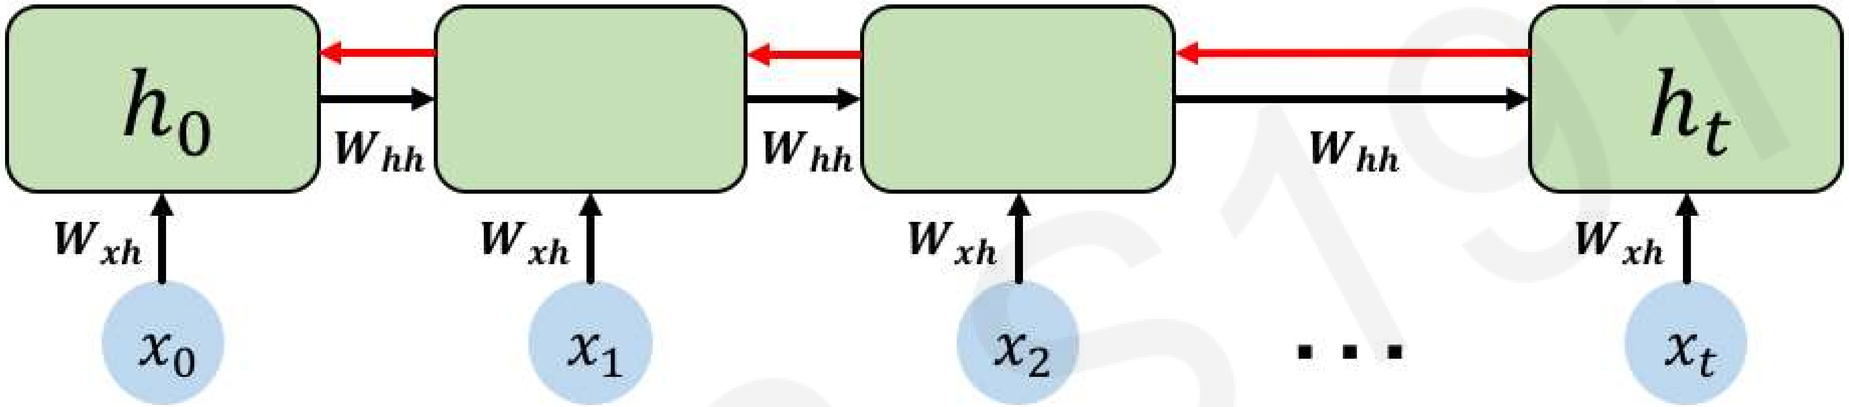

انفجار گرادیان

Computing the gradient wrt h_0 involves **many factors of W_{hh} + repeated gradient computation!**

Many values > 1:
exploding gradients

Gradient clipping to scale big gradients

$$\tilde{g} = \langle \text{sign}(g_i) \cdot \min(|g_i|, \text{threshold}) \rangle_{i=1,\dots,M}$$

$g = \langle 1, 4, -3, 5 \rangle$

$\tilde{g} = \langle 1, 2, -2, 2 \rangle$

threshold = 2

Traditional Standard RNN Gradient Flow: Vanishing Gradients

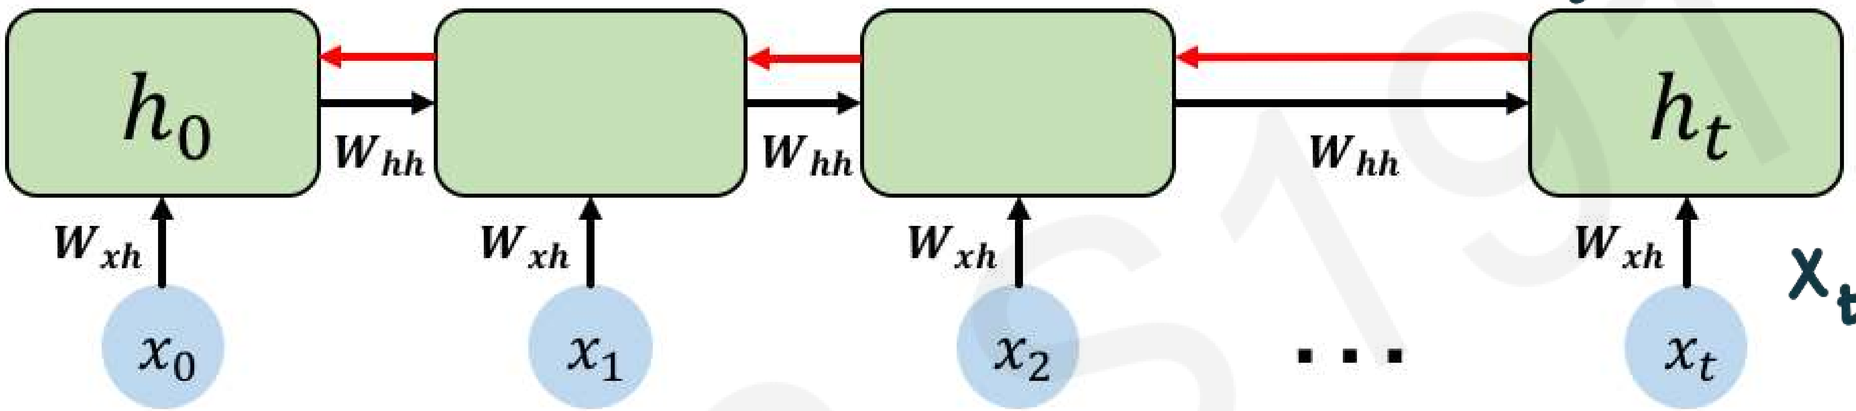

Computing the gradient wrt h_0 involves **many factors of W_{hh} + repeated gradient computation!**

Many values > 1:
exploding gradients

Gradient clipping to scale big gradients

Many values < 1:
vanishing gradients

1. Activation function
2. Weight initialization
3. Network architecture

The Problem of Long-Term Dependencies

Why are vanishing gradients a problem?

The Problem of Long-Term Dependencies

Why are vanishing gradients a problem?

Multiply many **small numbers** together

The Problem of Long-Term Dependencies

Why are vanishing gradients a problem?

Multiply many **small numbers** together

Errors due to further back time steps have smaller and smaller gradients

The Problem of Long-Term Dependencies

Why are vanishing gradients a problem?

Multiply many **small numbers** together

↓

Errors due to further back time steps have smaller and smaller gradients

↓

Bias parameters to capture short-term dependencies

The Problem of Long-Term Dependencies

"The clouds are in the ___"

Why are vanishing gradients a problem?

Multiply many **small numbers** together

Errors due to further back time steps have smaller and smaller gradients

Bias parameters to capture short-term dependencies

The Problem of Long-Term Dependencies

Why are vanishing gradients a problem?

Multiply many **small numbers** together

↓

Errors due to further back time steps have smaller and smaller gradients

↓

Bias parameters to capture short-term dependencies

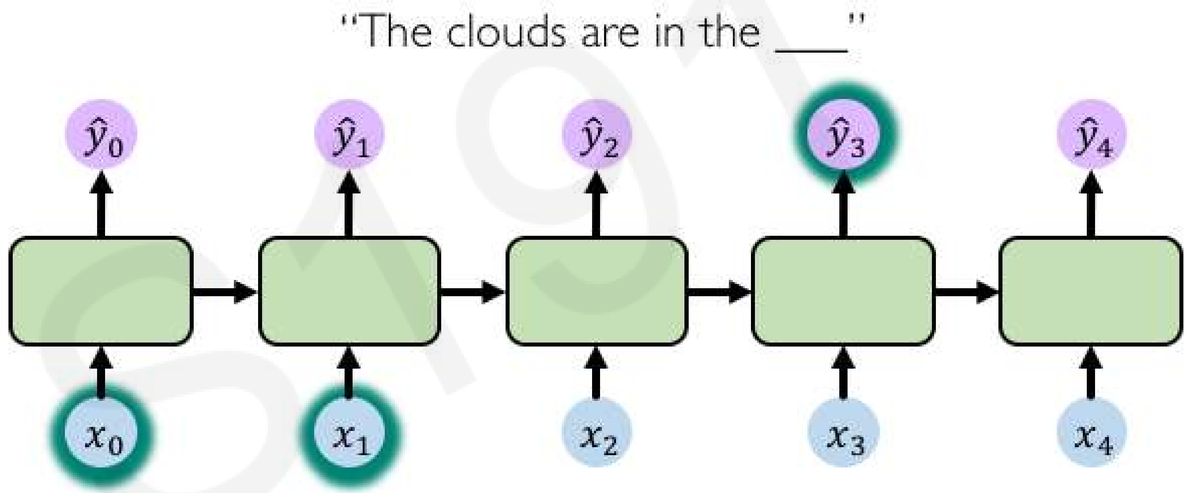

The Problem of Long-Term Dependencies

Why are vanishing gradients a problem?

Multiply many **small numbers** together

↓

Errors due to further back time steps have smaller and smaller gradients

↓

Bias parameters to capture short-term dependencies

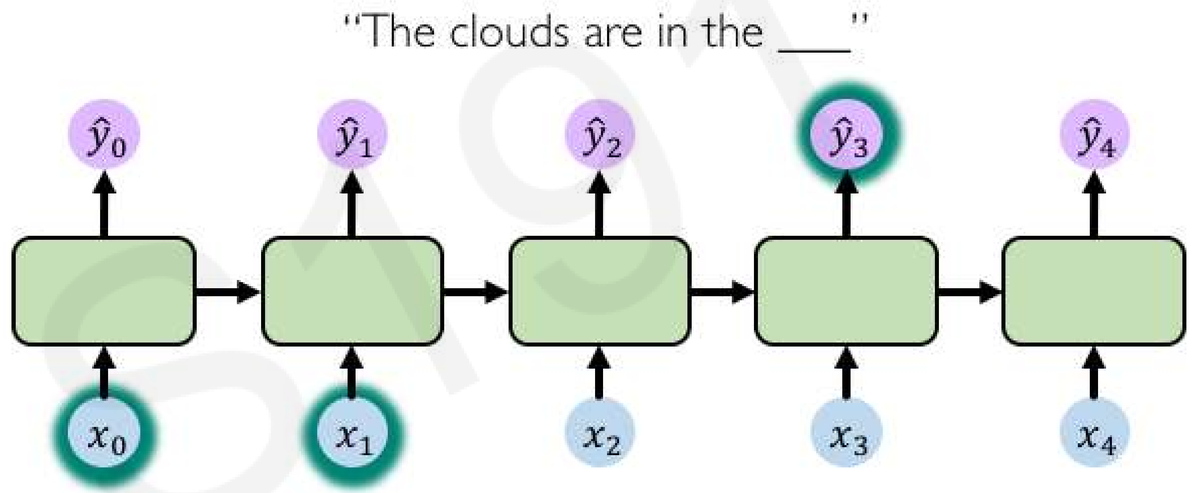

The Problem of Long-Term Dependencies

Why are vanishing gradients a problem?

Multiply many **small numbers** together

↓

Errors due to further back time steps have smaller and smaller gradients

↓

Bias parameters to capture short-term dependencies

"The clouds are in the ___"

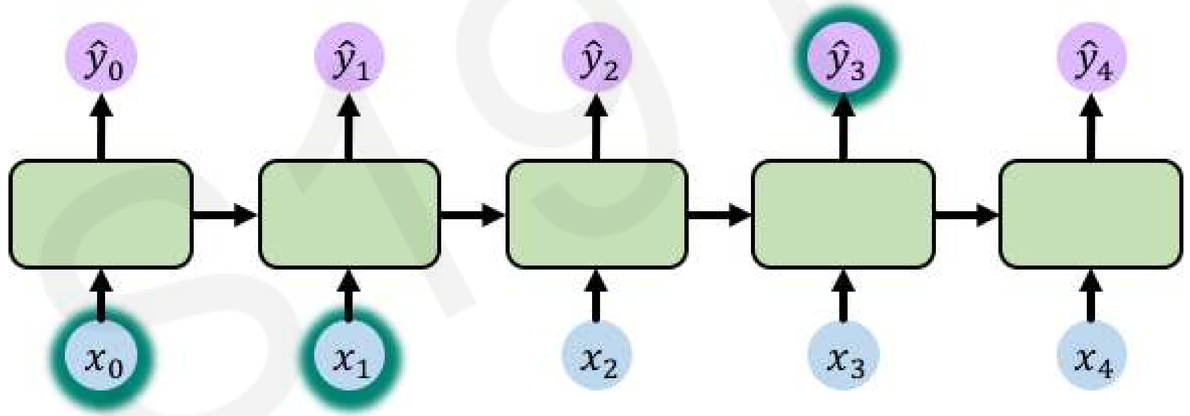

"I grew up in France, … and I speak fluent___ "

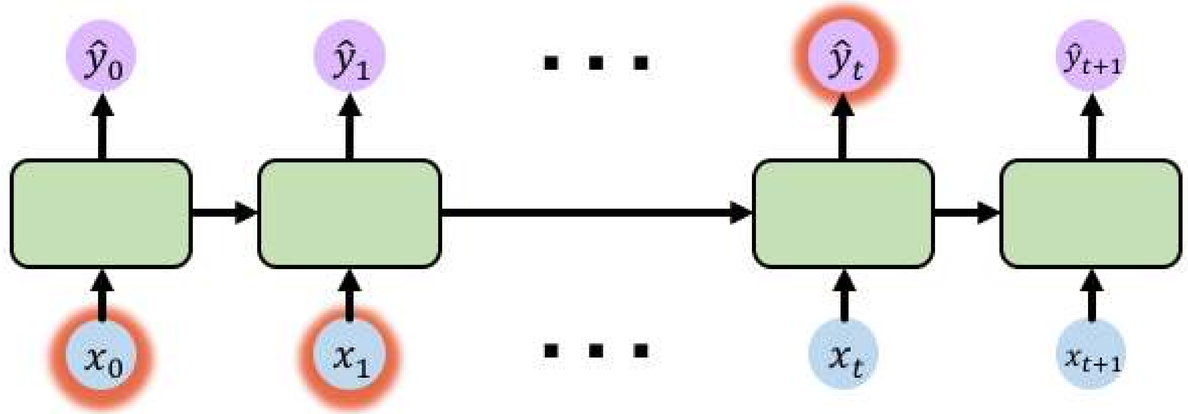

Trick #1: Activation Functions

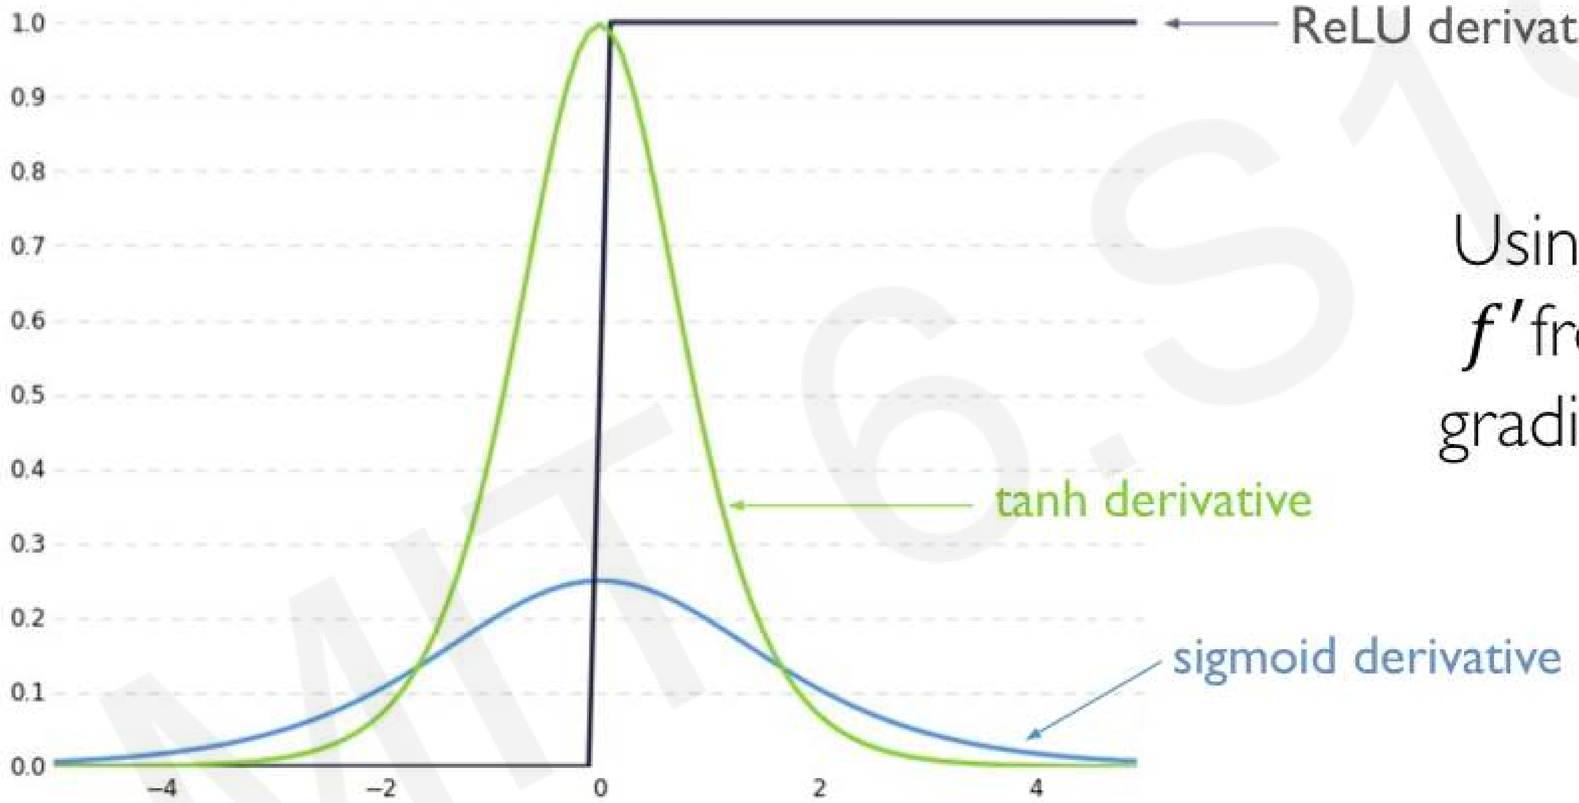

Using ReLU prevents f' from shrinking the gradients when $x > 0$

Trick #2: Parameter Initialization

Initialize **weights** to identity matrix

Initialize **biases** to zero ✓✓

$$I_n = \begin{pmatrix} 1 & 0 & 0 & \cdots & 0 \\ 0 & 1 & 0 & \cdots & 0 \\ 0 & 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & 0 & \cdots & 1 \end{pmatrix}$$

This helps prevent the weights from shrinking to zero.

Solution #3: Gated Cells

Idea: use a more **complex recurrent unit with gates** to control what information is passed through

حافظہ طولانی کوتاہ مدت

Long Short Term Memory (LSTMs) networks rely on a gated cell to track information throughout many time steps.

Long Short Term Memory (LSTM) Networks

Standard RNN

In a standard RNN, repeating modules contain a **simple computation node**

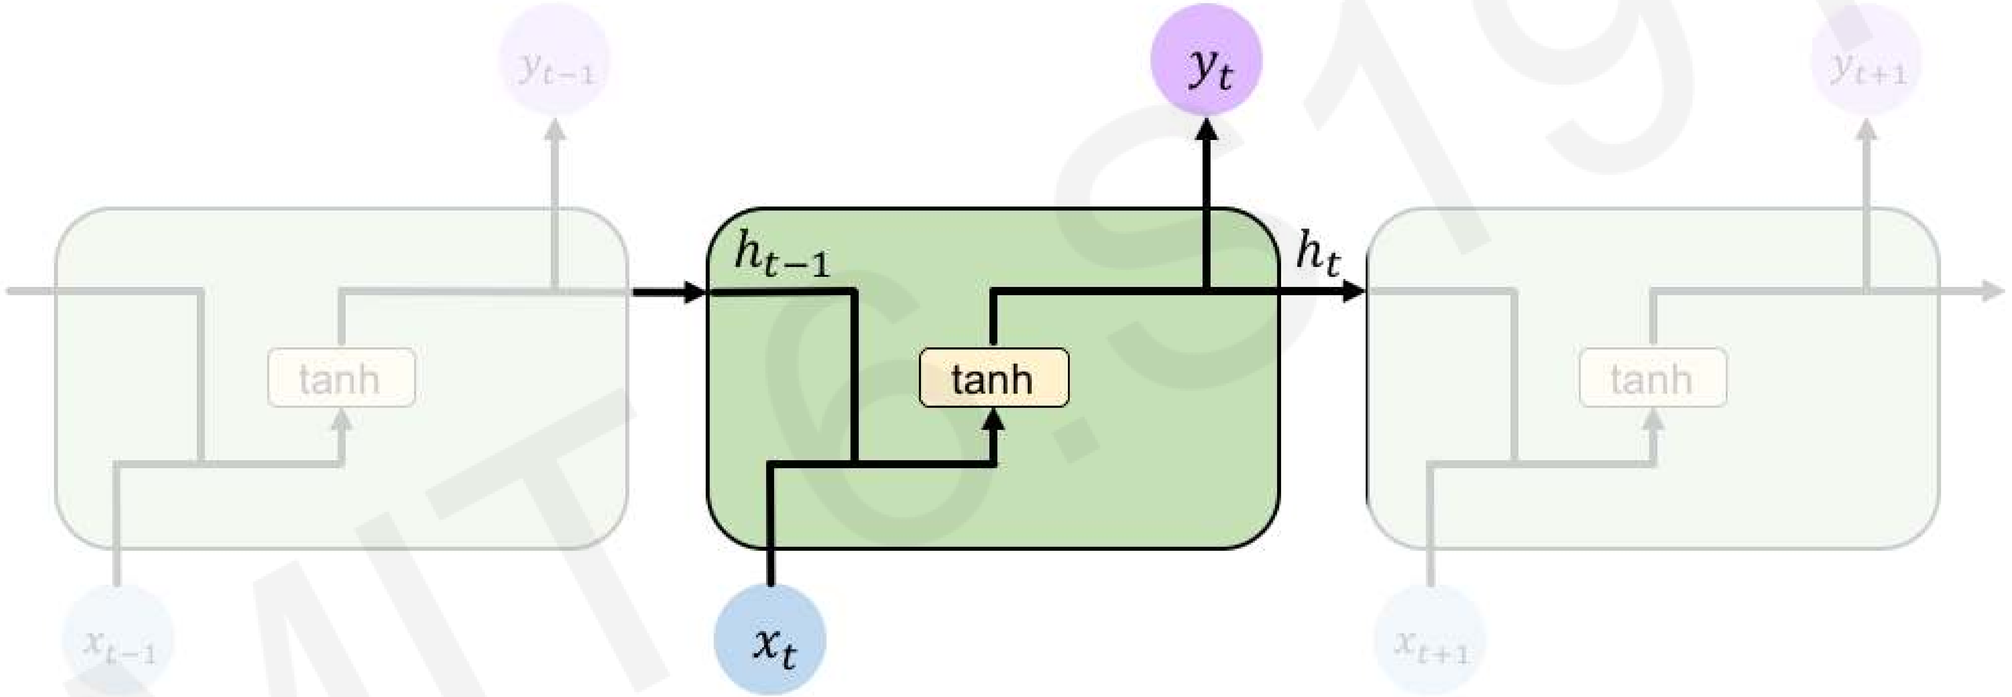

Long Short Term Memory (LSTMs)

LSTM modules contain **computational blocks** that **control information flow**

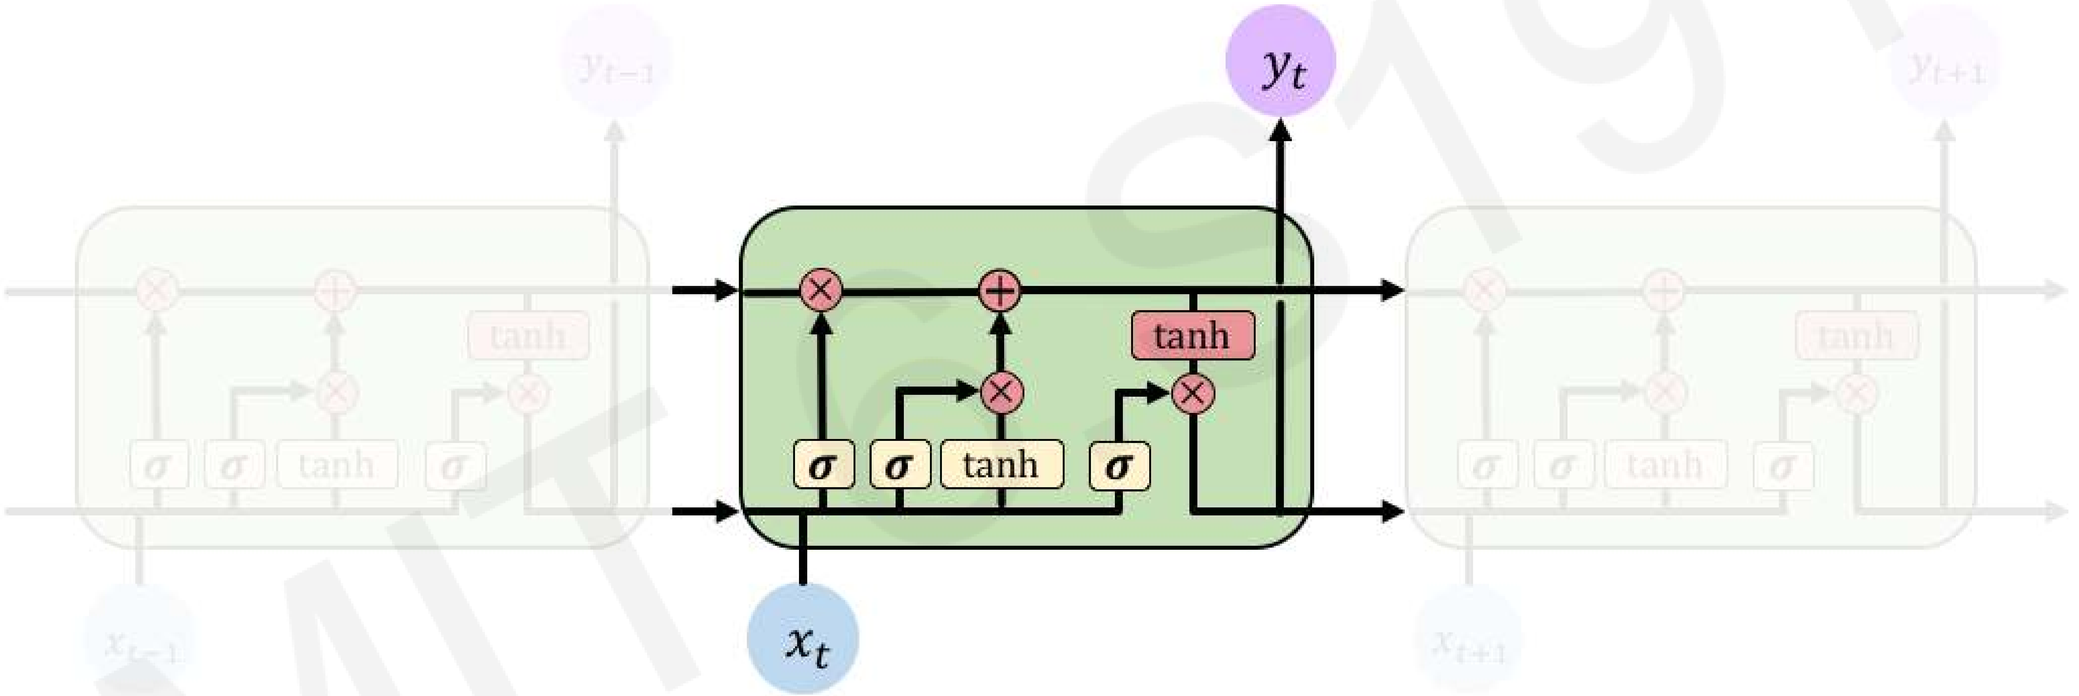

LSTM cells are able to track information throughout many timesteps

```
tf.keras.layers.LSTM(num_units)
```

Long Short Term Memory (LSTMs)

Information is **added** or **removed** through structures called **gates**

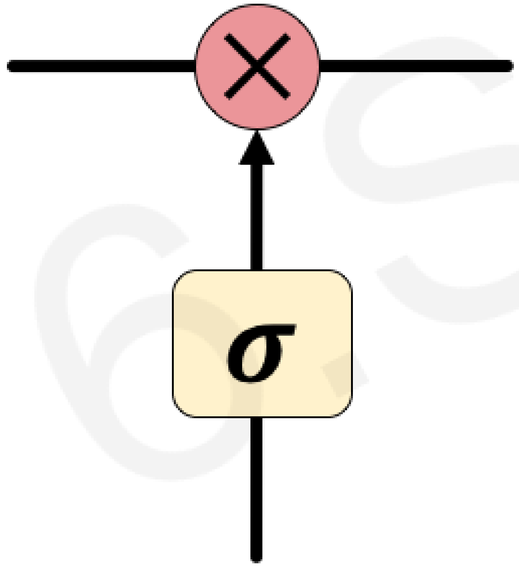

Gates optionally let information through, for example via a sigmoid neural net layer and pointwise multiplication

Long Short Term Memory (LSTMs)

How do LSTMs work?

1) Forget 2) Store 3) Update 4) Output

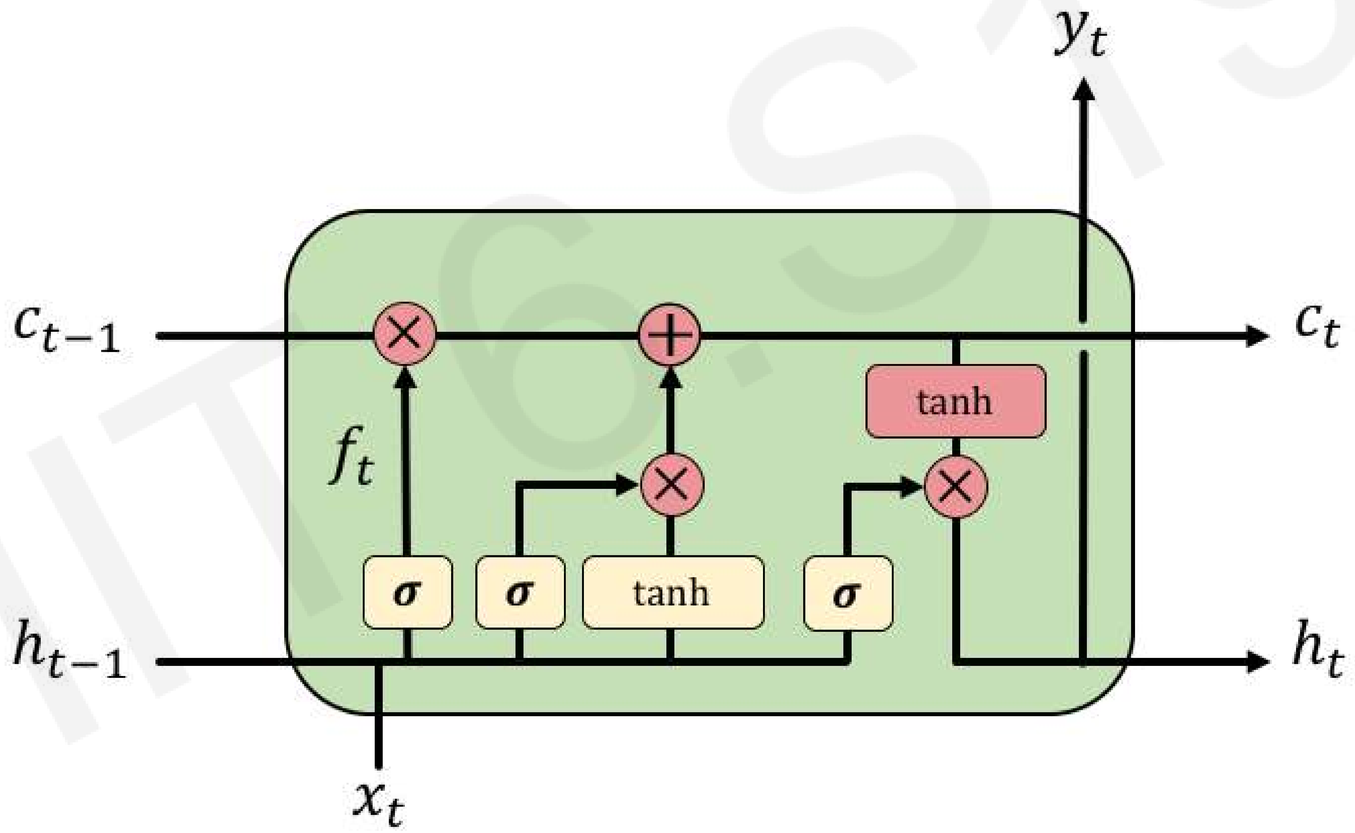

Long Short Term Memory (LSTMs)

1) Forget 2) Store 3) Update 4) Output

LSTMs **forget irrelevant** parts of the previous state

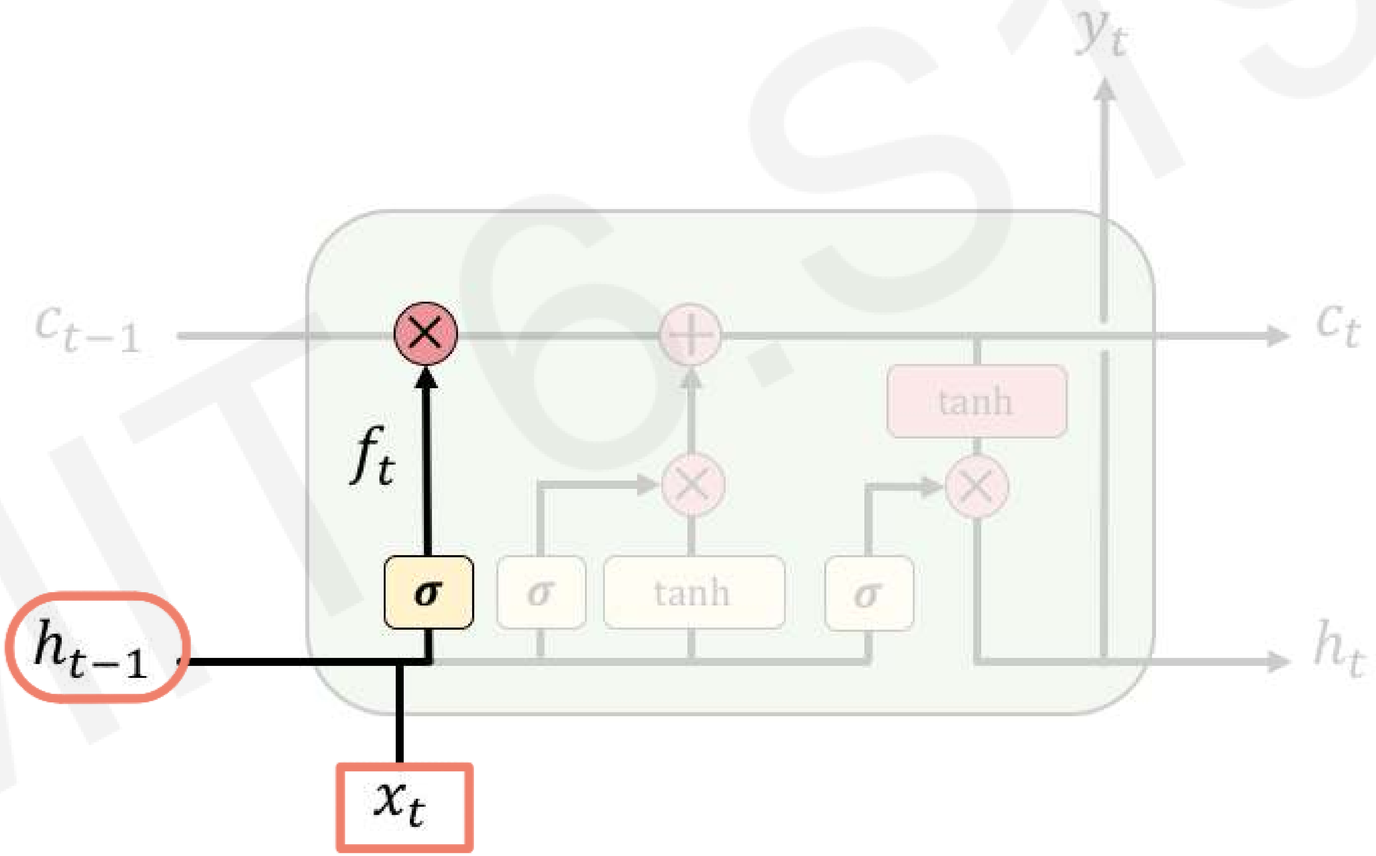

Long Short Term Memory (LSTMs)

1) Forget **2) Store** 3) Update 4) Output

LSTMs **store relevant** new information into the cell state

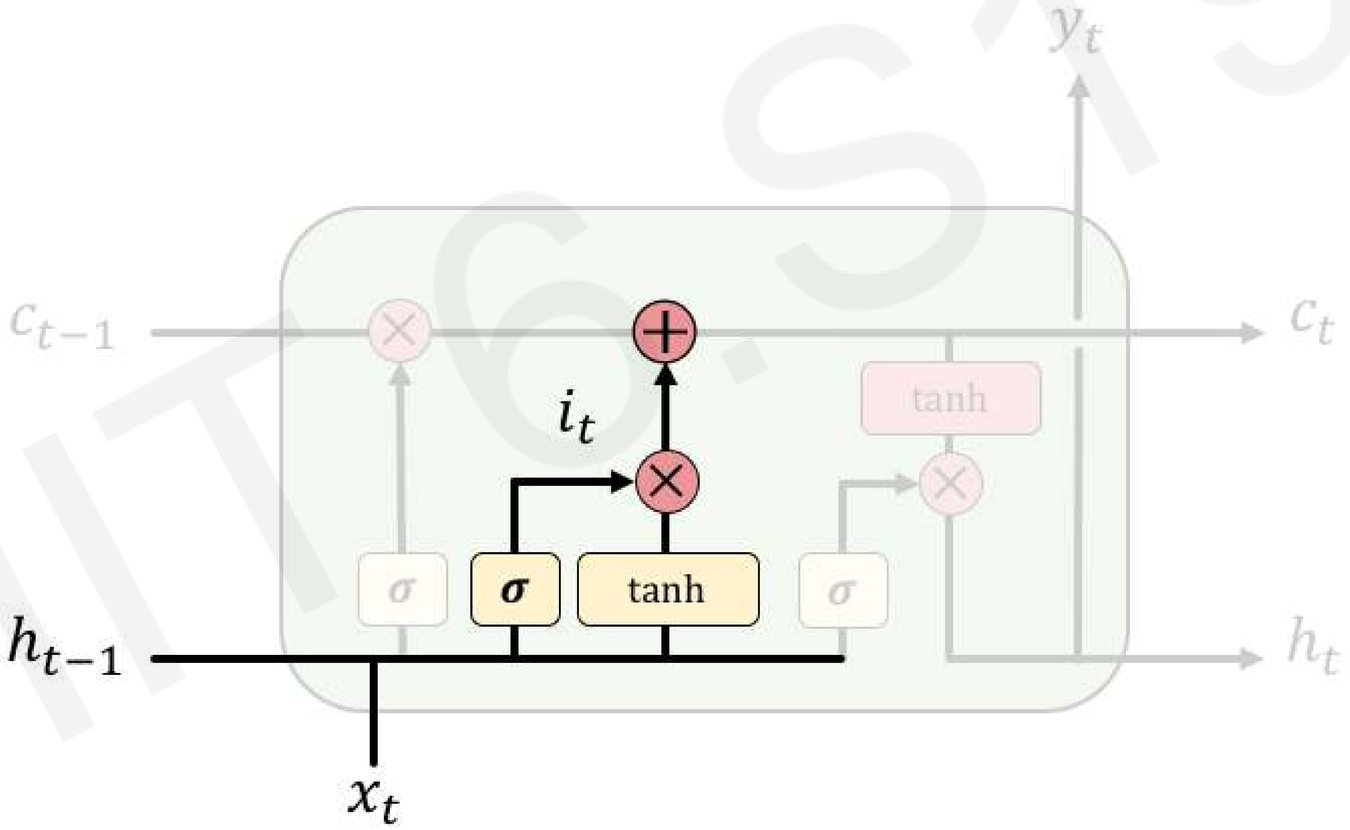

Long Short Term Memory (LSTMs)

1) Forget 2) Store **3) Update** ✓ 4) Output

LSTMs **selectively update** cell state values

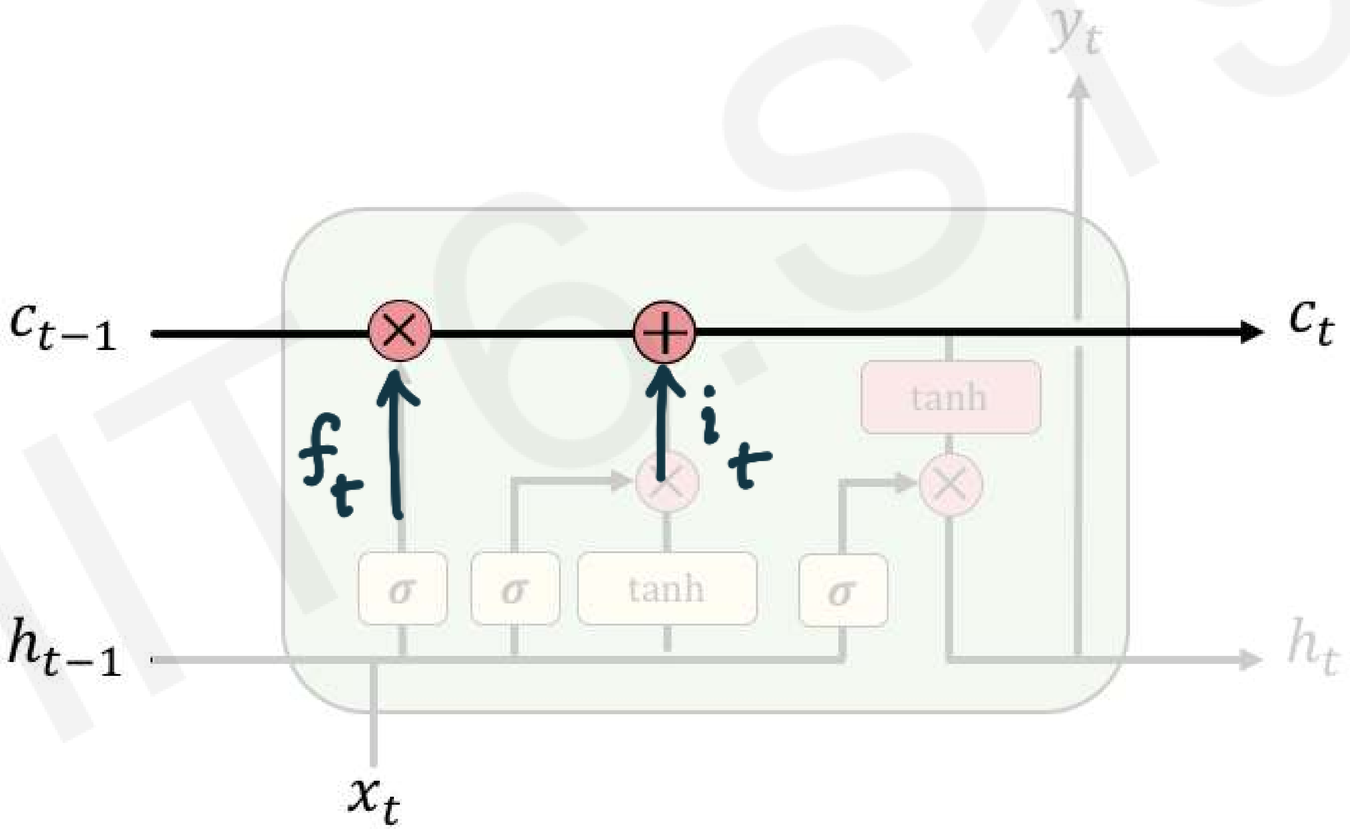

Long Short Term Memory (LSTMs)

1) Forget 2) Store 3) Update **4) Output** ✓

The **output gate** controls what information is sent to the next time step

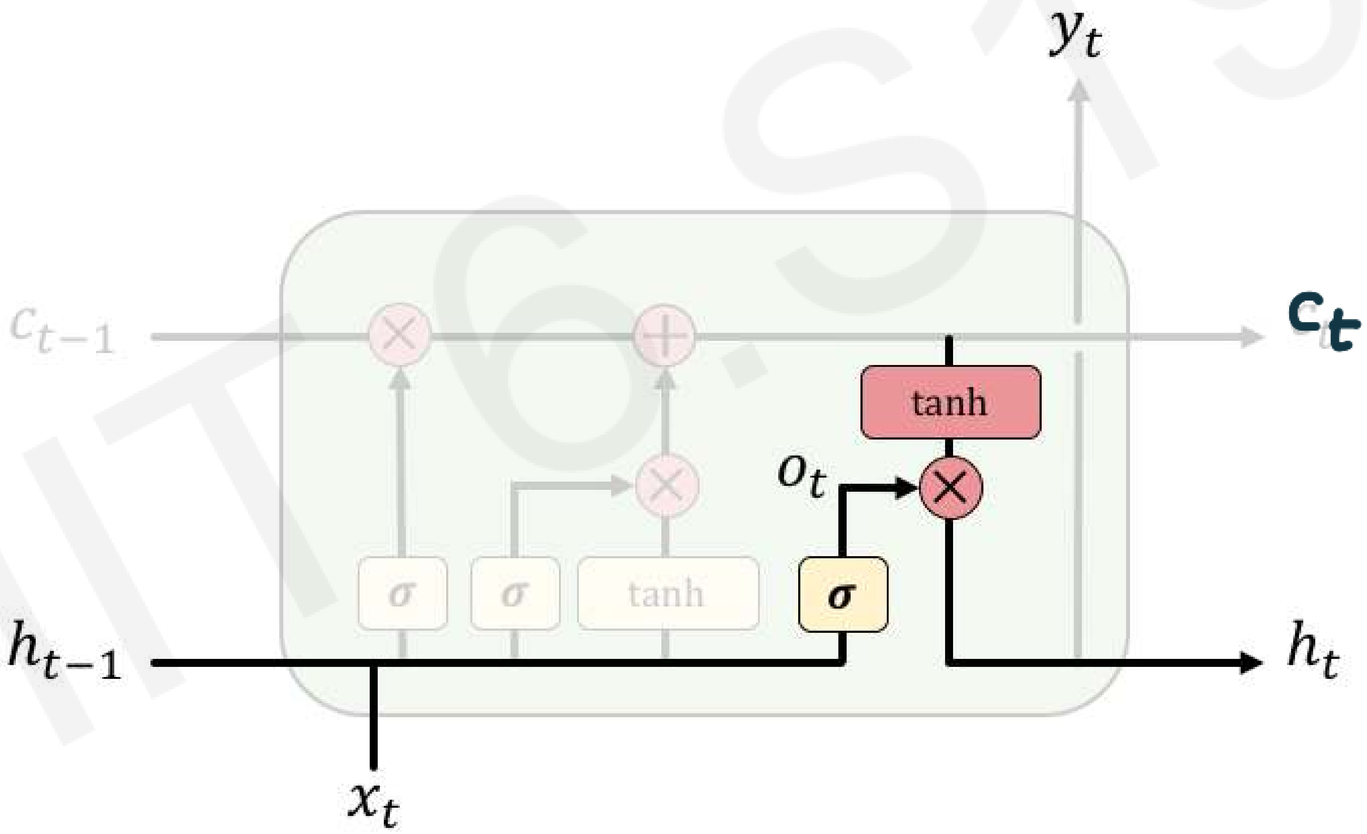

Long Short Term Memory (LSTMs)

1) Forget 2) Store 3) Update 4) Output

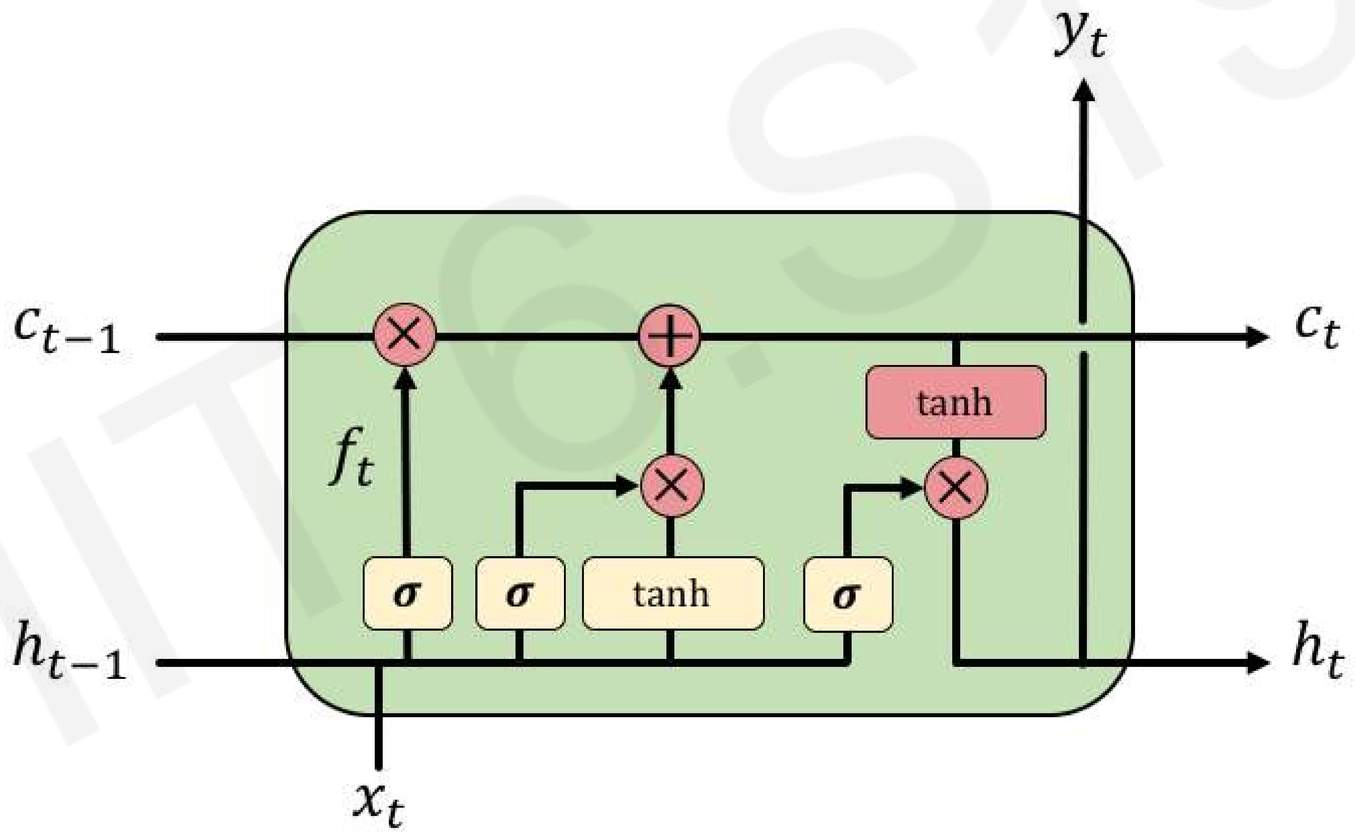

LSTM Gradient Flow

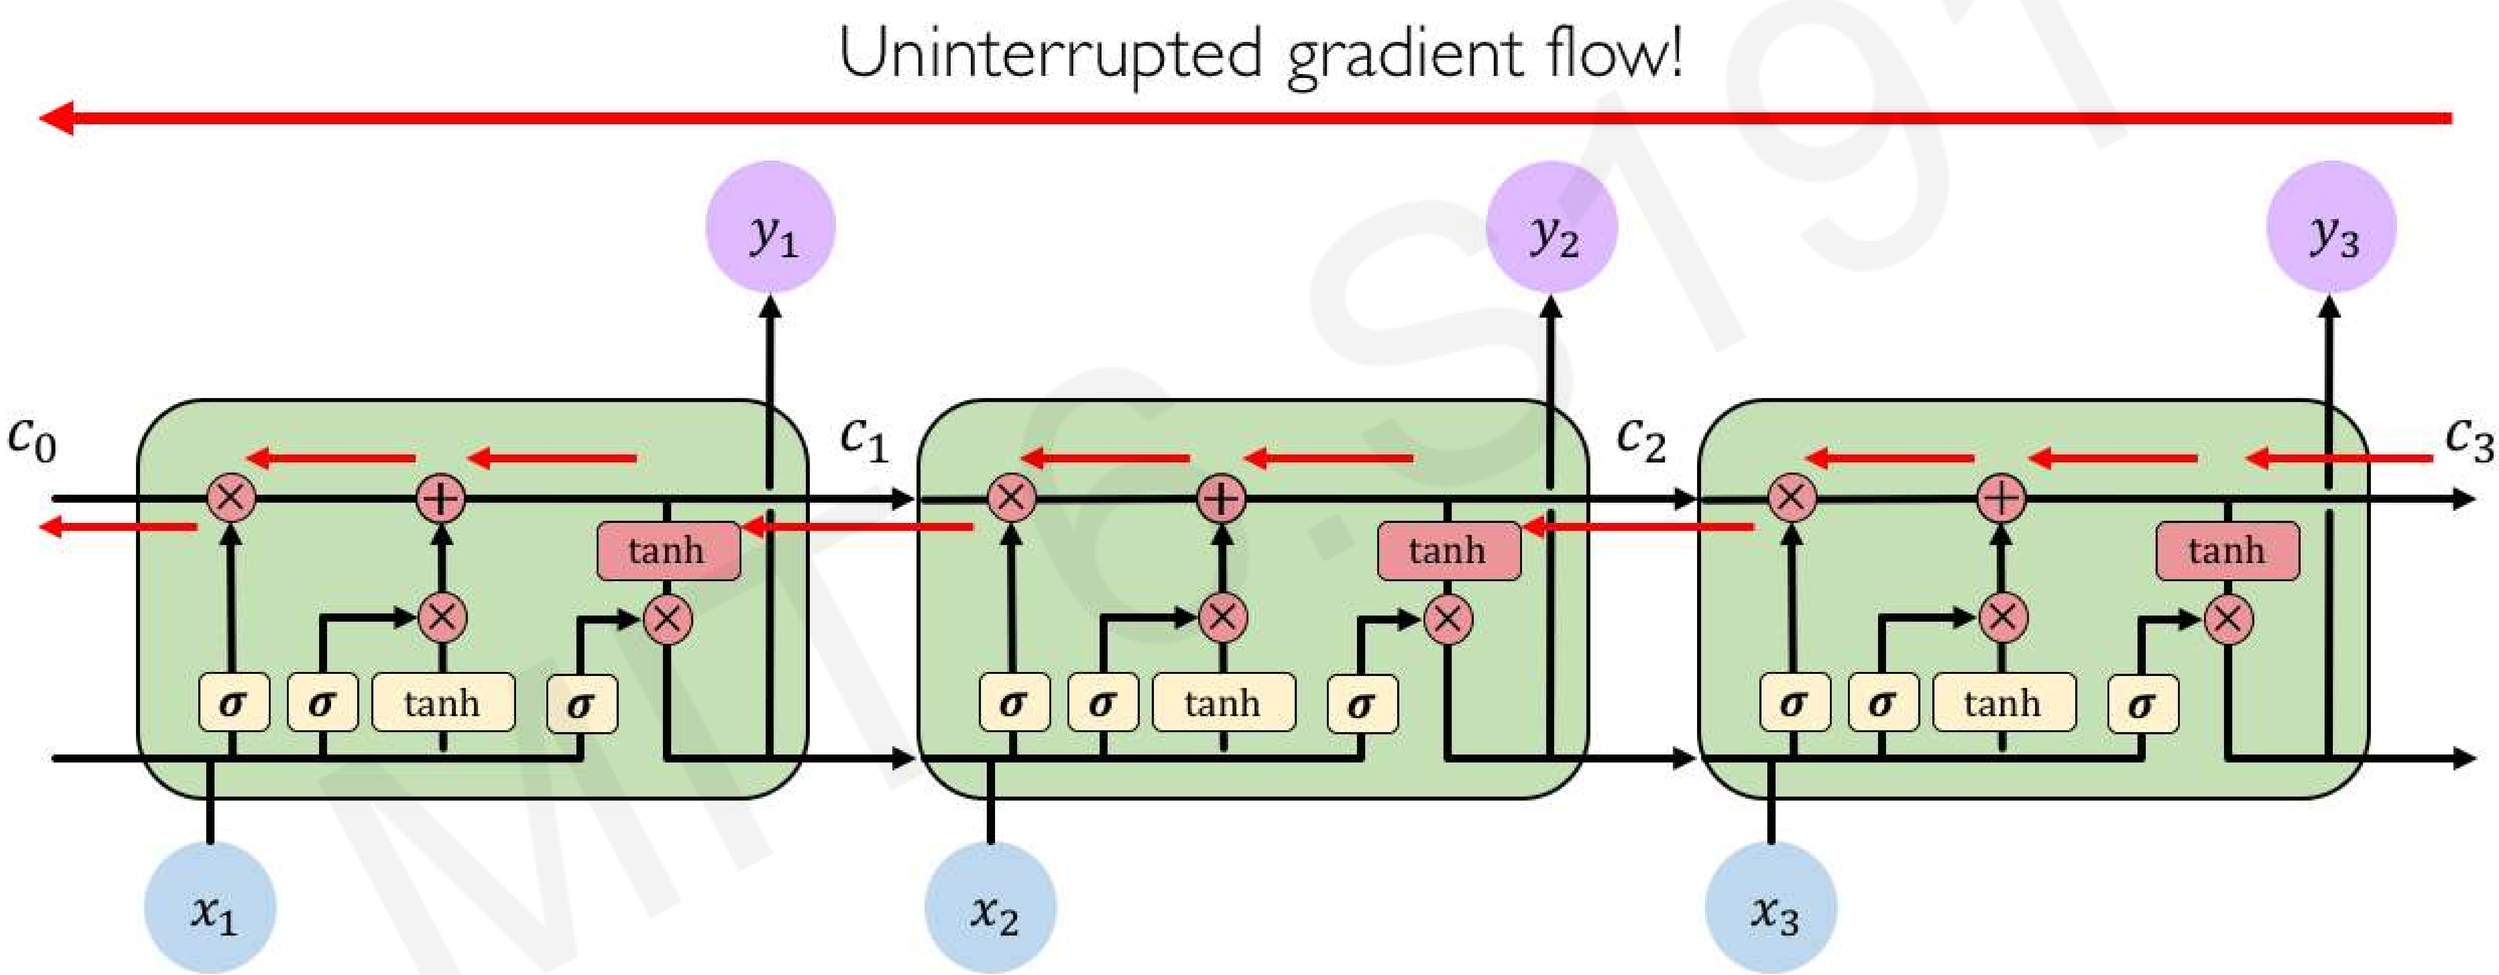

LSTMs: Key Concepts

1. Maintain a **separate cell state** from what is outputted
2. Use **gates** to control the **flow of information**
   - **Forget** gate gets rid of irrelevant information
   - **Store** relevant information from current input
   - Selectively **update** cell state
   - **Output** gate returns a filtered version of the cell state
3. Backpropagation through time with **uninterrupted gradient flow**

RNN Applications

Example Task: Music Generation

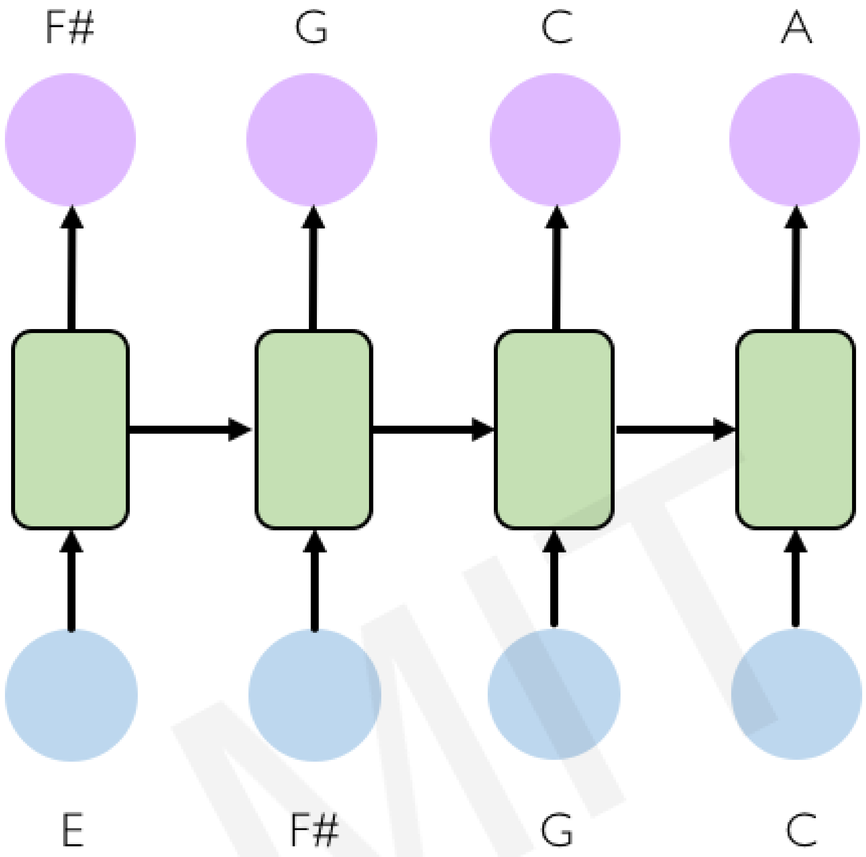

Input: sheet music

Output: next character in sheet music

6.S191 Lab!

Example Task: Sentiment Classification

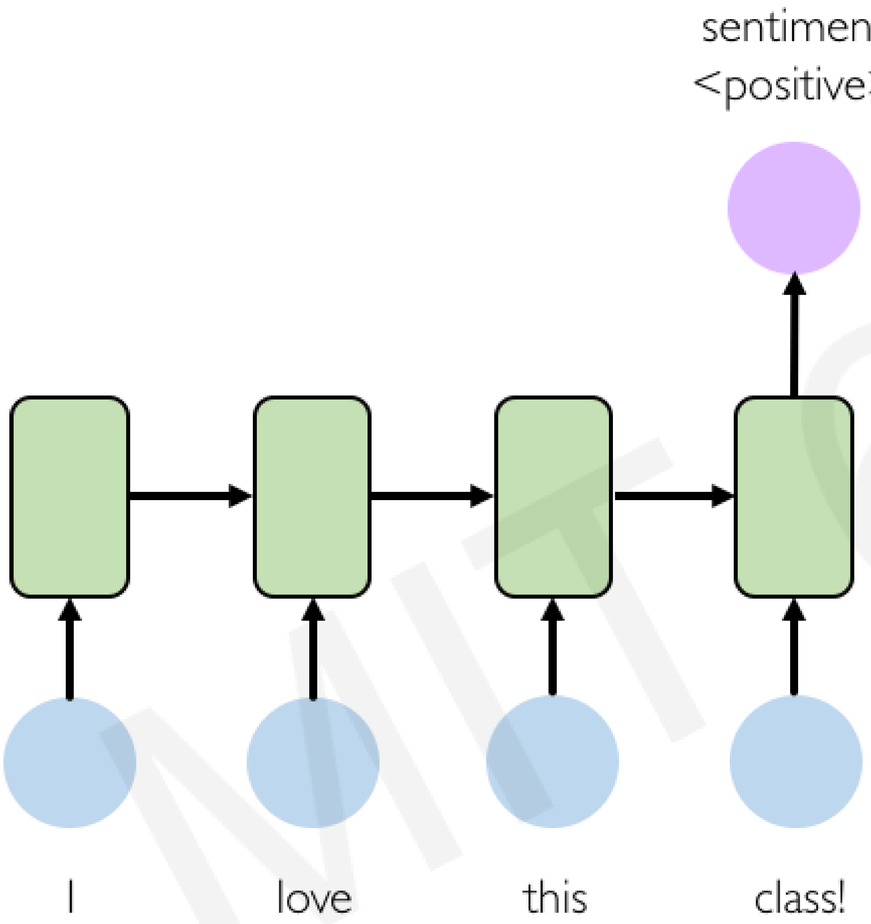

Input: sequence of words

Output: probability of having positive sentiment

```
loss = tf.nn.softmax_cross_entropy_with_logits(y, predicted)
```

Example Task: Sentiment Classification

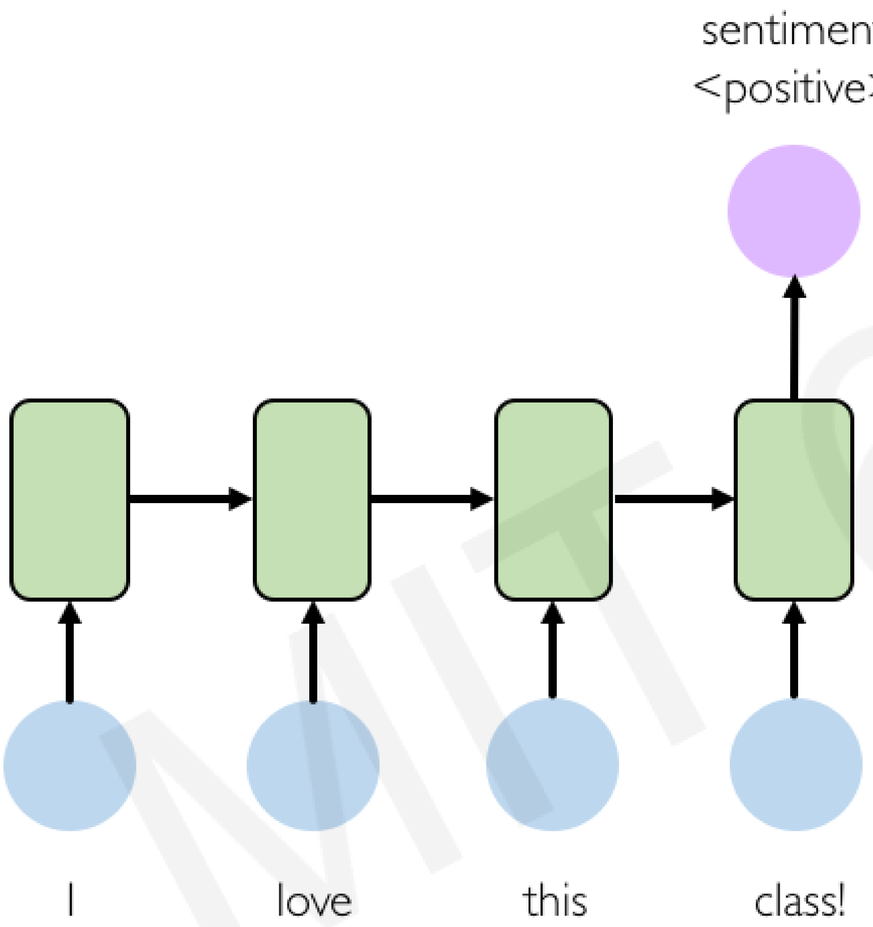

Tweet sentiment classification

Example Task: Machine Translation

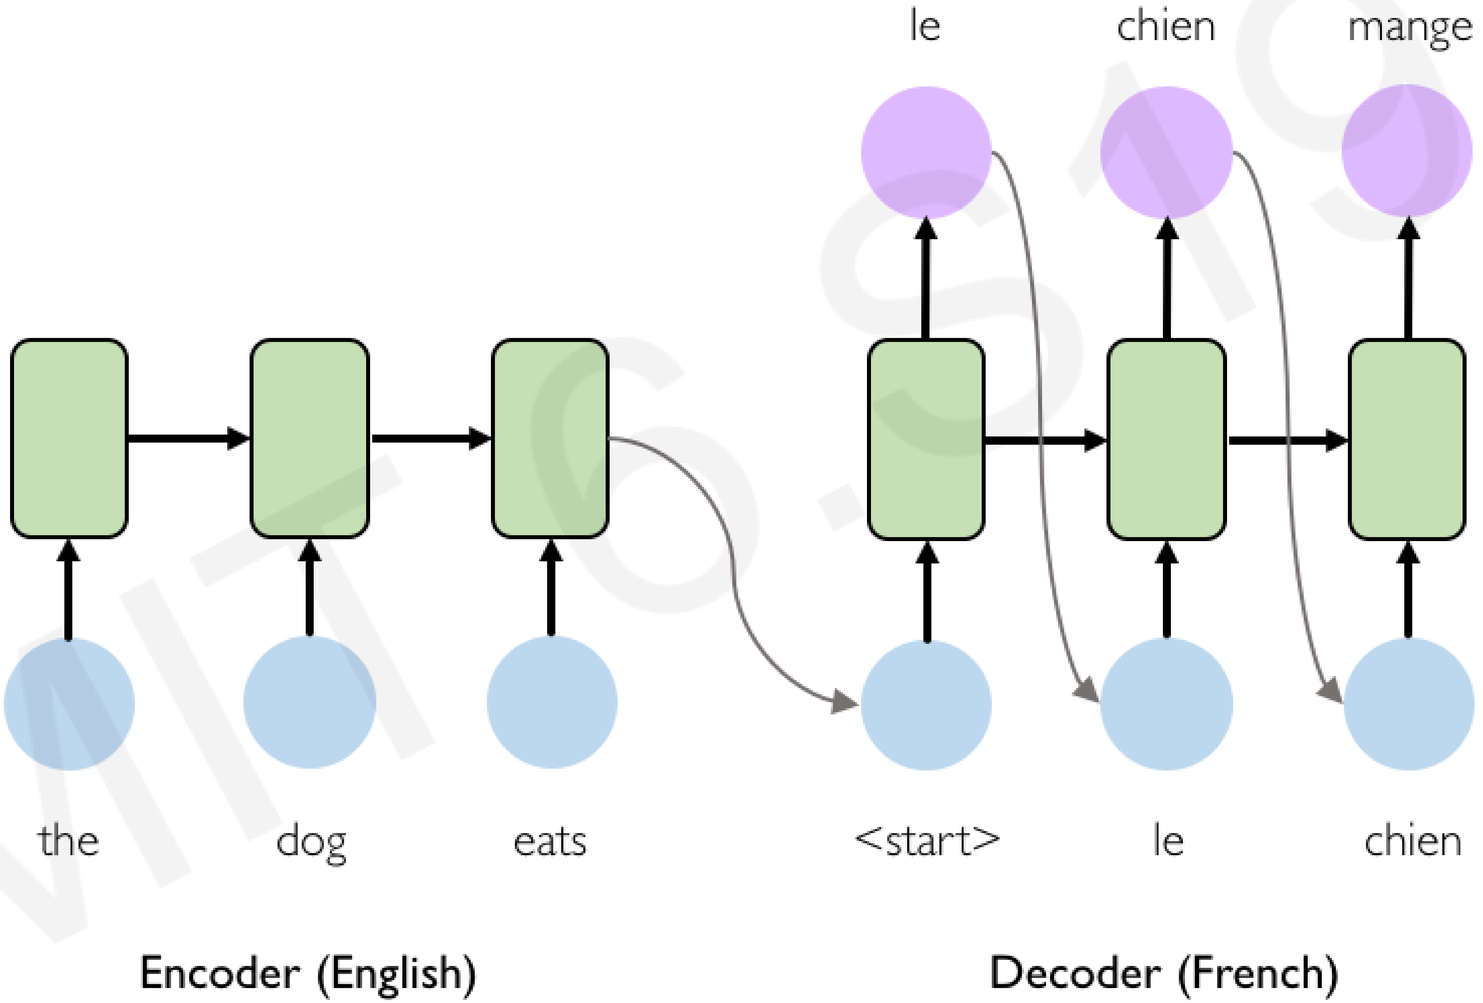

Example Task: Machine Translation