

# Introdução à Aprendizagem Automática (IAA)

---

SUSANA BRÁS

SUSANA.BRAS@UA.PT

---

# IAA – L11

---

## Sequence models

### Time Series decomposition

- Trend, Seasonality, Residual

### Time Series forecasting

- Principles
- Classical models, ML models, DL models
- Metrics

### Recurrent Neural Networks (RNN)

- Main characteristic
- Basic unit
- Optimization
- Memory
- Vanishing gradient problem

### Long-Short Term Memory (LSTM)

- Why and what
- Forget gate
- Update gate
- Output gate

### Gated Recurrent Unit (GRU)

- Why and what
- Gated recurrent unit

# Data Matrix

---

matrix X (mxn)	feature $x_1$	feature $x_2$	....	feature $x_n$	Target Y
Example 1	$x_1^{(1)}$	$x_2^{(1)}$		$x_n^{(1)}$	$y^{(1)}$
Example 2	$x_1^{(2)}$	$x_2^{(2)}$		$x_n^{(2)}$	$y^{(2)}$
...					
Example i	$x_1^{(i)}$	$x_2^{(i)}$		$x_n^{(i)}$	$y^{(i)}$
...					
...					
Example m	$x_1^{(m)}$	$x_2^{(m)}$		$x_n^{(m)}$	$y^{(m)}$

$x$  – input vector of features, attributes

$y$  – output vector of labels, ground truth, target

$m$  - number of training examples

$n$  – number of features

$h_{\theta}(x)$  - model (hypothesis)

$\theta$  - vector of model parameters

Training set: data matrix X (m rows, n columns)

# Sequence Models

# Deep Learning

---

Sequence models are machine learning models that process data organized as sequences: text, audio, video, or time series data.

They are designed to capture the dependencies between elements within a sequence.

The key point for sequence models is that the data we are processing are no longer independently and identically distributed (i.i.d.) samples; the data carry some dependency due to their sequential order.

**Sequential Data:** the order of elements matters, and the past influences the future.

# Sequence data

---

Speech recognition



“The quick brown fox jumped over the lazy dog.”

Sentiment classification

“There is nothing to like in this movie.”



DNA sequence analysis

AGCCCCTGTGAGGAACTAG

AG**CCCCTGTGAGGAACTAG**


Video activity recognition



Running

# Why not a standard neural network?

---



Problems:

- 1) Inputs and outputs can have different lengths in different examples.
- 2) Doesn't share features learned across different positions of text.
- 3) High input dimension (e.g.  $T_x$ \* dimension of dictionary) => high dimensional problem, implying a high number of parameters to optimize.

# Time Series Forecasting





# Time Series

---

A sequence of random variables defined at fixed sampling intervals is sometimes referred as a discrete-time stochastic process.

This is also known as a time series.

A time series represents a variable that is indexed by time.





# Sampling in Time Series


---

The variable can be measured over time or at a fixed interval (sampling interval).

If the data represented in the time series is sampled, it is important that the interval is sufficiently small to allow for a proper approximation of the continuous signal when the data is interpolated.




# Time Series Decomposition



By breaking down the data into its constituent parts, such as trend, seasonality, and residual, analysts can gain valuable insights into the patterns and behavior of the series

# Time Series Decomposition

---



**Trend** - represents the long-term movement or direction of the data. It captures the persistent behavior of the series, which may be:


- ✓ Increase
- ✓ Decrease
- ✓ Stability.

Key characteristics:

- **Long-Term Behavior:** Trends are typically observed over a long period.
- **Smooth and Gradual Changes:** Trend components are often relatively smooth, exhibiting gradual changes over time.
- **Ignore Short-Term Fluctuations:** The trend component filters out short-term, noise-like fluctuations.

# Time Series Decomposition

---



**Seasonality** - regular, recurring patterns evidenced at fixed time intervals.


Patterns are often associated with calendar time:

- Daily
- Weekly
- Monthly
- Yearly

Key characteristics of the seasonality component include:

- **Regular Periodicity:** Seasonal patterns occur at consistent time intervals.
- **Predictable Patterns:** Seasonality can be anticipated and may exhibit similar patterns in different periods, making it valuable for forecasting.

# Time Series Decomposition



**Residual component** - the random fluctuations that cannot be explained by the trend or seasonality.


Residuals are the unexplained variability within the data, often reflecting the influence of unpredictable or external factors.

Key characteristics of the residual component include:

- **Random Variability:** Residuals are typically random and uncorrelated, lacking any systematic patterns.
- **Variations Not Accounted For:** The residual component represents the differences between observed values and the values predicted by the trend and seasonality components.

# Time Series Decomposition

---



Understanding the residual component is essential for assessing the goodness of fit of a time series model. If the residuals exhibit a pattern or structure, the model might not adequately capture all the underlying patterns in the data.

# Time Series Decomposition

---

$$x_t = f(T_t, S_t, E_t)$$

Where:

$T_t$  time series trend;

$S_t$  time series seasonality;

$E_t$  time series residuals.

Additive Decomposition

$$x_t = T_t + S_t + E_t$$

This model is appropriate when the magnitude of the seasonal fluctuations and residuals do not depend on the level of the series.

Multiplicative Decomposition

$$x_t = T_t \cdot S_t \cdot E_t$$

This model is appropriate when the magnitude of the seasonal and residual variations increases with the level of the time series — i.e., the fluctuations grow proportionally as the values increase.

# Time Series Forecasting

Time series forecasting is a technique that predicts future values in a sequence of data points measured over time.


It involves analyzing past data to identify **patterns, trends,** and **seasonality**, and then using **statistical** or **machine learning** models to make predictions about future events.

Forecasting allows the:

- Understanding of future trends
- Making informed decisions
- Optimizing processes

## Data Dependence:

Time series data is characterized by its sequential nature, where each observation is influenced by previous observations.




# Time Series Forecasting

Stationarity refers to the condition where the statistical properties of a time series, such as its mean, variance, and autocorrelation, remain consistent over time.

Stationarity is significant because:

- Simplified Analysis
- Reliable Forecasts
- Statistical Assumptions
- Trend and Seasonality Analysis



# Time Series Forecasting

Time series forecasting utilizes statistical or machine learning models:

- Classical/Statistical Models: Moving averages, exponential smoothing, ARIMA, SARIMA, etc.
- Machine Learning: Linear regression, XGBoost, Random Forest, etc.
- Deep Learning: RNN, LSTM.

Key Principles:

**Data Collection:** Collecting reliable and relevant historical data.

**Data Analysis:** Analyzing trends, patterns, and seasonality in the data.

**Model Selection:** Choosing the appropriate model based on the characteristics of the data and the forecasting goal.

**Evaluation:** Evaluating the performance of the model using appropriate metrics to ensure its reliability.



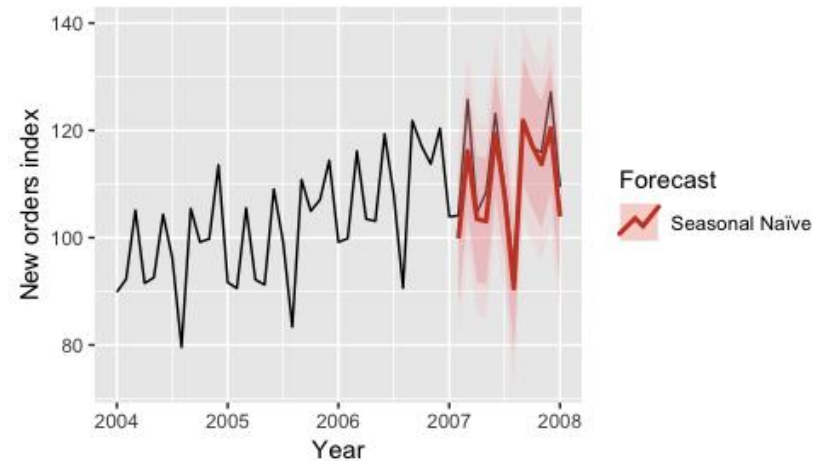

# Time Series Forecasting – Classical Models

1) Naïve model : the forecasts for every horizon (h) = last observed value:

$$\hat{Y}(t+h|t) = Y(t)$$

2) SNaïve (Seasonal Naïve) model: TS has a seasonal component with period of seasonality  $T$ :

$$\hat{Y}(t+h|t) = Y(t+h-T)$$




# Time Series Forecasting – Classical Methods

## 3) Exponential smoothing

Weighted average of past observations, weights decrease exponentially as go back in time.

$$\hat{Y}(t+h|t) = \alpha Y(t) + \alpha(1-\alpha)Y(t-1) + \alpha(1-\alpha)^2 Y(t-2) + \dots, 0 < \alpha < 1 \text{ (basic model)}$$



# Time Series Forecasting – Classical Methods

---

## 4) ARIMA - Auto-Regressive (Integrated) Moving Average

ARMA (stationary signals), ARIMA (both stationary & nonstationary)

Auto-Regressive model - linear combination of past values of TS.

Moving-Average model – lin. combination of past forecasting errors.

Integrated (use differences) :  $Y(t) = Y(t) - Y(t-1)$

ARIMA = AR + I + MA components

$$\hat{Y}(t+h|t) = a_1Y(t) + a_2Y(t-1) + \dots + a_pY(t-p) + b_1E(t) + b_2E(t-1) + \dots + b_qE(t-q)$$

## 5) SARIMA model (Seasonal ARIMA) extends ARIMA by adding a linear combination of seasonal past values and forecast errors.


ARIMA (p,d,q)

p=> autoregressive lags; q=> moving average lags;

d=> difference in the order

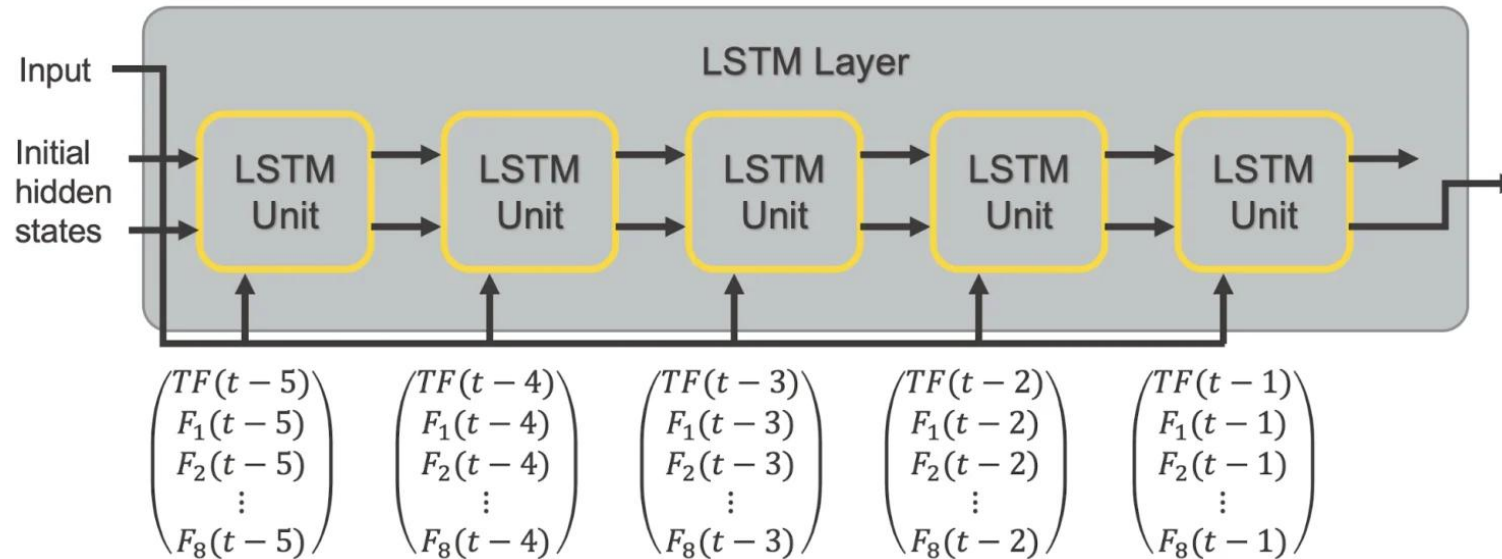
# Time Series Forecasting – ML Approach

Non-linear Auto Regressive Neural Network (NARN)  
Classical fully connected shallow (3 layers) ANN  
d- lag hyperparameter (time lag, model memory)



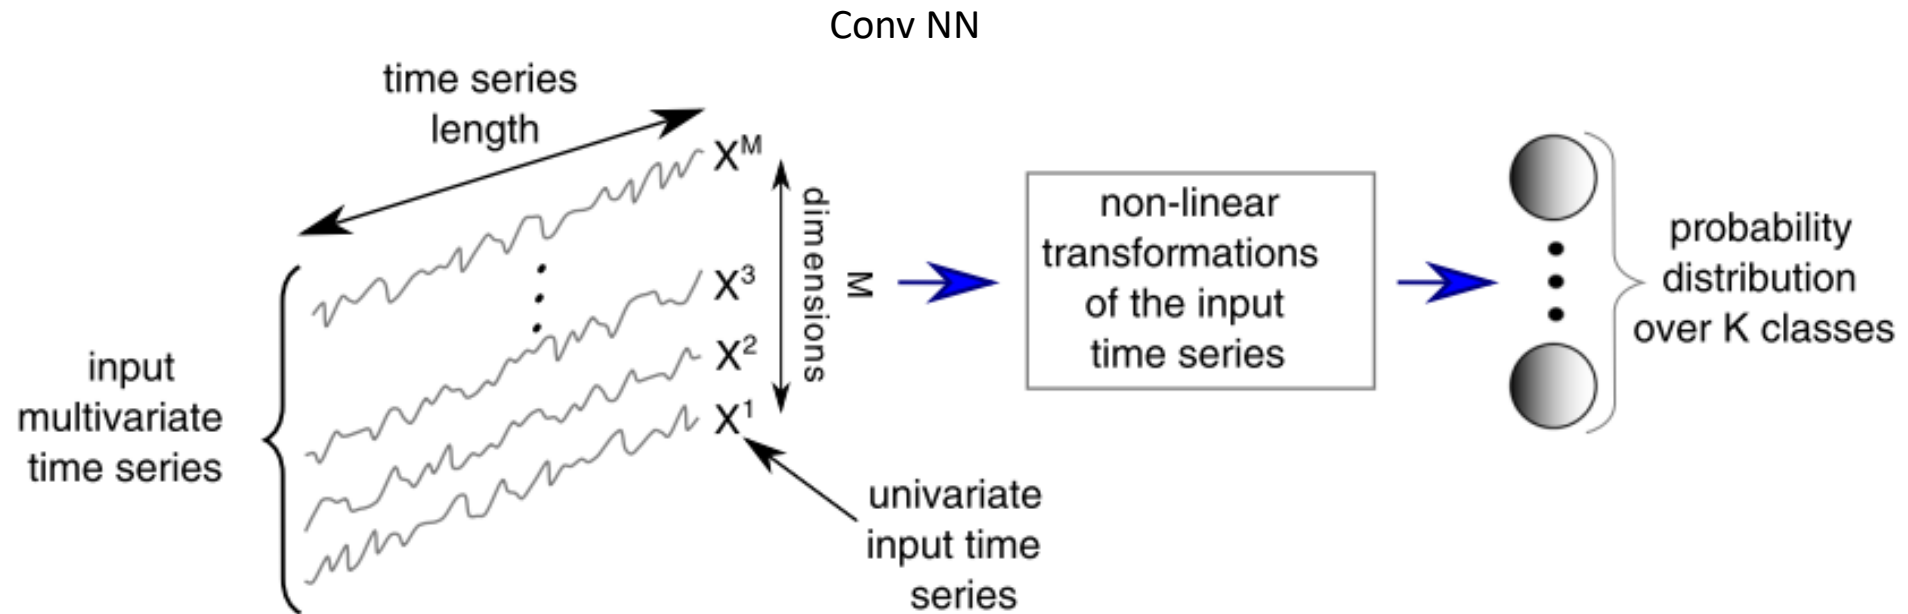
# Time Series Forecasting – DL Approach

Long Short Term Memory (LSTM) architecture  
or  
GRU (Gated Recurrent Unit) architecture



Multivariate prediction model  
Predict  $TF(t)$  – the output  
 $TF, F_1, F_2, \dots, F_8$  – predictors/features  
ex. lag  $d=5$

# Multivariate Time Series



# Time Series Forecasting

---

## Select the right model:

- Consider: data characteristics, the presence of trends or seasonality, and the forecasting requirements
- Begin with simple models like AR, MA, or ARMA and measure their performance. If the data shows clear patterns or dependencies, more complex models like ARIMA or SARIMA may be appropriate.
- Consider Seasonality: If the data shows seasonal patterns, models like SARIMA.
- Evaluate Performance: Use appropriate evaluation metrics and cross-validation techniques to compare the performance of different models.
- Consider any domain-specific knowledge or insights that can guide you in choosing a suitable model. Expert knowledge can help in identifying relevant variables, incorporating external factors, or applying specific modeling techniques.

## Metrics:

- Mean Absolute Error (MAE)
- Root Mean Squared Error (RMSE)
- Mean Absolute Percentage Error (MAPE)

A vertical bar on the left side of the slide, consisting of a wide red section and a thin blue section on its right edge.

# **Recurrent Neural Networks (RNN)**

# What is the difference between CNN and RNN?

---

**CNN** (Convolutional Neural Network) and **RNN** (Recurrent Neural Network) are both popular types of neural networks, but they serve different purposes and are designed to handle different types of data.

## CNN

- Primarily processes grid-like data, such as images
- Extracts local features using convolutional layers
- Excels at detecting spatial patterns and relationships
- No explicit memory of past inputs
- Treats each input independently
- Suitable for tasks like image recognition and computer vision
- Takes advantage of parallel processing
- Designed to capture spatial hierarchies and patterns
- Utilizes convolution and pooling layers
- Processes grid-like data with local spatial relationships
- Does not inherently capture temporal information
- Suitable for tasks where the order of data points is not significant
- Enables efficient computations on parallel hardware

## RNN

- Specifically designed for sequential data, like time series or natural language
- Captures temporal dependencies with recurrent connections
- Well-suited for capturing sequential patterns and long-term dependencies
- Has memory of previous inputs through hidden state
- Maintains information flow through time
- Commonly used in natural language processing, speech recognition, and time series analysis
- Sequential nature limits parallel processing capabilities
- Capable of modeling temporal hierarchies and patterns
- Employs recurrent connections for information persistence
- Processes sequential data with temporal dependencies
- Handles tasks where order of data points matters
- Sequential dependency limits parallel processing capabilities

# RNN


---

Recurrent Neural Networks (**RNNs**) are a type of neural network designed to process **sequential** data, such as text, speech, or time series, where **the order of elements is important**. They achieve this by **maintaining a hidden state** that **remembers previous inputs**, allowing them **to learn long-term dependencies and patterns**.

While traditional deep learning networks assume that inputs and outputs are independent of each other, the output of RNN depends on the prior elements within the sequence.

Another distinguishing characteristic of RNN is that they share parameters across each layer of the network. While feedforward networks have different weights across each node, RNN share the same weight parameter within each layer of the network.

The **core concept of RNNs** revolves around the **propagation of information** from one time step to the next. At each time step, the network updates its hidden state using both the current input and the previous hidden state. This recurrent structure imbues RNNs with the ability to model dependencies and relationships within sequential data.



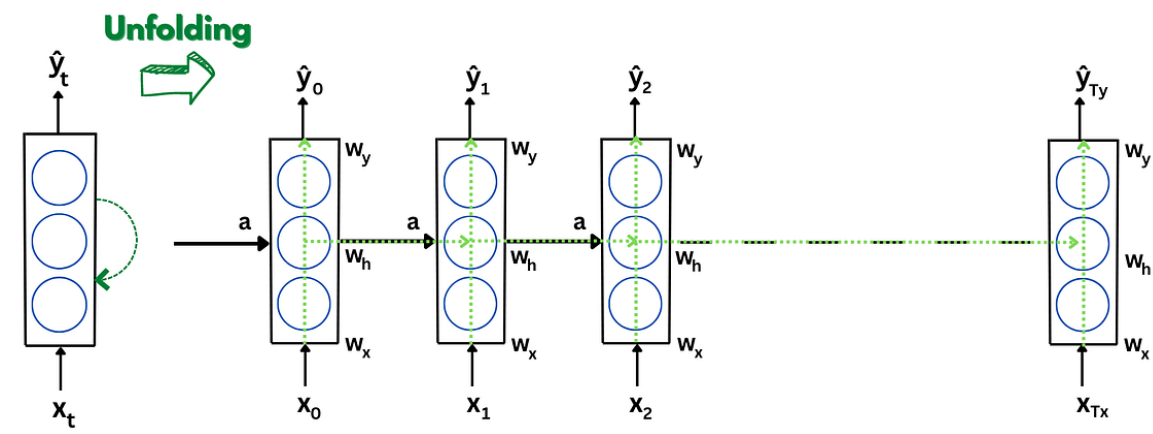
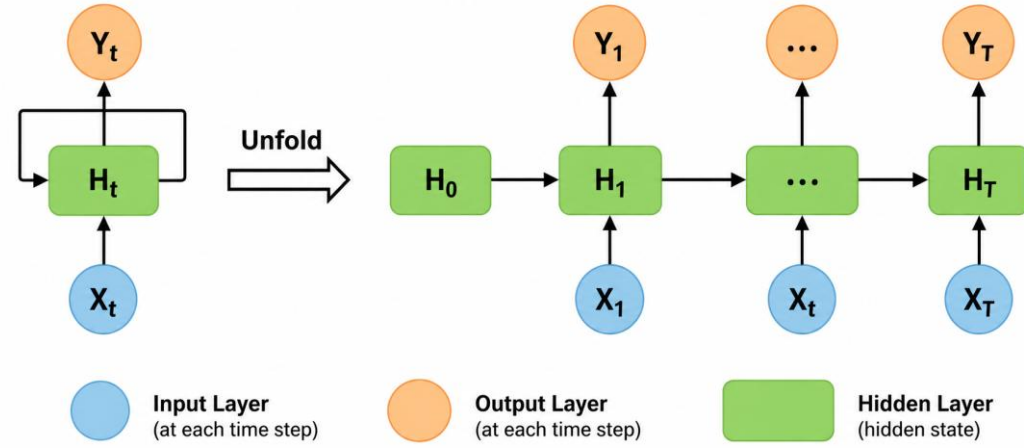
# Recurrent Neural Networks (RNN)

Read the sentence word by word (one time step per word).  
 Activation produced by 1<sup>st</sup> word is considered when read 2<sup>nd</sup> word.

Notation:

Inputs	$x_t$ ou $x^t$
Outputs	$y_t$ ou $y^t$
Hidden states	$a_t$ ou $a^t$
Initial state (at time 0) (usually vector of zeros)	$a_0$ ou $a^0$

The idea:  
 Previous results are passed as inputs, implying the transfer of context.




# Basic RNN unit

Takes  $x^{(t)}$  (current input) and  $a^{(t-1)}$  (activation from previous step) as inputs, and computes  $a^{(t)}$  (current activation).

$a^{(t)}$  is then used to predict  $y^{(t)}$  (current output) and passed forward to the next RNN unit (next time step).

$W_{aa}$ ,  $W_{ax}$ ,  $W_{ya}$ ,  $b_a$ ,  $b_y$  – the same RNN weights used in all time steps.



# Backpropagation through time

# BPTT

---

Backpropagation is a process by which a model learns from its mistakes (by updating weights) and tries to minimize the loss.

In RNN, the difference is that the backpropagation is THROUGH TIME, meaning the loss is backpropagated to each model and across the model because it is a SEQUENCE model.

Training RNNs involves a special kind of backpropagation called Backpropagation Through Time (BPTT)

Unlike traditional backpropagation, BPTT extends across time: it unfolds the entire sequence of data, applying backpropagation at each timestep.

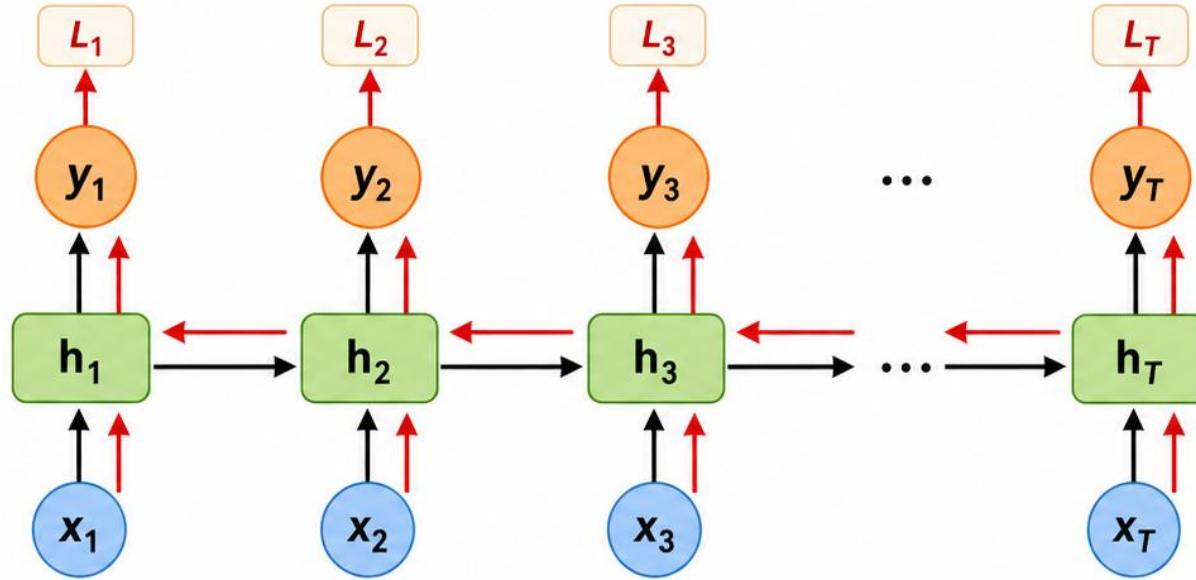
This method calculates gradients for each output, which are then used to adjust the weights and reduce the overall loss.

However, BPTT can be complex and resource-intensive, and it's prone to issues such as vanishing and exploding gradients, which can interfere with the network's ability to learn from data over longer sequences.

# Backpropagation Through Time (BPTT)


**1** Backpropagation updates weights to reduce loss.

**2** In RNNs, backpropagation happens through time because the model processes sequences.



**3** BPTT unfolds the RNN over the full sequence and computes gradients at each timestep.

**4** These gradients update the shared weights to minimize overall loss.



**Challenge:** BPTT can be computationally expensive and may suffer from vanishing or exploding gradients, especially for long sequences.

**Vanishing Gradients**  
Gradients shrink over time  
(harder for early steps to learn)

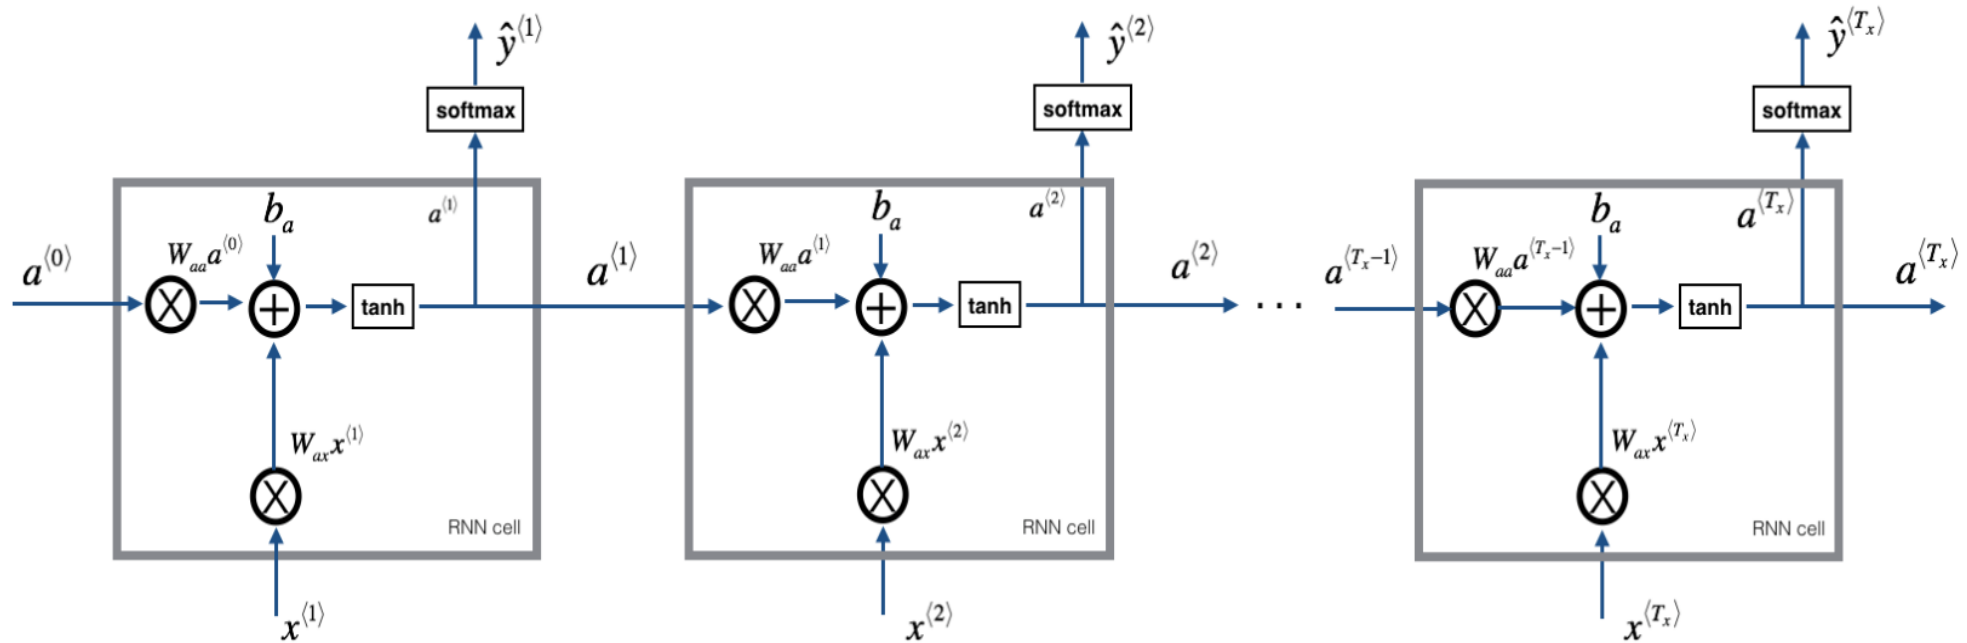
**Exploding Gradients**  
Gradients grow over time  
(unstable updates)

# RNN forward pass

RNN is a sequence of basic RNN cells.

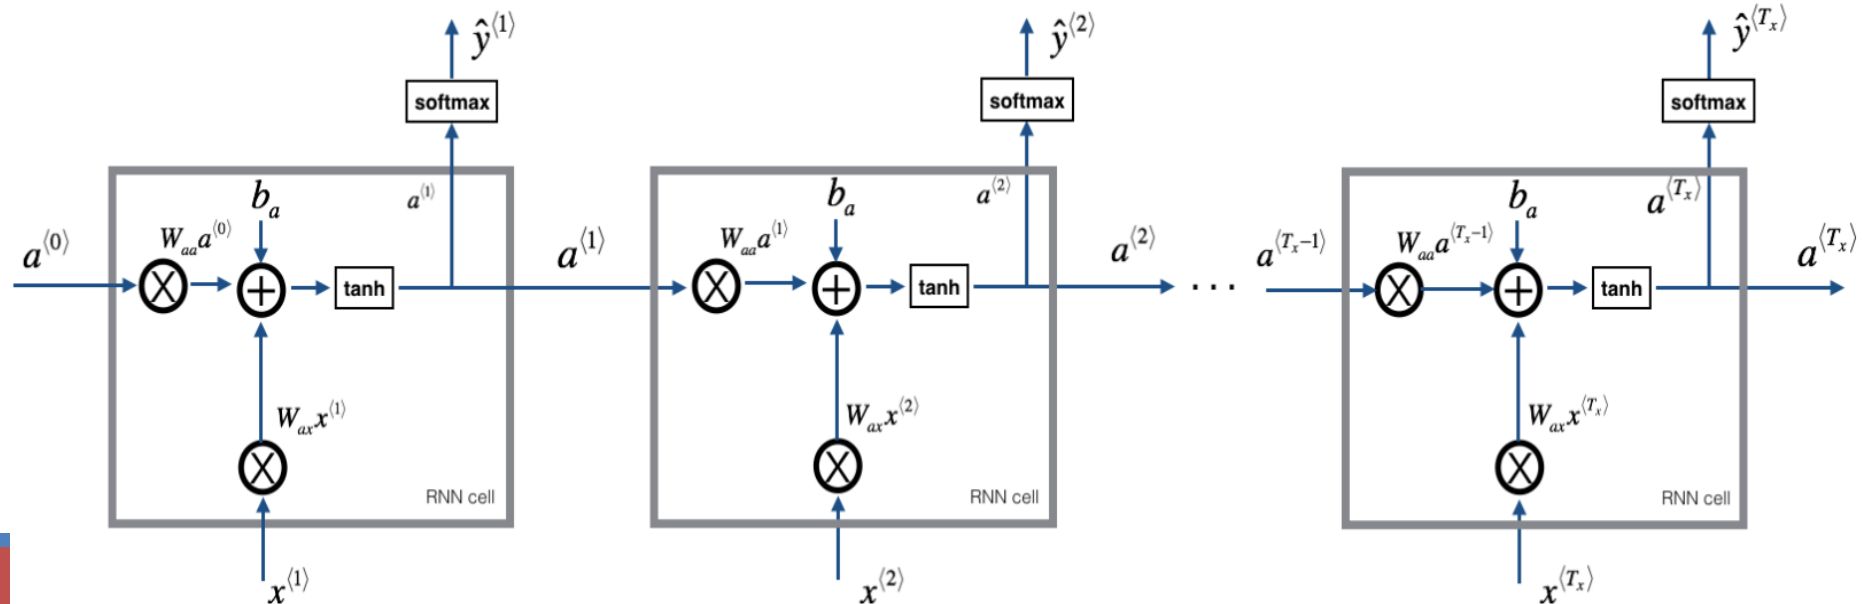
If input sequence is  $x=[x^{<1>}, x^{<2>}, \dots, x^{<T_x>}]$ , the RNN cell is copied  $T_x$  times.

RNN outputs  $y=[y^{<1>}, y^{<2>}, \dots, y^{<T_x>}]$



# BPTT

1. Forward propagation (for-prop): from time step 1 to time step  $T_x$  (from left to right), compute RNN predictions.
2. Compute the cost (loss) function for each output  $y^{(t)}$  ( $J^{(t)}$ )
3. Compute the sum of all cost (loss) functions ( $J_{all}$ )
4. Backward propagation (back-prop): to update RNN parameters compute the gradients of  $J_{all}$ , starting from the last time step  $y^{(T_x)}$  and going back through the previous time steps down to the first time step  $y^{(1)}$  (from right to left) (This is called *Backpropagation Through Time (BPTT)*.)
5. The programming framework usually computes BPTT automatically.



# RNN – vanishing/exploding gradients

---

RNN works well when each output  $y^{(t)}$  can be estimated using "local" context (i.e. information from inputs  $x^{(t')}$  where  $t'$  is not too far from  $t$ ).

Why ? – because RNN's suffers from vanishing gradient (like Deep NN).

Vanishing gradient - gradient values become very small. Layers that get small gradients stops learning. Those are usually the earlier layers in the sequence model. Because these layers don't learn, RNN's can forget what it seen in longer sequences, thus having a short-term memory.

Standard RNN is not very good in capturing long-range dependencies, ex.:

*The girl that entered the coffee shop with many friends was happy.*

*The girls that entered the coffee shop with many friends were happy.*

Need to remember "girl-was" or "girls-were".

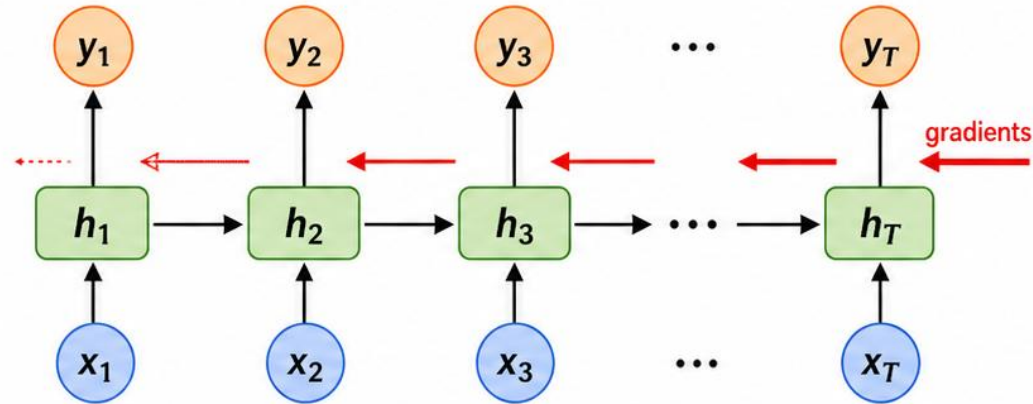
Long Short-Term Memory \* (LSTM) and Gated Recurrent Units (GRU): Solution for vanishing gradients.

# Standard RNNs: Short-Term Memory and the Vanishing Gradient Problem

## 1 When RNNs work well

RNNs work well when each output  $y^{(t)}$  depends mainly on local context.

That means the relevant inputs  $x^{(t')}$  are not too far away in time from  $t$ .



**Vanishing gradient:** gradients shrink as they are backpropagated through long sequences.

Earlier states learn less

RNN develops short-term memory

## 2 Why is this a problem?

- Very small gradients mean early steps stop learning effectively.
- Standard RNNs struggle to capture long-range dependencies.
- They may forget information seen many time steps earlier.

## 3 Example of long-range dependency


The **girl** that entered the coffee shop with many friends **was** happy.

The **girls** that entered the coffee shop with many friends **were** happy.

The model must remember: **girl** → **was**, **girls** → **were**

## 4 Solution

- ✓ LSTM (Long Short-Term Memory)
- ✓ GRU (Gated Recurrent Unit)
- ✓ Designed to reduce the vanishing gradient problem
- ✓ Better at capturing long-range dependencies



Gates control what to remember, what to forget, and what to output.



Standard RNNs are effective for local dependencies, but LSTM and GRU are better for longer sequences.

# RNN Challenges

---

## **Exploding gradients: When many gradient values are $>1$**

Occurs when large error gradients accumulate and there are very large updates to neural network weights.

Gradients are used during training to update the network weights.

When the magnitudes of the gradients accumulate, an unstable network is likely to occur, which can cause poor prediction results.

## **Vanishing gradients: When many gradient values are $<1$**


If the gradients are very small or zero, then little to no training occurs, leading to poor performance.

This also leads to capturing short-term dependencies instead of long-term dependencies.

# RNN Challenges

## Potential Solutions:

1. **Activation Functions:** Using ReLU prevents gradients from shrinking when  $x > 0$
2. **Parameter Initialization:** Initialize weights to identity matrix and biases to zero, preventing weights from shrinking to zero
3. **Gated Cells:** In the operation box (blue), use some logic function (i.e. gated cells).



Type of gate	Role	Used in
Update gate	How much past should matter now?	GRU, LSTM
Relevance gate	Drop previous information?	GRU, LSTM
Forget gate	Erase a cell or not?	LSTM
Output gate	How much to reveal of a cell?	LSTM

# Long-Short Term Memory (LSTM)

# LSTM

---

**LSTMs** are a type of RNN architecture built to address the limitations of traditional RNNs, especially the vanishing gradient problem.

**LSTMs have a dedicated memory cell that can store information for extended periods.**

Their effectiveness comes from the gate mechanisms they utilize:

- the input gate
- forget gate
- output gate

These **gates control the information flow**, determining what to retain or remove from memory and what to output at each step.

These gating mechanisms allow LSTMs to properly manage long-term dependencies.

# LSTMs: Managing Long-Term Dependencies

## 1 Why LSTMs?


LSTMs are a type of RNN designed to overcome limitations of standard RNNs.

They are especially useful for reducing the vanishing gradient problem.

## 2 Dedicated memory cell

LSTMs include a memory cell that can store information for long periods.

This helps the network preserve relevant context over time.



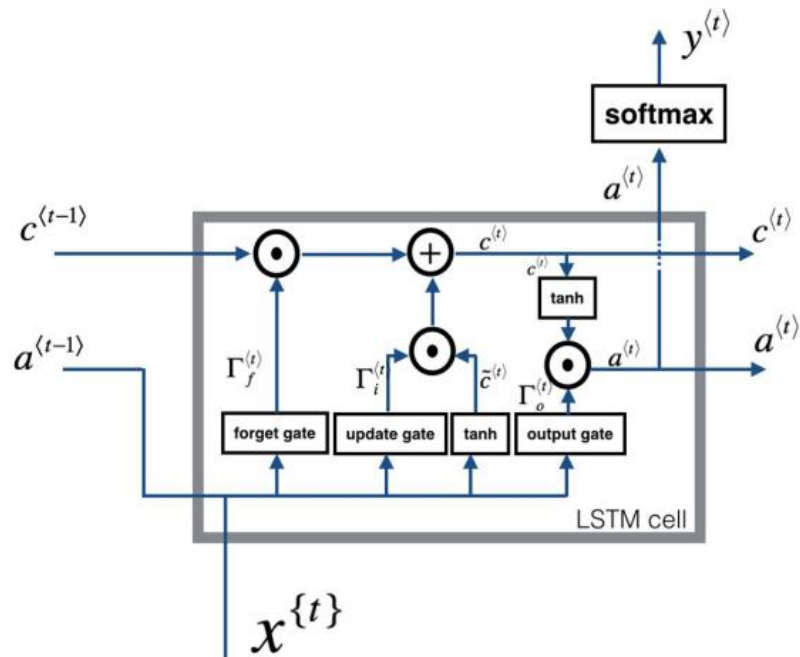
## 3 Gate mechanisms

- **Forget gate:** decides what information to remove from memory
- **Input gate:** decides what new information to store
- **Output gate:** decides what information to reveal at each step

## 4 Key takeaway

By controlling the flow of information, LSTMs can capture long-term dependencies more effectively than standard RNNs.

# LSTM



$$\tilde{c}^{<t>} = \tanh(W_c [a^{<t-1>}, x^{<t>}] + b_c) - \text{candidate}$$

$$\Gamma_u = \sigma(W_u [a^{<t-1>}, x^{<t>}] + b_u) - \text{update gate}$$

$$\Gamma_f = \sigma(W_f [a^{<t-1>}, x^{<t>}] + b_f) - \text{forget gate}$$

$$\Gamma_o = \sigma(W_o [a^{<t-1>}, x^{<t>}] + b_o) - \text{output gate}$$

$$c^{<t>} = \Gamma_u * \tilde{c}^{<t>} + \Gamma_f * c^{<t-1>} - \text{update memory cell}$$

$$a^{<t>} = \Gamma_o * \tanh(c^{<t>}) - \text{activation}$$

LSTM outputs both activation value and memory cell

Gates ( $\Gamma_f$ ,  $\Gamma_u$ ,  $\Gamma_o$ ) contain sigmoid activations, so values between 0 and 1.

Forget gate decides what is relevant to keep from prior steps.

Information from the previous hidden state and from the current input is passed through the sigmoid function. Output values between 0 and 1.

Closer to 0 means to forget, closer to 1 means to keep.

**Update gate** decides what inf is relevant to add from the current step.

**Output gate** determines the next hidden state based on the updated cell state. It filters the information that the LSTM will output based on the updated cell state.

# LSTM – Internal States

## Cell State:

Acts as the **memory** of the LSTM.

This runs straight down the entire chain of the LSTM, with minimal modifications.

Only interacts linearly with gates (forget, update, output).

It's the core differentiator in LSTMs that allows them to maintain and control long-term dependencies.

## Hidden State:

The output of the LSTM at each time step.

Derived from the cell state, modulated by the output gate.

Encodes short-term information for the current moment in the sequence.


Captures what the LSTM "knows" at that point.

**Gates**, the special neural layers that control the flow of information:

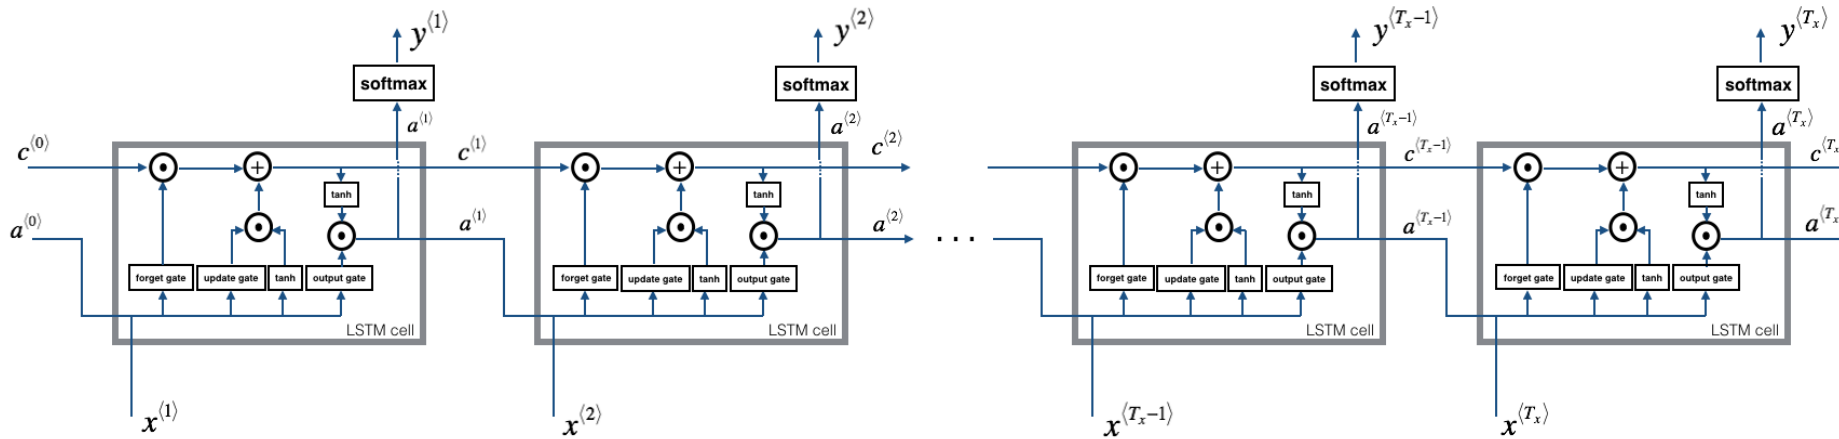
**Forget Gate:** Decides what to discard from the cell state.

**Update Gate:** Decides what new information to store.

**Output Gate:** Decides what to output from the cell state.



# LSTM network




An LSTM network is a temporal sequence of LSTM units.

There is a line at the top that shows how LSTM can memorise and pass certain values  $c^{<t>}$  through several temporal steps.

Other versions: the gate's values also depend on  $c^{<t-1>}$ . (peephole connection)

The gates are vectors (e.g., 100 elements ).

# Deep LSTM



# **Gated Recurrent Unit (GRU)**

# Gated Recurrent Units (GRUs)

---

GRUs are a streamlined variant of RNNs that modify the LSTM design to achieve similar results with less complexity.

GRUs merge the input and forget gates into a single “update gate” and combine the cell state and hidden state. This simpler structure reduces the model’s complexity, which can lead to quicker computations and lower memory usage.

GRUs are often favored over LSTMs for smaller datasets where LSTMs might overfit due to their complexity, or when faster training is essential. They perform well on many tasks that don’t involve very long-term dependencies.

# GRUs: Simpler Gated Recurrent Units

## 1 Why GRUs?

GRUs are a streamlined variant of RNNs designed to achieve results similar to LSTMs with less complexity.

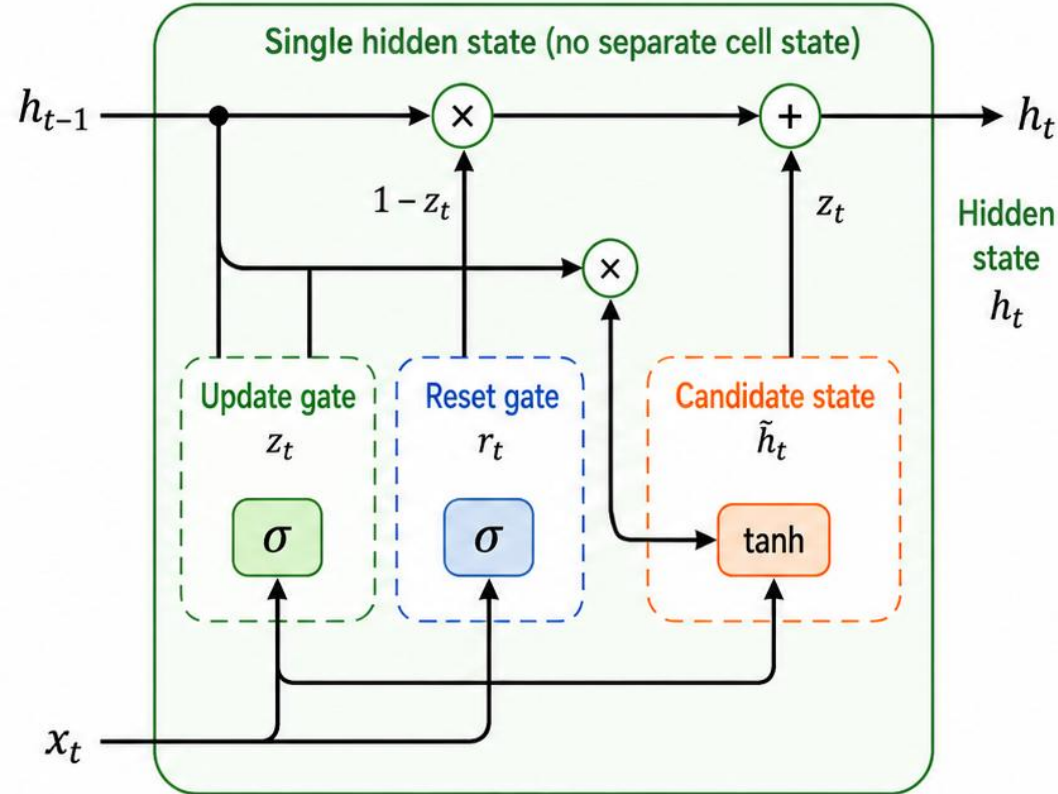
They simplify the LSTM architecture while still handling sequence dependencies effectively.

## 2 Simpler structure

GRUs combine the cell state and hidden state into a single state.

They merge the input and forget mechanisms into an update gate, reducing the number of parameters.

This usually means faster computation and lower memory usage.



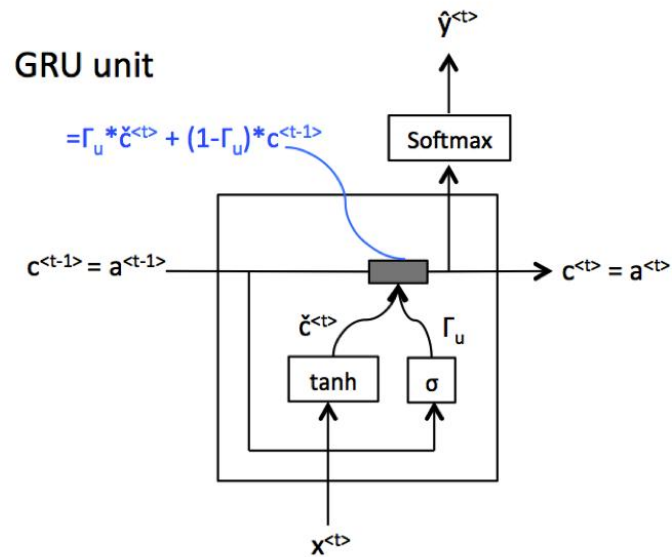
## 3 When are GRUs useful?

- Often preferred for smaller datasets
- Can be less prone to overfitting than more complex LSTMs
- Useful when faster training is important
- Perform well when very long-term dependencies are not critical

## 4 Key takeaway

GRUs offer a simpler gated architecture than LSTMs, often giving competitive performance with faster training and lower computational cost.

# GRU basic architecture



c - memory cell

$c^{<t>} = a^{<t>}$  (for LSTM they are different)

GRU outputs activation value = memory cell

$\tilde{c}^{<t>} = \tanh(w_c [c^{<t-1>}, x^{<t>}] + b_c)$  - candidate

$\Gamma_u = \sigma(w_u [c^{<t-1>}, x^{<t>}] + b_u)$  - update gate

$c^{<t>} = \Gamma_u * \tilde{c}^{<t>} + (1 - \Gamma_u) * c^{<t-1>}$  - update memory cell

$\Gamma_u \Rightarrow$  values between (0,1)

if  $\Gamma_u \Rightarrow 0$ , don't update  $c^{<t>}$ , if  $\Gamma_u \Rightarrow 1$ , update  $c^{<t>}$ .

$\Gamma_u=1 \quad \Gamma_u=0 \quad \Gamma_u=0 \quad \Gamma_u=0 \quad \dots \dots \Gamma_u=1$   
 $c^{<t>}= 1$ , keeps the same value of  $c^{<t>}$  long time .....

Ex. The girl that entered the coffee shop with many friends was happy.

**Memory cell c** memorises if the girl was singular or plural even if girl-was are separated with many many words. Efficient against vanishing gradients.

$c^{<t>}$  and  $\Gamma_u$  are vectors (e.g. 100 elements/bits to update)


# GRU vs LSTM

## Similarities:

- Both are advanced RNNs designed to handle sequential data.
- Aim to solve the vanishing gradient problem in traditional RNNs.
- Use gating mechanisms to control the flow of information.
- Capable of learning long-term dependencies.
- Commonly used in NLP, time series forecasting, and other sequence tasks.

Feature	LSTM	GRU
Architecture	3 gates (forget, update(or input), output)	2 gates (reset, update)
Cell State	Maintains a separate cell state (Ct)	No separate cell state, only ht
Complexity	More parameters, heavier computation	Fewer parameters, faster to train
Performance	May perform better on very long sequences	Often comparable, simpler to tune
Memory Control	Explicit memory via cell state	Merged memory & output in ht

# GRU vs LSTM



## Summary:

Use LSTM when fine control of long-term memory is needed.

Use GRU for faster training and when less memory overhead is desired.

### Advantages of RNNs:

- Handle sequential data: They are well-suited for processing data where the order of elements is crucial.
- Learn long-term dependencies: Their internal memory allows them to learn patterns over longer sequences.

### Limitations of RNNs:

- Vanishing Gradient Problem: RNNs can struggle to learn long-range dependencies due to the vanishing gradient problem, where the gradient of the loss function becomes too small to update the weights effectively.
- Parallelization Challenges: Due to their sequential nature, RNNs can be difficult to parallelize, which can slow down training.

### Advantages of LSTMs

The advantages of the LSTM are similar to RNNs, with the main benefit of their inner capability to capture patterns in the long-term and short-term of a sequence.

### Limitations of LSTMs

- Due to their more complex structure, LSTMs are computationally more expensive, leading to longer training times.
- As the LSTM also uses the backpropagation in time algorithm to update the weights, the LSTM suffers from the disadvantages of the backpropagation (e.g., dead ReLU elements, exploding gradients).

### Advantages of GRUs

- Due to the simpler architecture compared to LSTMs (i.e., two instead of three gates and one state instead of two), GRUs are computationally more efficient and faster to train as they need less memory.
- GRUs have proven to be more efficient for smaller sequences.

### Limitations of GRUs

As GRUs do not have a separate hidden and cell state, they might not be able to consider observations as far into the past as the LSTM.

Similar to the RNN and LSTM, the GRU also might suffer from the disadvantages of the backpropagation in time to update the weights.