



SNS COLLEGE OF TECHNOLOGY

(AN AUTONOMOUS INSTITUTION)

Approved by AICTE & Affiliated to Anna University
Accredited by NBA & Accredited by NAAC with 'A+' Grade,
Recognized by UGC saravanampatti (post), Coimbatore-641035.



Department of Biomedical Engineering

**Course Name: 19ECT303 & Artificial Intelligence and machine
learning**

III Year : V Semester

Unit V-DEEP LEARNING

Topic : Recurrent Neural Networks

19ECT303/Artificial Intelligence and Machine learning/Unit 5/Mr.
Karthik G. L. /AP/BME



What are recurrent neural networks?

A recurrent neural network (RNN) is a type of artificial neural network which uses sequential data or time series data. These deep learning algorithms are commonly used for ordinal or temporal problems, such as language translation, natural language processing (NLP), speech recognition, and image captioning; they are incorporated into popular applications such as Siri, voice search, and Google Translate.



Traditional NN Vs RNN



Like feedforward and convolutional neural networks (CNNs), recurrent neural networks utilize training data to learn. They are distinguished by their “memory” as they take information from prior inputs to influence the current input and output. While traditional deep neural networks assume that inputs and outputs are independent of each other, the output of recurrent neural networks depend on the prior elements within the sequence. While future events would also be helpful in determining the output of a given sequence, unidirectional recurrent neural networks cannot account for these events in their predictions.



Example of RNN



- Let's take an idiom, such as “feeling under the weather”, which is commonly used when someone is ill, to aid us in the explanation of RNNs.
- In order for the idiom to make sense, it needs to be expressed in that specific order.
- As a result, recurrent networks need to account for the position of each word in the idiom and they use that information to predict the next word in the sequence.



Weights in RNN



Another distinguishing characteristic of recurrent networks is that they share parameters across each layer of the network. While feedforward networks have different weights across each node, recurrent neural networks share the same weight parameter within each layer of the network. That said, these weights are still adjusted in the through the processes of backpropagation and gradient descent to facilitate reinforcement learning.



BPTT

- Recurrent neural networks leverage backpropagation through time (BPTT) algorithm to determine the gradients, which is slightly different from traditional backpropagation as it is specific to sequence data.
- The principles of BPTT are the same as traditional backpropagation, where the model trains itself by calculating errors from its output layer to its input layer.
- These calculations allow us to adjust and fit the parameters of the model appropriately. BPTT differs from the traditional approach in that BPTT sums errors at each time step whereas feedforward networks do not need to sum errors as they do not share parameters across each layer.

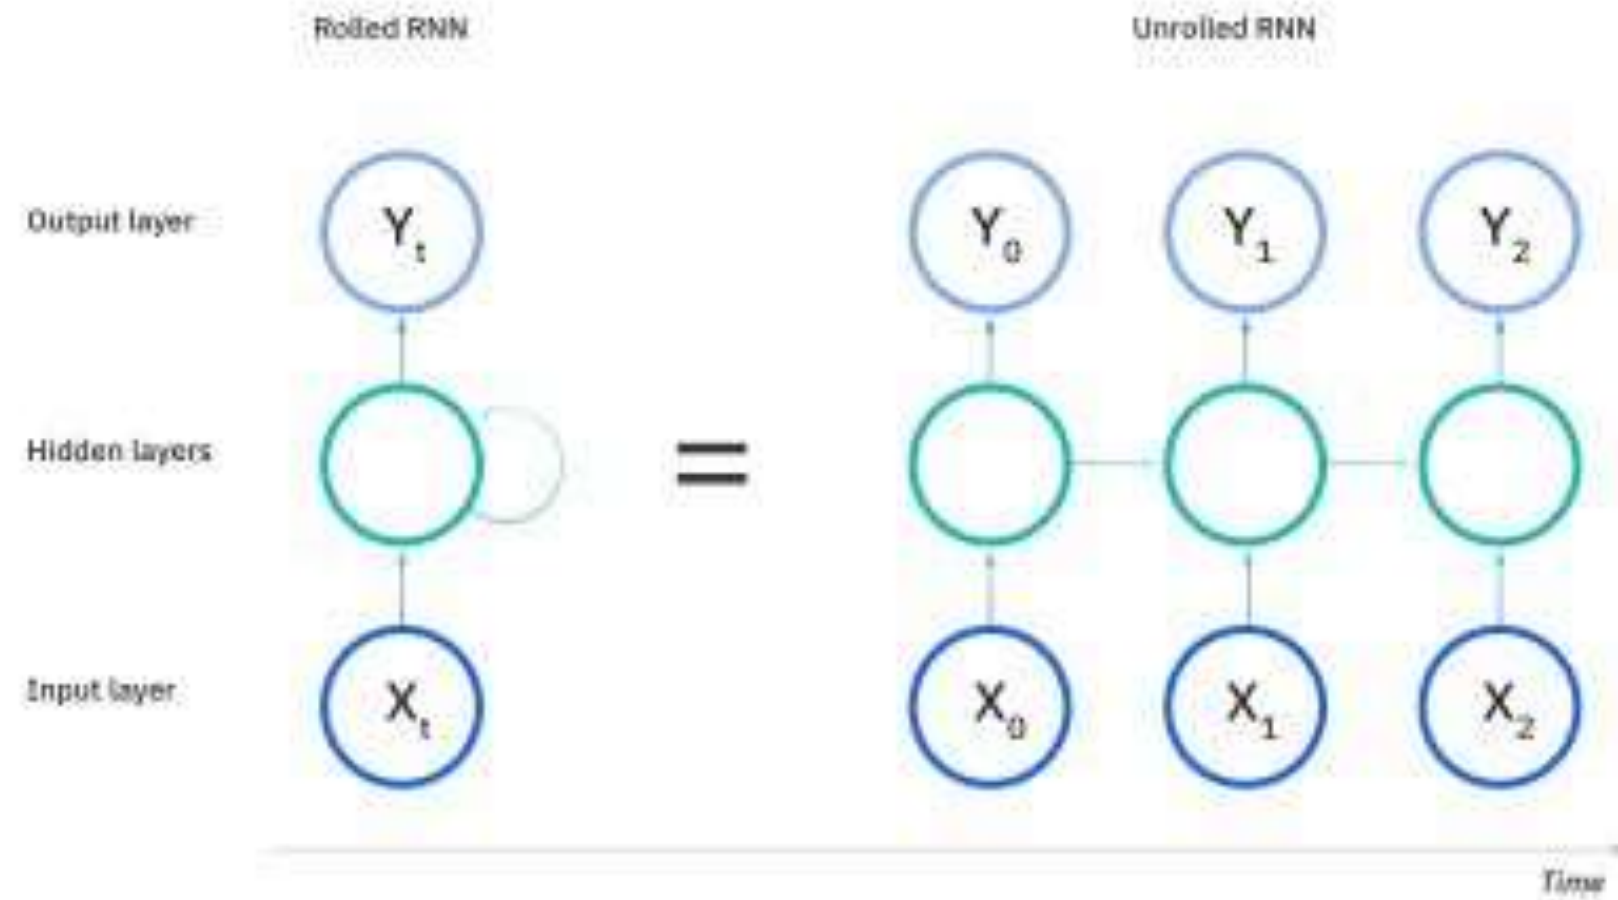
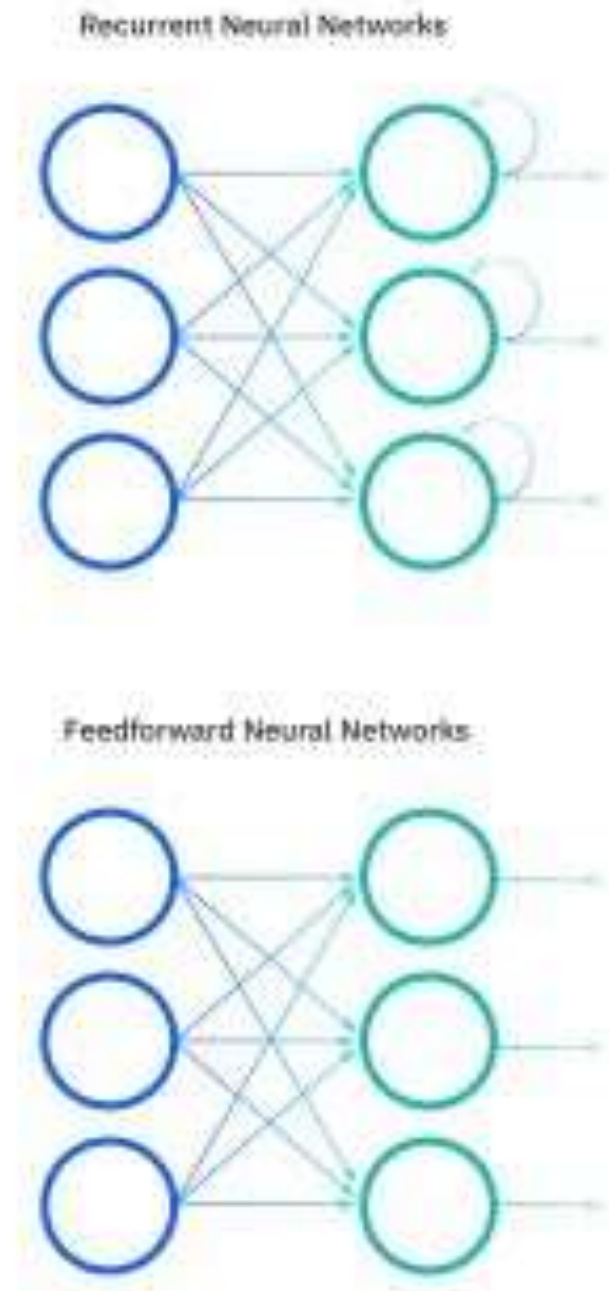


Exploding gradient and Vanishing gradient problem

- Through this process, RNNs tend to run into two problems, known as exploding gradients and vanishing gradients. These issues are defined by the size of the gradient, which is the slope of the loss function along the error curve. When the gradient is too small, it continues to become smaller, updating the weight parameters until they become insignificant—i.e. 0. When that occurs, the algorithm is no longer learning. **Exploding gradients** occur when the gradient is too large, creating an unstable model. In this case, the model weights will grow too large, and they will eventually be represented as NaN. One solution to these issues is to reduce the number of hidden layers within the neural network, eliminating some of the complexity in the RNN model.



RNN Architecture





Types of RNN



Feedforward networks map one input to one output, and while we've visualized recurrent neural networks in this way in the above diagrams, they do not actually have this constraint. Instead, their inputs and outputs can vary in length, and different types of RNNs are used for different use cases, such as music generation, sentiment classification, and machine translation.




Different configurations of RNN

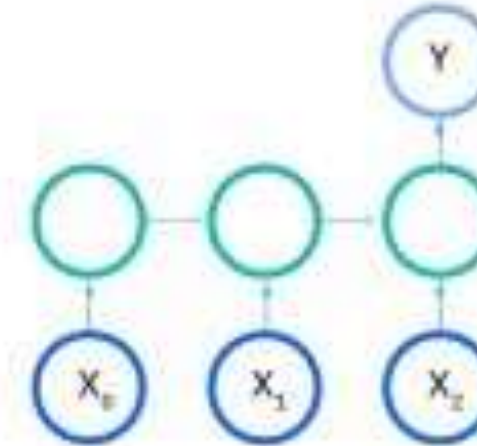
One-to-one




One-to-many




Many-to-one



Many-to-many



Many-to-many







Common activation functions



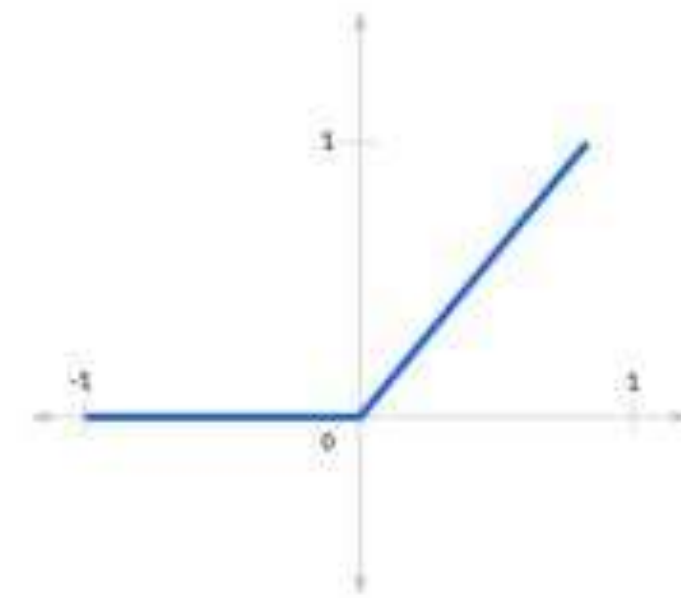
$$g(x) = \frac{1}{1 + e^{-x}}$$



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



$$g(x) = \max(0, x)$$





Variant RNN architectures



Bidirectional recurrent neural networks (BRNN): These are a variant network architecture of RNNs. While unidirectional RNNs can only draw from previous inputs to make predictions about the current state, bidirectional RNNs pull in future data to improve the accuracy of it. If we return to the example of “feeling under the weather” earlier in this article, the model can better predict that the second word in that phrase is “under” if it knew that the last word in the sequence is “weather.”



Variant RNN architectures

Long short-term memory (LSTM): This is a popular RNN architecture, which was introduced by Sepp Hochreiter and Juergen Schmidhuber as a solution to vanishing gradient problem. They work to address the problem of long-term dependencies. That is, if the previous state that is influencing the current prediction is not in the recent past, the RNN model may not be able to accurately predict the current state.



Variant RNN architectures

As an example, let's say we wanted to predict the italicized words in following, "Alice is allergic to nuts. She can't eat *peanut butter*." The context of a nut allergy can help us anticipate that the food that cannot be eaten contains nuts. However, if that context was a few sentences prior, then it would make it difficult, or even impossible, for the RNN to connect the information. To remedy this, LSTMs have "cells" in the hidden layers of the neural network, which have three gates—an input gate, an output gate, and a forget gate. These gates control the flow of information which is needed to predict the output in the network. For example, if gender pronouns, such as "she", was repeated multiple times in prior sentences, you may exclude that from the cell state.



Variant RNN architectures

Gated recurrent units (GRUs): This RNN variant is similar to the LSTMs as it also works to address the short-term memory problem of RNN models. Instead of using a “cell state” to regulate information, it uses hidden states, and instead of three gates, it has two—a reset gate and an update gate. Similar to the gates within LSTMs, the reset and update gates control how much and which information to retain.



THANK YOU !!!
HAPPINESS ISN'T OUTSIDE,
ITS WITHIN