


The background is a dark blue gradient with a subtle pattern of small white dots. On the left side, there are several overlapping circular elements. A large, semi-circular scale with tick marks and numbers (140, 150, 160, 170, 180, 190, 200, 210, 220, 230, 240, 250, 260) is visible. Other circles contain dashed lines and arrows, suggesting a sense of motion or a technical diagram.

RECURRENT NEURAL NETWORKS

INTRO TO RNN MODEL

- Recurrent neural networks (RNN) are the state-of-the-art algorithm for sequential data and are used by Apple's Siri and Google's voice search.
- It is the first algorithm that remembers its input, due to an internal memory, which makes it perfectly suited for machine learning problems that involve sequential data.
- It is one of the algorithms behind the scenes of the amazing achievements seen in deep learning over the past few years.

HISTORY


- Like many other deep learning algorithms, recurrent neural networks are relatively old. They were initially created in the 1980's, but only in recent years have we seen their true potential.
- An increase in computational power along with the massive amounts of data that we now must work with, and the invention of long short-term memory (LSTM) in the 1990s, has really brought RNNs to the foreground.

LONG SHORT-TERM MEMORY (LSTM)

- Long short-term memory networks (LSTMs) are an extension for recurrent neural networks, which basically extends the memory. Therefore, it is well suited to learn from important experiences that have very long-time lags in between.
- The units of an LSTM are used as building units for the layers of a RNN, often called an LSTM network.
- LSTMs enable RNNs to remember inputs over a long period of time. This is because LSTMs contain information in a memory, much like the memory of a computer. The LSTM can read, write and delete information from its memory.

LSTM

In an LSTM you have three gates: input, forget and output gate. These gates determine whether to let new input in (input gate), delete the information because it isn't important (forget gate), or let it impact the output at the current timestep (output gate). Below is an illustration of a RNN with its three gates:



RNN + LSTM

- A usual RNN has a short-term memory. In combination with a LSTM, they also have a long-term memory.
- Because of their internal memory, RNN's can remember important things about the input they received, which allows them to be very precise in predicting what's coming next.
- Therefore, they're the preferred algorithm for sequential data like time series, speech, text, financial data, audio, video, weather and much more.
- Recurrent neural networks can form a much deeper understanding of a sequence and its context compared to other algorithms.

WHEN TO USE RNN?


- “Whenever there is a sequence of data and that temporal dynamics that connects the data is more important than the spatial content of each individual frame.” – Lex Fridman (MIT)
- Since RNNs are being used in the software behind Siri and Google Translate, recurrent neural networks show up a lot in everyday life.

HOW RECURRENT NEURAL NETWORKS WORK


- To understand RNNs properly, you'll need a working knowledge of "normal" feed-forward neural networks and sequential data.
- Sequential data is basically just ordered data in which related things follow each other.
- Examples are financial data or the DNA sequence.
- The most popular type of sequential data is perhaps time series data, which is just a series of data points that are listed in time order.

RNN VS. FEED-FORWARD NEURAL NETWORKS

- In a feed-forward neural network, the information only moves in one direction — from the input layer, through the hidden layers, to the output layer. The information moves straight through the network and never touches a node twice.
- Feed-forward neural networks have no memory of the input they receive and are bad at predicting what's coming next. Because a feed-forward network only considers the current input, it has no notion of order in time. It simply can't remember anything about what happened in the past except its training.
- In a RNN the information cycles through a loop. When it decides, it considers the current input and what it has learned from the inputs it received previously.



Recurrent Neural Network




Feed-Forward Neural Network

A SCENARIO

Another good way to illustrate the concept of a recurrent neural network's memory is to explain it with an example:

Imagine you have a normal feed-forward neural network and give it the word "neuron" as an input, and it processes the word character by character. By the time it reaches the character "r," it has already forgotten about "n," "e" and "u," which makes it almost impossible for this type of neural network to predict which character would come next.

A recurrent neural network, however, can remember those characters because of its internal memory. It produces output, copies that output and loops it back into the network.



THE END