Recurrent Neural Networks

Recurrent Neural Networks (RNNs) are a type of artificial neural network designed to process sequential data, like text or time series, by maintaining a memory of previous inputs through feedback loops, allowing them to learn patterns and dependencies in sequential data.

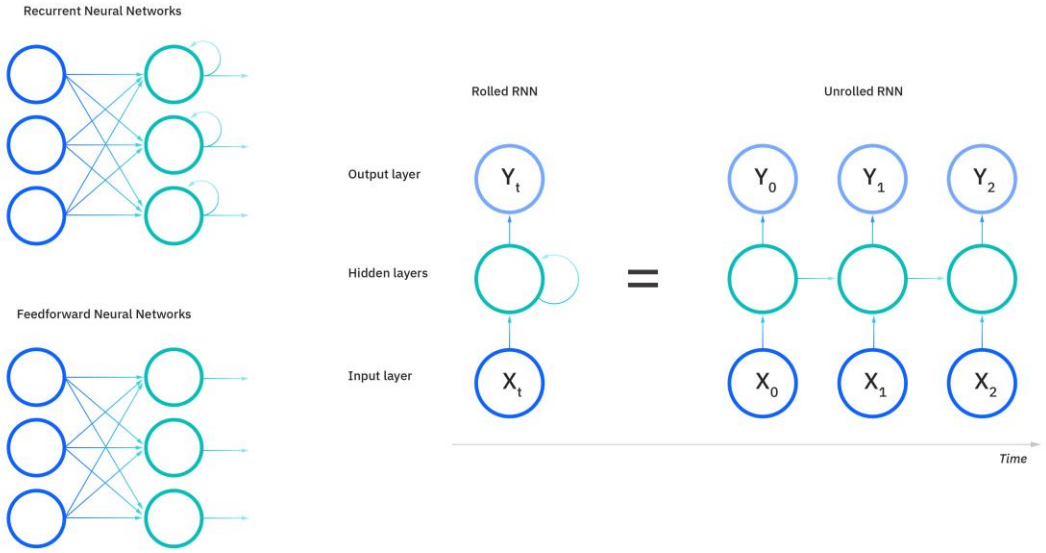

Here's a more detailed explanation:

- **Sequential Data Processing:**

Unlike traditional feedforward neural networks, RNNs are specifically designed to handle data where the order of elements matters, such as sentences, speech, or stock prices.

- **Feedback Loops:**

RNNs have connections that allow information to persist through time, creating a "memory" of past inputs that influences the current output.

- **Hidden State:**

RNNs maintain an internal state (hidden state) that captures information about the sequence processed so far, enabling them to make predictions based on context.

- **Applications:**

RNNs are widely used in various applications, including:

- **Natural Language Processing (NLP):** Machine translation, text generation, sentiment analysis, and language modeling.
- **Speech Recognition:** Understanding and transcribing spoken language.
- **Time Series Analysis:** Forecasting stock prices, predicting weather patterns, and analyzing sensor data.
- **Example:**

  Imagine an RNN trained on Shakespeare's works. It can learn the patterns and structures of his writing style and then generate new text that resembles his prose.

- **Limitations:**

  RNNs can struggle with very long sequences due to the vanishing gradient problem, where information from distant time steps can become lost.

- **Variations:**

  To address the limitations of standard RNNs, more advanced architectures like Long Short-Term Memory (LSTM) and Gated Recurrent Unit (GRU) networks have been developed.