MASTERY IN RECURRENT NEURAL NETWORK

Introduction

- Start your RNN journey today

What You'll Learn

This three day course includes **code examples, practical exercises, and links**.
This will help data scientists learn the theory behind RNNs and gain practical experience in implementing them for various tasks.

Highlights

- **Course Duration - 2 days (16 Hours)**
- **Number of Modules - 10**

Module 1: Introduction to Recurrent Neural Networks

- What are RNNs and their significance in sequential data analysis?
- Basic architecture and working principles of RNNs.
- Common use cases for RNNs in data science.

Module 2: Understanding Vanishing and Exploding Gradients

- Discuss the challenges of training RNNs, such as vanishing and exploding gradients.
- Explore how these issues affect the network's ability to learn long sequences.
- Introduce techniques like gradient clipping to mitigate these problems.

Module 3: Types of RNNs

- Introduction to different types of RNN architectures: Vanilla RNNs, LSTMs, and GRUs.
- Explain the advantages and disadvantages of each type.
- When to use each type based on specific use cases.

Module 4: Long Short-Term Memory (LSTM) Networks

- In-depth exploration of LSTM architecture.
- Explain how LSTMs address the vanishing gradient problem.
- Real-world applications where LSTMs shine.

Module 5: Gated Recurrent Unit (GRU) Networks

- Detailed discussion on GRU architecture and its simplicity compared to LSTMs.
- Highlight use cases and scenarios where GRUs are a good choice.

Module 6: Building Your First RNN with TensorFlow/Keras

- Hands-on tutorial on creating a simple RNN model using TensorFlow or Keras.
- Data preprocessing for sequential data and model evaluation.

Module 7: Text Generation with RNNs

- Step-by-step guide on training an RNN to generate text.
- Discuss the concept of character-level and word-level text generation.
- Tools, libraries, and datasets for text generation projects.

Module 8: Sequence-to-Sequence Models with RNNs

- Explore sequence-to-sequence tasks and their applications, like machine translation and chatbots.
- How to build Seq2Seq models using RNNs and attention mechanisms.

Module 9: Time Series Forecasting with RNNs

- Explain how RNNs can be used for time series forecasting.
- Discuss advanced topics like multivariate time series and windowed approaches.
- Hands-on example of building a time series forecasting model.

Module 10: Advanced RNN Techniques

- Cover advanced RNN techniques such as bidirectional RNNs, stacked RNNs, and attention mechanisms.
- Real-world case studies showcasing the impact of these advanced techniques.