

Artificial Intelligent training



Overview:

Artificial intelligence allows machines to replicate the capabilities of the human mind.

From the development of self-driving cars to the proliferation of smart assistants like Siri and Alexa, AI is a growing part of everyday life. As a result, many tech companies across various industries are investing in artificially intelligent technologies.

Introduction to Artificial Intelligence

- What is AI
- What is the AI future?

Introduction to Artificial Neural Network

- Artificial Neural Network
- Plan of Attack
- The Neuron
- The Activation Function
- How do Neural Networks work?
- How do Neural Networks learn?
- Gradient Descent

Convolutional Neural Network

- Plan of Attack
- What are Convolutional Neural Networks?
- Step 1 - The Convolution Operation
- Step 1 Bis - The ReLU Layer
- Step 2 - Pooling
- Step 3 - Flattening
- Step 4 - Full Connection
- Softmax & Cross-Entropy

AutoEncoder

- Plan of Attack
- What are AutoEncoders?
- Overcomplete Hidden Layers
- Sparse AutoEncoders
- Denoising AutoEncoders
- Contractive AutoEncoders
- Stacked AutoEncoders
- Deep AutoEncoders

Variational AutoEncoder

- What are Variational AutoEncoder
- Introduction to the VAE
- Variational AutoEncoders
- Reparameterization Trick

Implementing the CNN-VAE

- How to implement the CNN-VAE
- Initializing all the parameters and variables of the CNN-VAE class
- Building the Encoder part of the VAE
- Building the "V" part of the VAE
- Building the Decoder part of the VAE
- Implementing the Training operations
- The Keras Implementation

Recurrent Neural Network

- Introduction to Recurrent Neural Network
- Plan of Attack
- What are Recurrent Neural Networks?
- The Vanishing Gradient Problem
- LSTMs
- LSTM Practical Intuition
- LSTM Variations

Mixture Density Network

- Introduction to Mixture Density Network
- Introduction to the MDN-RNN
- Mixture Density Networks
- VAE + MDN-RNN Visualization

Implementing the MDN-RNN

- Introduction to Implementing the MDN-RNN
- Initializing all the parameters and variables of the MDN-RNN class
- Building the RNN - Gathering the parameters
- Building the RNN - Creating an LSTM cell with Dropout
- Building the RNN - Setting up the Input, Target, and Output of the RNN
- Building the RNN - Getting the Deterministic Output of the RNN
- Building the MDN - Getting the Input, Hidden Layer and Output of the MDN
- Building the MDN - Getting the MDN parameters
- Implementing the Training operations (Part 1)
- Implementing the Training operations (Part 2)
- The Keras Implementation

Reinforcement Learning

- Introduction to Reinforcement Learning
- What is Reinforcement Learning?
- A Pseudo Implementation of Reinforcement Learning for the Full World Model

Deep NeuroEvolution

- Introduction to Deep NeuroEvolution
- Evolution Strategies
- Genetic Algorithms

- Covariance-Matrix Adaptation Evolution Strategy (CMA-ES)
- Parameter-Exploring Policy Gradients (PEPG)
- OpenAI Evolution Strategy

The Final Run

- The Whole Implementation
- Download the whole AI Masterclass folder here
- Installing the required packages
- The Final Race: Human Intelligence vs. Artificial Intelligence

