Human Made

Unlocking AI

A guide to key AI concepts, algorithms, techniques & trends

Contents

AI Foundations

Definition of AI

Artificial Intelligence (AI) refers to the simulation of human intelligence in machines, enabling them to perform tasks that usually require human thinking, such as visual perception, speech recognition, decision-making, and translation.

At its core, AI involves the creation of algorithms that allow computers to perform tasks in a way that mimics human cognition.

Types of AI

Narrow AI: This is AI designed and trained for a particular task. Virtual personal assistants, like Siri or Alexa, are examples of Narrow AI.

General AI: Theorised as machines that can perform any intellectual task that a human being can, AGI is highly versatile and adaptive. Unlike Narrow AI, which is built for specific tasks, AGI can theoretically learn and apply knowledge in different domains.

Superintelligent AI: A future form of AI where its cognitive skills surpass human abilities across virtually all professions, from simple tasks to complex scientific reasoning. The concept of Superintelligence spurs both excitement and existential concerns.

AI Bytes

Gaming Achievement: In 2019, OpenAI's AI called "Dactyl" learned to solve a Rubik's Cube with a single robotic hand, showcasing advances in both robotics and AI.

History of AI

1943: McCulloch & Pitts Neural Model
First conceptual model of an artificial neuron.

1950: Turing Test
Alan Turing proposes a measure of a machine's ability to exhibit intelligent behaviour indistinguishable from that of a human.

1956: Dartmouth Conference
The term "Artificial Intelligence" is coined. The conference kickstarts interest and research into AI.

1965: ELIZA
An early natural language processing computer program, ELIZA mimics a psychotherapist and demonstrates the illusion of understanding.

1970: Expert Systems Boom
Rise of computers that mimic the decision-making abilities of a human expert. These systems use logical rules and became popular in various fields.

1980: Backpropagation
The backpropagation algorithm is introduced, which becomes fundamental for training neural networks.

1997: Deep Blue Defeats Kasparov
IBM's chess-playing computer, Deep Blue, defeats world chess champion Garry Kasparov, showcasing the potential of AI in specific tasks.

2012: AlexNet & Deep Learning
AlexNet wins the ImageNet Large Scale Visual Recognition Challenge. The deep learning model achieves a significant reduction in error rates for image recognition, leading to renewed interest and rapid advancements in deep learning.

2014: AlphaGo
Google DeepMind's AI, AlphaGo, defeats a world champion Go player in 2016. This is significant as Go is a complex game with more possible moves than atoms in the universe, highlighting the advancements in AI's problem-solving capabilities.

2018: BERT & Transformers
Introduction of the BERT model for natural language processing, utilising transformer architecture, reshaping the landscape of NLP tasks with its state-of-the-art performance.

AI Paradigms

1 Symbolic AI (or Good Old-Fashioned AI, GOFAI)

Overview: This paradigm is rooted in the use of symbols to represent knowledge and using rules to manipulate these symbols. It's largely based on logic.

Key Features:
- Knowledge Representation: Using facts, rules, and ontologies.
- Rule-Based Inference: Applying rules to derive new information.
- Logic-Based Reasoning: Using formal logic to solve problems.

Examples:
- Expert systems, SHRDLU (a natural language understanding computer program), Prolog (a programming language based on formal logic).

2 Connectionism (Neural Networks)

Overview: Connectionism models intelligence using artificial neural networks, which are inspired by the structure and functioning of biological neural networks in the brain.

Key Features:
- Nodes and Weights: Information processing units (neurons) connected by weighted links.
- Learning: Adjusting weights based on data to minimise prediction errors.
- Layers: Typically composed of input, hidden, and output layers.

Examples:
- Multilayer perceptrons, convolutional neural networks (CNN), recurrent neural networks (RNN).

3 Evolutionary Algorithms

Overview: These are optimisation algorithms inspired by the process of natural selection.

Key Features:
- Selection: Selecting the most optimal solutions from a pool.
- Crossover: Merging attributes of parent solutions to generate new ones.
- Mutation: Implementing minor alterations to solutions for diversity.
- Fitness Function: Assessing the effectiveness of a solution in addressing the given problem.

Examples:
- Genetic algorithms, genetic programming, evolutionary strategies.

4 Bayesian Networks

Overview: A type of probabilistic graphical model that represents a set of variables and their conditional dependencies via a directed acyclic graph (DAG).

Key Features:
- Nodes: Represent random variables.
- Edges: Indicate conditional dependencies.
- Probabilities: Assign probabilities to various events or states.
- Inference: Predict unknown variables given observed data.

Examples:
- Medical diagnosis systems, spam filters, recommendation systems.

5 Reinforcement Learning

Overview: An agent learns by interacting with an environment, receiving feedback in the form of rewards or penalties.

Key Features:
- Agent: Entity that takes actions.
- Environment: The world in which the agent operates.
- Actions: Decisions made by the agent.
- Rewards: Feedback received after taking an action.
- Policy: Strategy that the agent employs to decide actions.

Examples:
- Game playing (like AlphaGo), robotics, self-driving cars.

AI Bytes

AI's Winter Seasons: AI has gone through several "winters", or periods of reduced funding and interest, often because of overly optimistic promises and under-delivery. Symbolic AI, for instance, faced challenges that led to one of these winters in the late 20th century.

Neural Networks Basics

Neuron Model

Overview

A neuron, or a node, is the basic unit of a neural network. It receives one or more inputs and sums them to produce an output. The output is then transformed using an activation function.

Components

- **Input values** or **one-dimensional array** of data.
- **Weights:** Values that modulate the input.
- **Bias:** An additional parameter that allows the output to be shifted.
- **Net Sum:** Combined weighted input values and biases.
- **Activation Function:** A function that transforms the net sum into the neuron's output.

AI Bytes

The XOR Problem: In the early days of neural networks, the XOR problem was a big challenge. It's a simple problem where the network needs to distinguish between inputs [0,1], [1,0] vs. [0,0], [1,1]. A single-layer perceptron can't solve this, but a multi-layer one can, highlighting the importance of depth!

Loss Functions

Purpose

Measures the difference between the actual and predicted outputs. The goal during training is to minimise this difference.

Common Types

- **Mean Squared Error (MSE):** Average of the squared differences between the predicted and actual values. Used mainly for regression tasks.
- **Cross-Entropy:** Measures the difference between two probability distributions, commonly used for classification tasks.

Activation Functions

Purpose

They introduce non-linearity to the model, enabling neural networks to learn from error and make corrections, essential for learning complex patterns.

Common Types

- **Sigmoid:** Maps any input value to a value between 0 and 1.
- **ReLU (Rectified Linear Unit):** Allows positive values to pass through unchanged, but replaces negative values with zero.
- **Tanh (Hyperbolic Tangent):** Outputs values between -1 and 1.

Layers: Input, Hidden, Output

- **Input Layer:** This is the layer that directly receives the data for processing.
- **Hidden Layers:** Layers between input and output. They help in extracting features and learning patterns. There can be multiple hidden layers in deep networks.
- **Output Layer:** This layer produces the result, transforming intermediate representations from hidden layers into the desired format.

Forward & Backpropagation

Forward Propagation

The process of passing the input data through the network (from input to output layers) to generate a prediction.

Backpropagation

An optimisation algorithm used for minimising the error in the predictions. It works by calculating the gradient of the error function with respect to each weight by employing the chain rule and then adjusting the weights in the opposite direction of the gradient.

Deep Learning

CNNs

Neural networks designed for processing structured grid data like images.

RNNs

Neural networks designed for sequential data, like time series or natural language.

Transformers

Architectures that use attention mechanisms to draw global dependencies between input and output.

Transfer Learning

Leveraging knowledge from one task to improve performance on a different, yet related, task.

Regularisation

Techniques like **Dropout** and **Batch Normalisation** help in reducing overfitting and improving network convergence.

AI Bytes

Dreaming of Dogs: When deep learning researchers were trying to understand what individual neurons in their models were doing, they found that one neuron in Google's model was particularly good at recognizing... dog faces! This discovery led to the creation of Google's DeepDream, which often produces images filled with dog-like features.

Reinforcement Learning

Agent, Environment, Actions, Rewards

The agent takes actions in the environment, and these actions produce rewards. The agent's goal is to maximize its cumulative reward over time.

Q-learning & Deep Q Networks (DQN)

A method where the agent learns an action-value function to choose optimal actions. DQNs combine Q-learning with deep neural networks.

Exploration vs. Exploitation

The dilemma of whether to try out new actions (exploration) or stick with known beneficial actions (exploitation).

Policy Gradient Methods

These methods optimize the parameters of a policy by following the gradients toward higher rewards.

Common AI Algorithms

Supervised Learning

Algorithms learn from labeled training data, and this learning helps make predictions or decisions without human intervention. Examples: Linear Regression, Decision Trees, SVM.

Unsupervised Learning

Algorithms learn from data without explicit labels, typically identifying patterns or groupings. Examples: K-means clustering, PCA (dimensionality reduction), DBSCAN clustering.

Ensemble Methods

Combine multiple models to produce one optimal predictive model. Examples: Random Forests (ensemble of Decision Trees), Gradient Boosting Machines.

Optimisation

Techniques like Gradient Descent aim to minimise (or maximise) an objective function, helping models tune their parameters.

Essential Tools & Libraries

Python Libraries

- **TensorFlow** (by Google)
- **PyTorch** (by Facebook)
- **Keras** (runs on top of TensorFlow)
- **Scikit-learn** (for classical ML algorithms)

Platforms

- **Google AI Platform**
- **AWS SageMaker**
- **Azure Machine Learning**

Visualisation Tools

- **TensorBoard** (for TensorFlow)
- **Matplotlib**
- **Seaborn**

Tips and Tricks

Overfitting vs. Underfitting

Overfitting occurs when a model is too complex and starts to capture noise in the data, while underfitting means the model is too simple to capture underlying patterns.

Hyperparameter Tuning

The process of finding the best set of parameters (e.g., learning rate, batch size) for an AI/ML model.

Data Augmentation

Techniques to artificially increase the size of the training dataset by adding slightly modified copies of existing data.

Model Evaluation Metrics

Various metrics help gauge a model's performance, like Accuracy (classification), F1 Score (classification), ROC-AUC (classification), and Mean Absolute Error (MAE) for regression.

Trends & Emerging Topics

GANs

Two neural networks – a generator and a discriminator – are pitted against each other to generate new, synthetic instances of data.

Quantum Computing & AI

Leveraging the principles of quantum mechanics to build more powerful computation structures for AI.

Federated Learning

A method that trains AI models across multiple devices or servers while keeping the data localized, thus preserving privacy.

Neuromorphic Computing

Designing chips and hardware inspired by the human brain's architecture and functioning.

Resources

Online Courses & Websites

Coursera

Udacity

eDX

Udemy

hackr.io

YouTube Channels

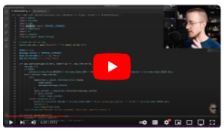

Sentdex

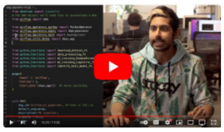

Siraj Raval

Data School

DeepLearningAI

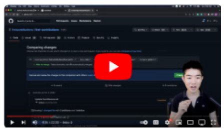

CS Dojo

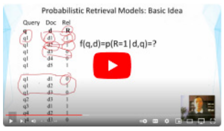

Artificial Intelligence - All in One

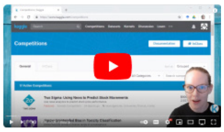

Kaggle

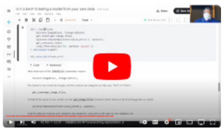

Jeremy Howard

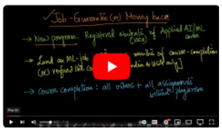

Applied AI Course

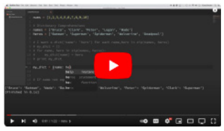

Corey Schafer

Books

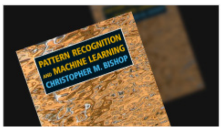

"Pattern Recognition and Machine Learning" by Bishop

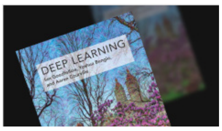

Deep Learning by Goodfellow, Bengio, and Courville

Superintelligence: Paths, Dangers, Strategies by Nick Bostrom

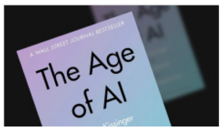

The Age of AI: And Our Human Future by Henry A Kissinger, Eric Schmidt, Daniel Huttenlocher

Atlas of AI by Kate Crawford

Communities and Forums

Stack Overflow

Reddit's **Machine Learning**

Hugging Face

Kaggle

Dev.to