Content : Deep Learning and Reinforcement Learning Training (2 days)

Audience

Knowledge on AI Technology

Content of the Training:

Day 1: Deep Learning

- Introduction to Deep Learning
    - What is deep learning?
    - Applications of deep learning
    - Basic terminology (neural networks, layers, activation functions, etc.)
    - Neural Networks and Activation Functions
- Structure of a neural network
    - Types of layers (input, hidden, output)
    - Common activation functions (sigmoid, ReLU, etc.)
    - Hands-on exercise: Implementing a simple neural network from scratch
- Training Neural Networks
    - Loss functions and optimization algorithms (gradient descent, backpropagation)
    - Regularization techniques (dropout, weight decay)
    - Tips for effective training (learning rate, batch size, etc.)
    - Hands-on exercise: Training a neural network using a deep learning framework (e.g., TensorFlow, PyTorch)
- Convolutional Neural Networks (CNNs)
    - Motivation and architecture of CNNs
    - Convolutional layers, pooling layers
    - Common CNN architectures (LeNet, AlexNet, etc.)
    - Hands-on exercise: Implementing a CNN for image classification
- Recurrent Neural Networks (RNNs)
    - Introduction to sequential data processing
    - Structure and functioning of RNNs
    - Applications of RNNs (sequence generation, language modeling, etc.)
    - Hands-on exercise: Building a simple RNN for text generation
- Transfer Learning and Pretrained Models
    - Leveraging pretrained models for transfer learning
    - Fine-tuning and feature extraction
    - Hands-on exercise: Using a pretrained model for a different task
- Generative Adversarial Networks (GANs)
    - Introduction to generative models
    - GAN architecture and training process
    - Generating realistic images with GANs
    - Hands-on exercise: Creating a basic GAN for image generation
- Applications of Deep Learning
    - Overview of deep learning applications in various domains (computer vision, natural language processing, etc.)
    - Recent advancements and future trends in deep learning

Day 2: Reinforcement Learning

- Introduction to Reinforcement Learning (RL)
    - What is RL and why is it important?
    - Key concepts: agent, environment, states, actions, rewards
    - RL terminology (Markov Decision Processes, policy, value functions, etc.)
- Markov Decision Processes (MDPs)
    - Formal definition of MDPs
    - Value iteration and policy iteration algorithms
    - Hands-on exercise: Solving a simple MDP problem
- Q-Learning and Temporal Difference (TD) Learning
    - Introduction to Q-Learning
    - Exploration vs. exploitation trade-off
    - TD learning and the Q-Learning algorithm
    - Hands-on exercise: Implementing Q-Learning for a simple RL problem
- Deep Q-Networks (DQNs)
    - Challenges of using Q-Learning with large state spaces
    - Introduction to DQNs and the role of deep neural networks
    - Experience replay and target networks
    - Hands-on exercise: Building a DQN for a RL task using a deep learning framework
- Policy Gradient Methods
    - Introduction to policy gradients
    - Policy parameterization and the policy gradient theorem
    - REINFORCE algorithm
    - Hands-on exercise: Implementing a basic policy gradient