Kubeflow

Course Outline

- Introduction to Kubernetes
- Overview of Kubeflow Features and Architecture
- Kubeflow on AWS vs on-premise vs on other public cloud providers
- Setting up a Cluster using AWS EKS
- Setting up an On-Premise Cluster using Microk8s
- Deploying Kubernetes using a GitOps Approach
- Data Storage Approaches
- Creating a Kubeflow Pipeline
- Triggering a Pipeline
- Defining Output Artifacts
- Storing Metadata for Datasets and Models
- Hyperparameter Tuning with TensorFlow
- Visualizing and Analyzing the Results
- Multi-GPU Training
- Creating an Inference Server for Deploying ML Models
- Working with JupyterHub
- Networking and Load Balancing
- Auto Scaling a Kubernetes Cluster
- Troubleshooting
- Summary and Conclusion