

Advanced Deep Learning and Computer Vision

Course Content

Lesson 1 - Introduction to Deep Computer Vision

- Overview of CV applications
- Pipeline for vision projects
- Image data handling

Lesson 2 - Advanced CNN Architectures

- VGG, ResNet, Inception
- Transfer learning revisited
- Model selection

Lesson 3 - Object Detection and Localization

- Bounding boxes
- YOLO and SSD frameworks
- Metrics for detection

Lesson 4 - Image Segmentation Techniques

- Semantic vs instance segmentation
- U-Net architecture
- Applications in healthcare

Lesson 5 - Generative Adversarial Networks (GANs)

- GAN architecture
- Training stability
- Image generation

Lesson 6 - Transfer Learning for Vision

- Domain adaptation
- Layer freezing
- Fine-tuning best practices

Lesson 7 - Attention Mechanisms

- Self-attention
- Transformers for vision
- Visual transformers

Lesson 8 - Video Analytics with Deep Learning

- Video classification
- Action detection
- Temporal modeling

Lesson 9 - Ethics and Fairness in Vision AI

- Bias in datasets
- Transparency and explainability
- Regulatory frameworks

Lesson 10 - Vision AI Applications in Industry

- Manufacturing
- Retail
- Security and surveillance

Capstone Project

- Build an end-to-end computer vision solution for a real-world dataset.