

APPLIED COMPUTER VISION ESSENTIALS

Course Code: 840043

Learn to build, deploy, and evaluate modern computer vision systems—from classical techniques to cutting-edge deep learning

Applied Computer Vision Essentials is a hands-on course designed for professionals eager to deepen their understanding of modern computer vision techniques. Whether you're transitioning from classical image processing or already working with deep learning models, this course offers a structured path to mastering the tools and concepts that power today's most advanced visual systems. From edge detection and feature extraction to segmentation and multimodal pipelines, learners will explore the full spectrum of computer vision applications through practical labs and real-world scenarios.

Participants will gain experience with cutting-edge frameworks like YOLOv9, SAM 2, and DINOv2, while building and deploying models in a GPU-enabled Ubuntu environment. The course emphasizes not just technical proficiency but also ethical considerations, including bias auditing and production monitoring. With a curriculum that blends theory, demos, and capstone projects, learners will leave equipped to tackle challenges in domains ranging from industrial automation to health tech and retail analytics.

Ideal for software engineers, data scientists, and MLOps professionals, this course bridges the gap between foundational knowledge and applied expertise. Whether you're optimizing models for edge deployment or integrating vision with language models for safety reporting, Applied Computer Vision Essentials provides the skills and confidence to build robust, scalable solutions.

What You'll Learn

- Apply classical computer vision techniques for edge detection, feature extraction, and lane detection
- Analyze color spaces, histogram equalization, and contrast enhancement methods for image quality improvement
- Create data augmentation pipelines and fine-tune CNN architectures like EfficientNet for classification
- Evaluate object detection performance using mAP and IoU metrics with TIDE error analysis
- Implement YOLO training workflows for safety compliance with hyperparameter optimization

- Compare segmentation approaches from traditional methods to modern promptable SAM 2
- Construct Vision Transformer solutions using DINOv2 and self-supervised learning principles
- Synthesize multimodal pipelines integrating detection, CLIP embeddings, and language models for alt-text generation
- Optimize models for production through ONNX conversion, INT8 quantization, and edge deployment
- Assess computer vision systems for bias and fairness while implementing production monitoring with Prometheus

Who Needs to Attend

Sample learning personas:

- Rajesh Singh – Senior software engineer, industrial-automation firm, Bengaluru, India. Uses classical OpenCV; needs a roadmap for defect and lane detection with deep learning.
- Maria Alvarez – Data scientist, retail supply-chain analytics, Guadalajara, Mexico. Comfortable with PyTorch classifiers; wants hands-on object detection and edge deployment for PPE compliance.
- Esther Ndiaye – Machine-learning engineer, health-tech start-up, Dakar, Senegal. NLP background; seeks robust instrument segmentation and guidance on regulatory alignment.
- Lucas Chen – DevOps engineer moving into MLOps, Toronto, Canada. Strong in Docker and CI/CD; aims to learn model quantisation, monitoring, and bias auditing for a vision API.

Prerequisites

- Working knowledge of Python 3.9+: functions, classes, virtual-environment management (venv or conda), package install with pip.
- Familiarity with NumPy arrays and tensor concepts; ability to write a simple forward pass in PyTorch or TensorFlow.
- Experience running a supervised-learning loop: dataset split, loss calculation, back-prop, checkpoint save.
- Basic shell skills on Linux (navigate directories, edit config files, run git clone).
- Git fundamentals: clone, branch, commit, push, pull-request workflow.
- JupyterLab usage: open notebooks, run cells, inspect GPU memory.
- Awareness of GPU vs CPU execution; can read nvidia-smi output or fallback to CPU when GPUs are unavailable.
- Introductory linear-algebra and probability: matrix multiply, softmax, cross-entropy.
- Ability to read JSON/YAML config files and tweak hyper-parameters.
- Laptop or desktop with stable broadband (⚡ 10 Mbps down / 2 Mbps up) and a modern browser that reaches Skillable lab URLs over HTTPS.
- Company VPN, proxy, or security policy allows outbound WebSocket traffic for JupyterLab (ports 8888/8443) and VS Code Server if used.

- Optional but helpful: basic Docker commands (docker build, docker run) and REST API testing with curl or Postman.

APPLIED COMPUTER VISION ESSENTIALS

Course Code: 840043

VIRTUAL CLASSROOM LIVE

\$3,195 USD

4 Day

Virtual Classroom Live Outline

Foundations & Classical Computer Vision

- Pixels, color spaces, convolution filters
- Lane-finding with Canny + Hough
- Histogram equalisation & CLAHE
- Low-light rescue with CLAHE
- Feature extraction: classical descriptors
- Image matching: ORB vs SIFT
- CVAT annotation + COCO export
- Wrap-up: bridging classical to modern CV

Deep Learning for Computer Vision

- Classical to deep transition
- CNN architectures & evolution
- Data-augmentation strategies
- AutoAugment & RandAugment demo
- Fine-tune EfficientNet-V2-S + Grad-CAM
- Intro to object detection & YOLO family
- YOLOv11-nano training start
- Detection metrics & interpretation; TIDE taxonomy
- Model robustness discussion

Advanced Vision: Segmentation & Transformers

- From detection to segmentation
- Segmentation approaches
- SAM 2: promptable segmentation
- SAM 2 segmentation vs YOLO masks
- Vision Transformers revolution
- Video processing fundamentals
- Attention rollout visualisation
- Self-supervised learning
- Fine-tune DINOv2-tiny
- Modern CV landscape

- Capstone prep

Modern Applications & Integration

- Recap: CV evolution journey
- Vision-language models
- Image & video generation
- Detector ? CLIP ? LLM safety report
- Model deployment essentials
- ONNX conversion & optimization
- Production monitoring demo
- Adversarial robustness
- Ethics in Computer Vision
- Wrap-up; Q&A
- Capstone demos

Aug 17 - 20, 2026 | 8:30 AM - 4:30 PM EDT

Sep 21 - 24, 2026 | 8:30 AM - 4:30 PM EDT

Oct 26 - 29, 2026 | 8:30 AM - 4:30 PM EDT

Nov 30 - Dec 3, 2026 | 8:30 AM - 4:30 PM EST

Visit us at www.globalknowledge.com or call us at 1-866-716-6688.

Date created: 6/17/2026 5:24:16 AM

Copyright © 2026 Global Knowledge Training LLC. All Rights Reserved.