SAGIVTECH COURSES

Where AI & Computer Vision Meet

ABOUT SAGIVTECH

SagivTech was established in 2009 as a courses and projects company.

Our main focus is in Computer Vision and Deep Learning.

We provide R&D services in algorithms development and software implementation.

We give courses for AI and Computer Vision since 2009!

For more information: masha@sagivtech.com

TABLE OF CONTENTS

- CUDA Programming for Beginners
- Advanced CUDA Programming
- Classical Computer Vision
- Advanced Topics In Computer Vision
- Deep Learning for Computer Vision
- Classical Computer Vision For Everyone
- AI In Computer Vision For Everyone
- Deep Learning with Python

CUDA PROGRAMMING COURSE FOR BEGINNERS

Course Overview: This 6-day course is designed to introduce participants to CUDA programming for parallel computing. Through a mix of theory and hands-on sessions, participants will learn to develop efficient GPU-accelerated applications using NVIDIA's CUDA platform.

- Introduction to CUDA and Parallel Programming Histograms based operations

- CUDA Basics and Memory Management

- Parallel Programming Concepts

- Advanced CUDA Features

- CUDA Libraries and Integration

- Project Development and Optimization

ADVANCED CUDA PROGRAMMING

Course Overview: This course covers advanced topics in parallel computing using CUDA (Compute Unified Device Architecture) developed by NVIDIA. It builds upon the fundamentals of CUDA programming and delves into more complex techniques for optimizing performance, managing memory, and leveraging advanced features of CUDA. The course aims to equip students with the skills necessary to develop highly efficient and scalable parallel applications for various domains, including scientific computing, machine learning, and high-performance computing.

- CUDA Programming Model Review
- Performance Optimization Techniques
- Memory Management
- Advanced CUDA Features

- CUDA Libraries and Frameworks
- CUDA and Heterogeneous Computing
- CUDA Debugging and Profiling
- Advanced Topics and Applications

CLASSICAL COMPUTER VISION

- Sampling & Quantization

- Histograms based operations

- Convolutions

- Fourier

- Image Denoising

- Image Segmentation

- Image Restoration

- Classification

- Morphology

Target Audience: people who are new to image processing & computer vision, algorithms developers, engineers, SW developers

Duration: 24 academic hours

ADVANCED TOPICS IN COMPUTER VISION

- SIFT and friends

- Viola Jones

- Geodesic Active Contours

- HOG

- 3D Vision

- Image Formation

- Gabor & Wavelets

- Laplacian Gaussian Pyramid

- Super Resolution

- Blur

- HDR

Target Audience: people who are familiar with basics of image processing & computer vision, algorithms developers, engineers, SW developers

Duration: course is modular and can be tuned to the needs of the group

DEEP LEARNING FOR VISION

- The building blocks of Neural Networks (NN)

- NN for Classification

- NN for Segmentation

- NN for Object Detection

- Metrics for Deep Learning

- How to train NN

- Object Detection: the RNN Family, SSD and YOLO

- Segmentation:
FCN, Unet, MassRCNN, Deep Lab, Segment All

- Autoencoders & VAE

- RNN, LSTM, Transformers

- Diffusion Models

- GANS

Target Audience: people who are familiar with deep learning basics, algorithms developers, engineers, SW developers

Duration: 24 academic hours. Course is modular and can be tuned to the needs of the group

CLASSICAL COMPUTER VISION FOR EVERYONE

- What is classical image processing and computer vision ?
- Sampling & Quantization
- Histograms and pixel wise operations
- Convolutions for denoising and edge detection
- Non linear operations: median, bilateral filter, morphology
- Fourier, Gabor and Wavelets
- Segmentation Methods

- Denoising Methods
- Super resolution
- HDR
- 3D Imaging
- Deconvolution
- Snakes & Variational approaches
- Viola Jones
- SIFT and friends
- HOG

The course aims to review classical image processing and computer vision methods to people who are new to this field

Duration: 2*4 hours sessions

AI IN COMPUTER VISION FOR EVERYONE

- A short history of AI

- The basics of CNN and how to train them

- Architectures for classification, detection and segmentation

- Metrics to evaluate AI

- Explainability

- Gans, SSD, Unet, RCNN Family, YOLO, DeepLab

- Autoencoders

- Transformers

- Diffusion Models

- Latest & Greatest

The course aims to review Deep Learning and AI in Computer Vision to people who are new to this field.

Duration: 2*4 hours sessions

DEEP LEARNING WITH PYTHON

Module A - Deep Learning Foundations (Hands-On)

Hands-On Topics

- Backpropagation from scratch (NumPy)
- Gradient checking
- Vanishing/exploding gradients
- Weight initialization experiments

Labs

- Build a neural net without PyTorch
- Visualize loss landscapes
- Show gradient norms across layers

Module B — CNNs & Vision (Hands-On)

Labs

- CNN from scratch (PyTorch)
- Transfer learning with ResNet
- CAM / Grad-CAM visualization
- Overfitting & augmentation experiments

Datasets

- CIFAR-10
- Histopathology patches (very strong for your audience)

DEEP LEARNING WITH PYTHON

Module C - RNNs, LSTMs & Attention (Hands-On)

Labs

- RNN vs LSTM on sequence task
- Truncated BPTT experiments
- Attention mechanism from scratch
- Seq2Seq with and without attention

Module D - Transformers & Foundation Models (Hands-On)

Labs

- Self-attention from scratch
- Multi-head attention implementation
- Fine-tuning a small transformer
- Embeddings + similarity search