

# Computer Vision



## Overview:

This course provides an introduction to computer vision including fundamentals of image formation, camera imaging geometry, feature detection and matching, Multiview geometry including stereo, motion estimation and tracking, and classification.

We'll develop basic methods for applications that include finding known models in images, depth recovery from stereo, camera calibration, image stabilization, automated alignment (e.g. panoramas), tracking, and action recognition. We focus less on the machine learning aspect of CV as that is really classification theory best learned in an ML course.

## Topics to cover:

### Face detection

- Images and pixels
- Cascade classifier - intuition
- Loading and pre-processing the image
- Face detection with Haarcascade and OpenCV
- Haarcascades parameters
- Eye detection with haarcascades
- HOG (Histograms of Oriented Gradients) – intuition
- Face detection with HOG and Dlib
- Face detection with CNN and Dlib
- Anaconda and PyCharm
- Face detection on the webcam

### Face recognition

- LBPH algorithm - intuition
- Loading the faces dataset
- Evaluating the LBPH classifier
- LBPH parameters
- Detecting facial points
- Detecting facial descriptors
- Calculating distances between faces
- Recognizing faces with Dlib
- Face recognition on the webcam

### Object tracking

- Object tracking vs. object detection
- KCF and CSRT algorithms
- Object tracking with KCF
- Object tracking with CSRT

### Neural networks for image classification

- Biological fundamentals
- Artificial neuron
- Perceptron
- Weight update with back propagation
- Introduction to multilayer neural networks
- Activation functions
- Output layer activation
- Error calculation (loss function)
- Basic algorithm
- Gradient descent and derivative
- Output layer delta
- Backpropagation and learning rate
- Weight update with backpropagation
- Bias, error and multiple outputs
- Stochastic gradient descent
- Deep learning
- Pixels and neural networks
- Importing the libraries
- Extracting pixels from images
- Normalizing the data
- Creating the train and test sets
- Building and training the neural network
- Evaluating the neural network
- Extracting features from images
- Feature extraction with OpenCV
- Creating the train and test sets
- Building and training the neural network
- Evaluating the neural network
- Saving, loading and classifying one single image

#### Convolutional neural networks for image classification

- Introduction to convolutional neural networks
- Convolutional operation
- Pooling
- Flattening
- Dense neural network
- Importing the libraries
- Loading the images
- Creating the train and test dataset
- Building and training the neural network
- Evaluating the neural network

#### Transfer learning and fine tuning

- Transfer learning – intuition
- Importing the libraries and dataset
- Creating the train and test dataset
- Pre-trained neural network
- Creating the custom dense layer

- Building and training the neural network
- Evaluating the neural network
- Fine tuning – intuition
- Fine tuning – implementation and evaluation
- Saving, loading and classifying one single image

#### Neural networks for classification of emotions

- Importing the libraries and images
- Creating the train and test dataset
- Building and training the neural network
- Saving and loading the model
- Evaluating the neural network

#### Autoencoders

- Autoencoders – intuition
- Importing the libraries and dataset
- Visualizing the images
- Building and training a linear autoencoder
- Encoding the images
- Decoding the images
- Encoding and decoding the test images
- Convolutional autoencoders

#### Object detection with YOLO

- YOLO – intuition
- Downloading and compiling Darknet
- Testing the detector
- Darknet and GPU
- Threshold and ext\_output parameters

#### Recognition of gestures and actions

- Gestures and actions recognition – intuition
- Importing the libraries and the image
- Loading the pre-trained neural network
- Predicting body points
- Detecting gestures in images
- Detecting gestures

#### Deep dream

- Deep dream – intuition
- Loading the Inception Net network
- Loading and pre-processing the image
- Getting the activations
- Calculating the loss
- Gradient ascent
- Generating images
- Plan of attack
- Style transfer – intuition

- Loading VGG19 network
- Building the neural network
- Visualizing the result
- Plan of attack
- GANs – intuition
- Loading the dataset

#### Image segmentation

- Image segmentation – intuition
- Downloading the repository
- Importing the libraries
- Loading the pre-trained neural network
- Detecting objects
- Removing the background
- Segmentation

