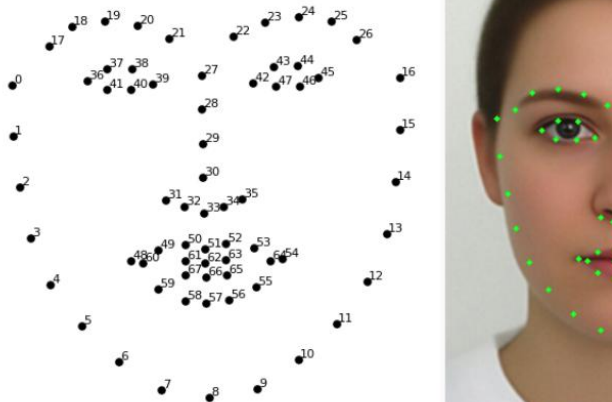


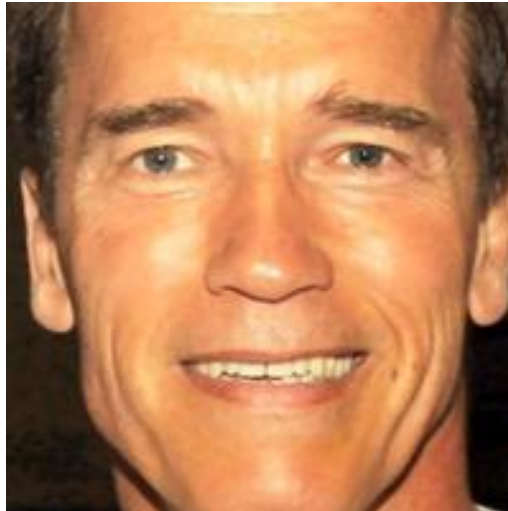
## Facial Landmark Detection (keypoint detection)

- Face Recognition
- Driver Analysis
- Face Swap
- Head Pose Estimation
- Facial Region Extraction
- Emotion Detection



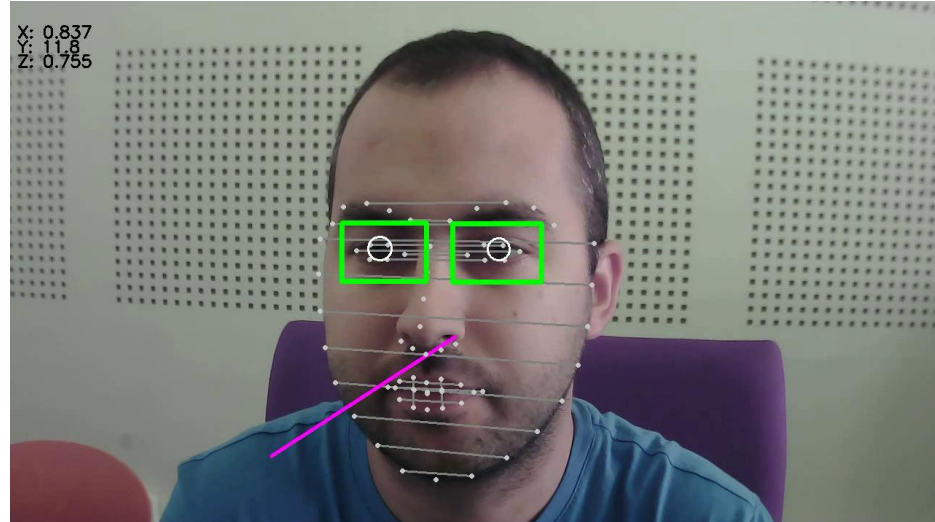
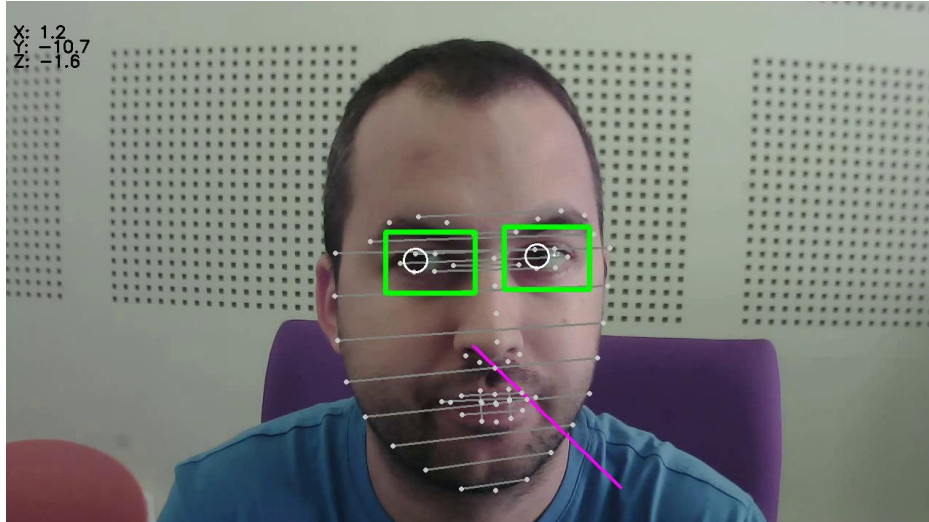
## Applications of Facial landmark detection (keypoint detection)

- Facial landmarks can be used to align facial images to improve face recognition



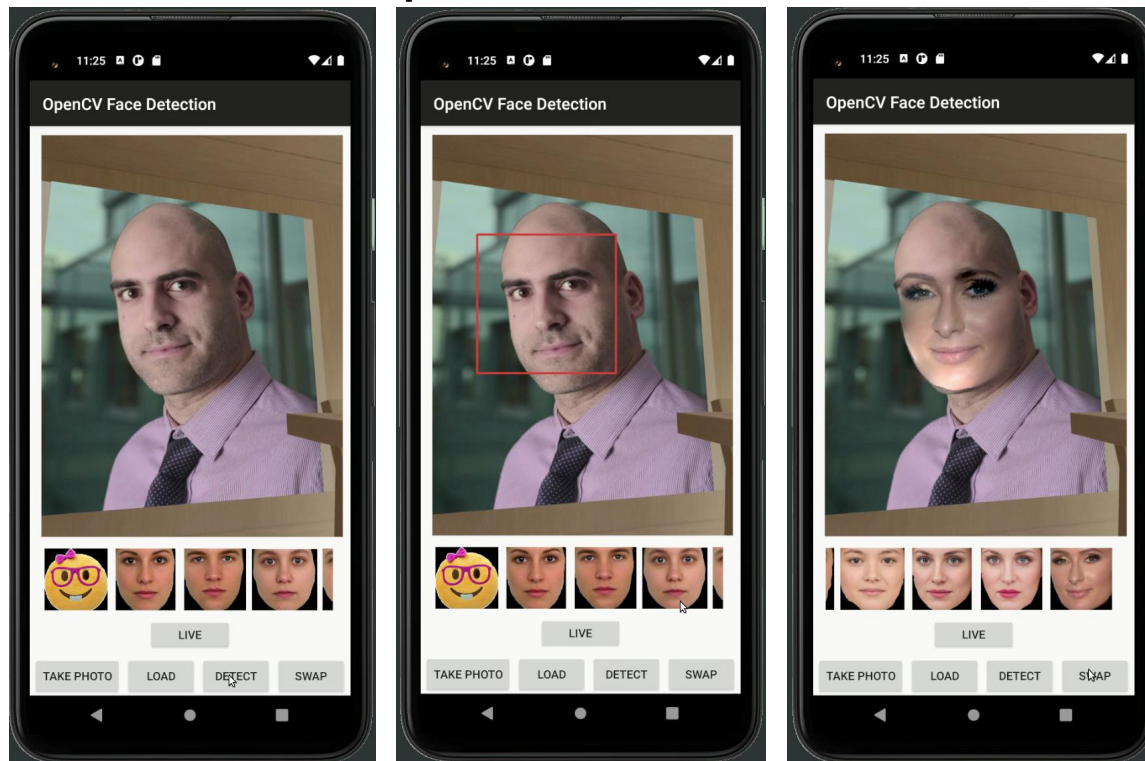
## Applications of Facial landmark detection (keypoint detection)

- We can estimate Head pose - where the person is looking



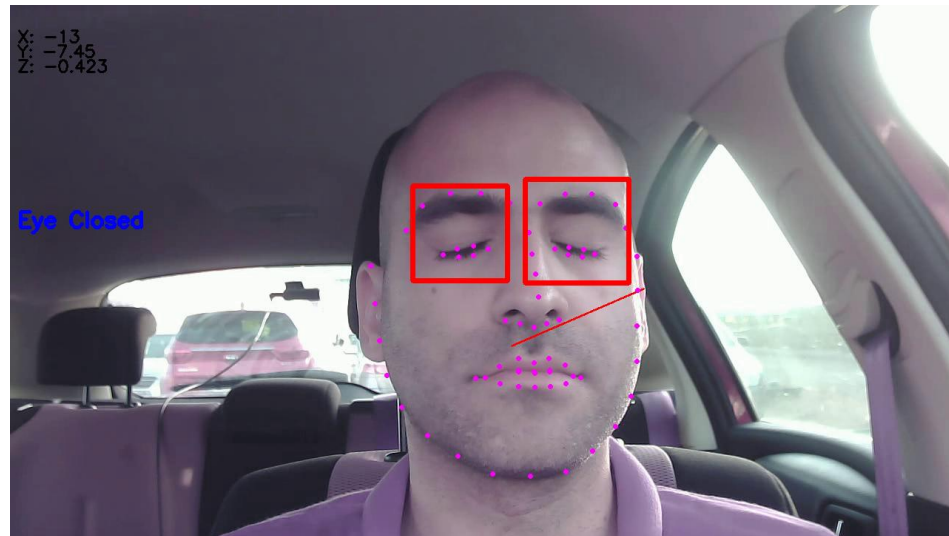
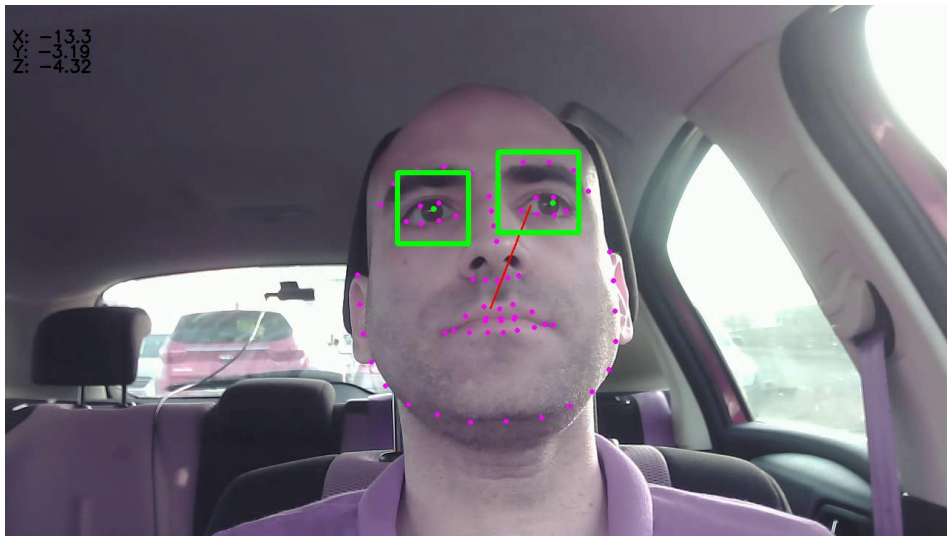
## Applications of Facial landmark detection (keypoint detection)

- Face Replacement/Face Swap



## Applications of Facial landmark detection (keypoint detection)

- Eye Blink Detection



## Challenges (can have a significant influence on the final detection)



(a) Pose

(b) Occlusion

(c) Expression

(d) Illumination

**Pose:** Faces can have several different poses (e.g. frontal, profile, from the bottom)

**Occlusion:** For example, glasses, hair or other items can cover the faces

**Expression:** Different facial expressions can create deformation of parts of the face

**Illumination:** Different Illumination Intensity (covering certain parts of the face, shadow)

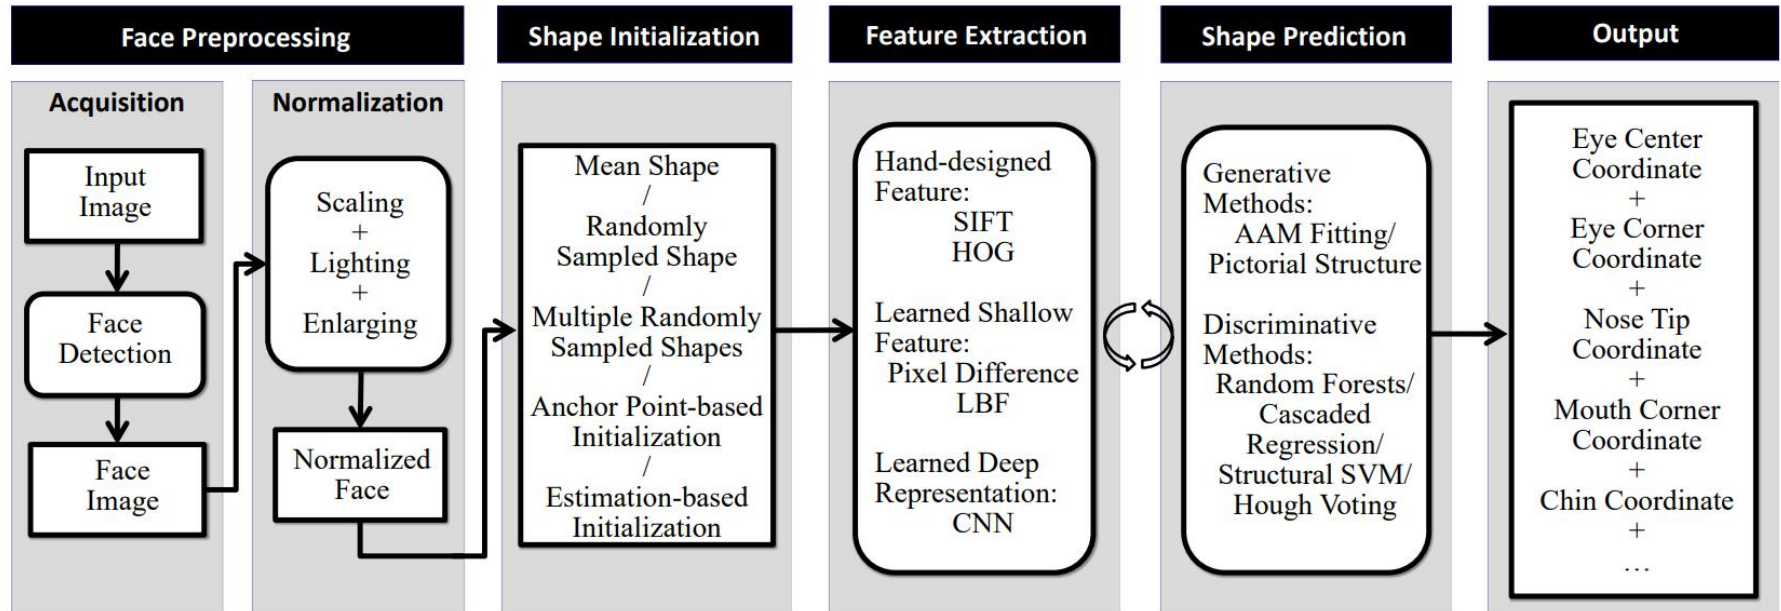


Figure 7: A global system architecture for face alignment.

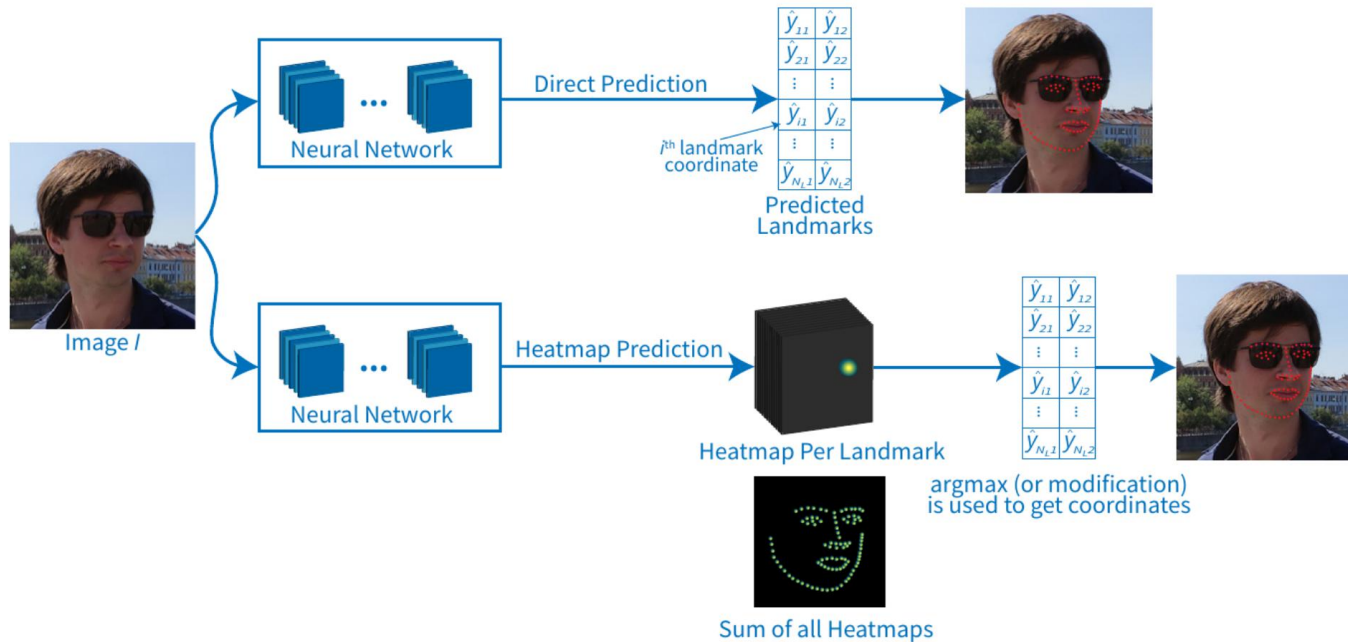


Figure 2. Direct landmark regression (upper row). The problem is solved in a form of regression, where actual landmark coordinates  $(x, y)$  are predicted directly by the algorithm. Heatmap-based (bottom row). The algorithm predicts probability distributions of landmark locations in a form of heatmaps. One heatmap per each of the landmarks is formed. Argmax (or its modification) is used to get each landmark coordinates.



Table 1: Categorization of the popular approaches for face alignment.

Approach	Representative works
<b>Generative methods</b>	
<i>Active appearance models (AAMs)</i>	
Regression-based fitting	Original AAM [16]; Boosted Appearance Model [17]; Nonlinear discriminative approach [18]; Accurate regression procedures for AMMs [19]
Gradient descent-based fitting	Project-out inverse compositional (POIC) algorithm [20]; Simultaneous inverse compositional (SIC) algorithm [21]; Fast AAM [22]; 2.5D AAM [23]; Active Orientation Models [24]
<i>Part-based generative deformable models</i>	
	Original Active Shape Model (ASM) [25]; Gauss-Newton deformable part model [26]; Project-out cascaded regression [27]; Active pictorial structures [28]
<b>Discriminative methods</b>	
<i>Constrained local models (CLMs)<sup>a</sup></i>	
PCA shape model	Regularized landmark mean-shift [29]; Regression voting-based shape model matching [30]; Robust response map fitting [31]; Constrained local neural field [32]
Exemplar shape model	Consensus of exemplar [11]; Exemplar-based graph matching [33]; Robust Discriminative Hough Voting [34]
Other shape models	Gaussian Process Latent Variable Model [35]; Component-based discriminative search [36]; Deep face shape model [37]
<i>Constrained local regression</i>	
	Boosted regression and graph model [38]; Local evidence aggregation for regression [39]; Guided unsupervised learning for model specific models [40]
<i>Deformable part models (DPMs)</i>	
	Tree structured part model [41]; Structured output SVM [42]; Optimized part model [43]; Regressive Tree Structured Model [44]
<i>Ensemble regression-voting</i>	
	Conditional regression forests [12]; Privileged information-based conditional regression forest [45]; Sieving regression forest votes [46]; Nonparametric context modeling [47]
<i>Cascaded regression</i>	
Two-level boosted regression	Explicit shape regression [48]; Robust cascaded pose regression [49]; Ensemble of regression trees [50]; Gaussian process regression trees [51];
Cascaded linear regression	Supervised descent method [52]; Multiple hypotheses-based regression [53]; Local binary feature [54]; Incremental face alignment [55]; Coarse-to-fine shape search [56]
<i>Deep neural networks<sup>b</sup></i>	
Deep CNNs	Deep convolutional network cascade [57]; Tasks-constrained deep convolutional network [58]; Deep Cascaded Regression [59]
Other deep networks	Coarse-to-fine Auto-encoder Networks (CFAN) [60]; Deep face shape model [37]

<sup>a</sup> Classic Constrained Local Models (CLMs) typically refer to the combination of local detector for each facial point and the parametric Point Distribution Model [61, 62, 29]. Here we extend the range of CLMs by including some methods based on other shape models (i.e., exemplar-based model [11]). In particular, we will show that the exemplar-based method [11] can also be interpreted under the conventional CLM framework.

<sup>b</sup> We note that some deep learning-based systems can also be placed in other categories. For instance, some systems are constructed in a cascade manner [60, 59, 63], and hence can be naturally categorized as cascaded regression. However, to highlight the increasing important role of deep learning techniques for face alignment, we organize them together for more systematic introduction and summarization.

## Generative :

These methods typically construct parametric models (or statistical models, **Active shape models**) and try to find the optimal parameters that gives best fit of the shape model to the test image.

## Discriminative:

These methods typically use independent local detector or regressor for each facial point. In many cases, this requires more training data. As time goes on, larger and larger datasets are being created, and methods based on deep learning have achieved very promising results and play a dominant role in this area in recent years.

Table 4: A list of sources of wild databases for face alignment.

Databases	Year	#Images	#Training	#Test	#Point	Links
LFW [96]	2007	13,233	1,100	300	10	<a href="http://www.dantone.me/datasets/facial-features-lfw/">http://www.dantone.me/datasets/facial-features-lfw/</a>
LFPW [11]	2011	1,432 <sup>a</sup>	-	-	35 <sup>b</sup>	<a href="http://homes.cs.washington.edu/~neeraj/databases/lfpw/">http://homes.cs.washington.edu/~neeraj/databases/lfpw/</a>
AFLW [112]	2011	25,993	-	-	21	<a href="http://lrs.icg.tugraz.at/research/aflw">http://lrs.icg.tugraz.at/research/aflw</a>
AFW [41]	2012	205	-	-	6	<a href="http://www.ics.uci.edu/~xzhu/face/">http://www.ics.uci.edu/~xzhu/face/</a>
HELEN [84]	2012	2,330	2,000	300	194	<a href="http://www.ifp.illinois.edu/~vuongle2/helen/">http://www.ifp.illinois.edu/~vuongle2/helen/</a>
300-W [113]	2013	3,837	3,148	689	68	<a href="http://ibug.doc.ic.ac.uk/resources/300-W/">http://ibug.doc.ic.ac.uk/resources/300-W/</a>
COFW [49]	2013	1,007	-	-	29	<a href="http://www.vision.caltech.edu/xpburgos/ICCV13/">http://www.vision.caltech.edu/xpburgos/ICCV13/</a>
MTFL [58]	2014	12,995	-	-	5	<a href="http://mmlab.ie.cuhk.edu.hk/projects/TCDCN.html">http://mmlab.ie.cuhk.edu.hk/projects/TCDCN.html</a>
MAFL [114]	2016	20,000	-	-	5	<a href="http://mmlab.ie.cuhk.edu.hk/projects/TCDCN.html">http://mmlab.ie.cuhk.edu.hk/projects/TCDCN.html</a>

<sup>a</sup> LFPW is shared by web URLs, but some URLs are no longer valid.

<sup>b</sup> Each face image in LFPW is annotated with 35 points, but only 29 points defined in [11] are used for the face alignment.

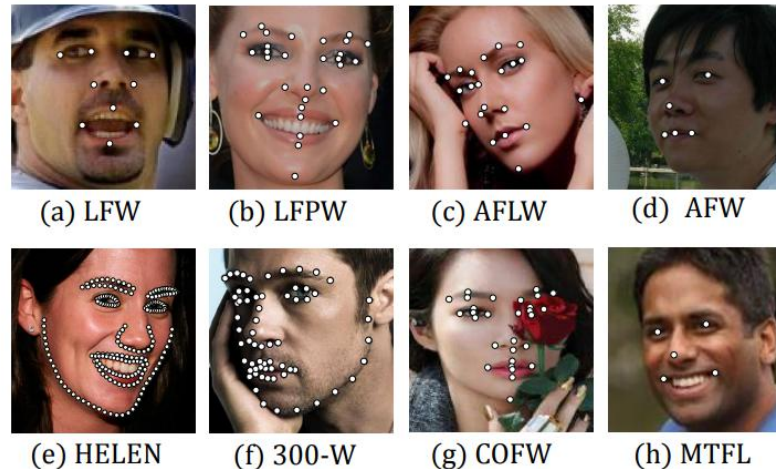


Figure 8: Illustration of the example face images from eight wide face databases with original annotation.

Table 1: Categorization of the popular approaches for face alignment.

Approach	Representative works
<b>Generative methods</b>	
<i>Active appearance models (AAMs)</i>	
Regression-based fitting	Original AAM [16]; Boosted Appearance Model [17]; Nonlinear discriminative approach [18]; Accurate regression procedures for AAMs [19]
Gradient descent-based fitting	Project-out inverse compositional (POIC) algorithm [20]; Simultaneous inverse compositional (SIC) algorithm [21]; Fast AAM [22]; 2.5D AAM [23]; Active Orientation Models [24]
<i>Part-based generative deformable models</i>	
	Original Active Shape Model (ASM) [25]; Gauss-Newton deformable part model [26]; Project-out cascaded regression [27]; Active pictorial structures [28]
<b>Discriminative methods</b>	
<i>Constrained local models (CLMs)<sup>a</sup></i>	
PCA shape model	Regularized landmark mean-shift [29]; Regression voting-based shape model matching [30]; Robust response map fitting [31]; Constrained local neural field [32]
Exemplar shape model	Consensus of exemplar [11]; Exemplar-based graph matching [33]; Robust Discriminative Hough Voting [34]
Other shape models	Gaussian Process Latent Variable Model [35]; Component-based discriminative search [36]; Deep face shape model [37]
<i>Constrained local regression</i>	Boosted regression and graph model [38]; Local evidence aggregation for regression [39]; Guided unsupervised learning for model specific models [40]
<i>Deformable part models (DPMs)</i>	Tree structured part model [41]; Structured output SVM [42]; Optimized part model [43]; Regressive Tree Structured Model [44]
<i>Ensemble regression-voting</i>	Conditional regression forests [12]; Privileged information-based conditional regression forest [45]; Sieving regression forest votes [46]; Nonparametric context modeling [47]
<i>Cascaded regression</i>	
Two-level boosted regression	Explicit shape regression [48]; Robust cascaded pose regression [49]; Ensemble of regression trees [50]; Gaussian process regression trees [51];
Cascaded linear regression	Supervised descent method [52]; Multiple hypotheses-based regression [53]; Local binary feature [54]; Incremental face alignment [55]; Coarse-to-fine shape search [56]
<i>Deep neural networks<sup>b</sup></i>	
Deep CNNs	Deep convolutional network cascade [57]; Tasks-constrained deep convolutional network [58]; Deep Cascaded Regression [59]
Other deep networks	Coarse-to-fine Auto-encoder Networks (CFAN) [60]; Deep face shape model [37]

<sup>a</sup> Classic Constrained Local Models (CLMs) typically refer to the combination of local detector for each facial point and the parametric Point Distribution Model [61, 62, 29]. Here we extend the range of CLMs by including some methods based on other shape Models (i.e., exemplar-based model [11]). In particular, we will show that the exemplar-based method [11] can also be interpreted under the conventional CLM framework.

<sup>b</sup> We note that some deep learning-based systems can also be placed in other categories. For instance, some systems are constructed in a cascade manner [60, 59, 63], and hence can be naturally categorized as cascaded regression. However, to highlight the increasing important role of deep learning techniques for face alignment, we organize them together for more systematic introduction and summarization.

Table 2: Overview of the six classes of discriminative methods in our taxonomy.

	Appearance model	Shape model	Highlights of the method
<i>Constrained local models</i>	Independently trained local detector that computes a pseudo probability of the target point occurring at a particular position.	Point Distribution Model; Exemplar model, etc. <sup>a</sup>	The local detectors are first correlated with the image to yield a filter response for each facial point, and then shape optimization is performed over these filter responses.
<i>Constrained local regression</i>	Independently trained local regressor that predicts a distance vector relating to a patch location.	Markov Random Fields to model the relations between relative positions of pairs of points.	Graph model is used to constrain the search space of local regressors by exploiting the constellations that facial points can form.
<i>Deformable part models</i>	Part-based appearance model that computes the appearance evidence for placing a template for a facial part.	Tree-structured models that are easier to optimize than dense graph structures.	All parameters of the appearance model and shape model are discriminatively learned in a max-margin structured prediction framework; efficient dynamic programming algorithms can be used to find globally optimal solutions.

<i>Ensemble regression-voting</i>	Image patches to cast votes for all facial points relating to the patch centers; Local appearance features centered at facial points.	Implicit shape constraint that is naturally encoded into the multi-output function (e.g., regression tree).	Votes from different regions are ensemble to form a robust prediction for the face shape.
<i>Cascaded regression</i>	Shape-indexed feature that is related to current shape estimate (e.g., concatenated image patches centered at the facial points).	Implicit shape constraint that is naturally encoded into the regressor in a cascaded learning framework.	Cascaded regression typically starts from an initial shape (e.g., mean shape), and refines the holistic shape through sequentially trained regressors.
<i>Deep neural networks</i>	Whole face region that is typically used to estimate the whole face shape jointly; Shape-indexed feature <sup>b</sup>	Implicit shape constraint that is encoded into the networks since all facial points are predicted simultaneously.	Deep network is a good choice to model the nonlinear relationship between the facial appearance and the shape update. Among others, deep CNNs have the capacity to learn highly discriminative features for face alignment.

<sup>a</sup> Constrained Local Models (CLMs) typically employ a parametric (PCA-based) shape model [29], but we will show that the exemplar-based method [11] can also be derived from the CLM framework. Furthermore, we extend the range of CLMs by including some methods that combine independently local detector and other face shape model [35, 36, 37].

<sup>b</sup> Some deep network-based systems follow the cascaded regression framework, and use the shape-indexed feature [60].

## Cascaded Regression

- This method starts from the average face shape.
- Estimates the amount of movement based on the image features for each detected point.
- Iteratively moves the detected points to obtain the predicted face shape.

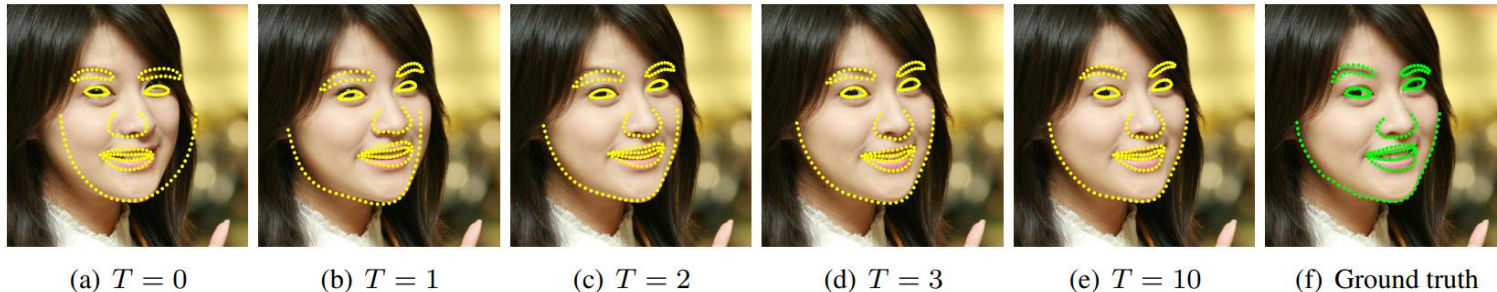


Figure 2. Landmark estimates at different levels of the cascade initialized with the mean shape centered at the output of a basic Viola & Jones[17] face detector. After the first level of the cascade, the error is already greatly reduced.

In the training phase, the stage regressors ( $\mathcal{R}^1, \dots, \mathcal{R}^T$ ) are sequentially learnt to reduce the alignment errors on training set, during which geometric constraints among points are *implicitly* encoded.

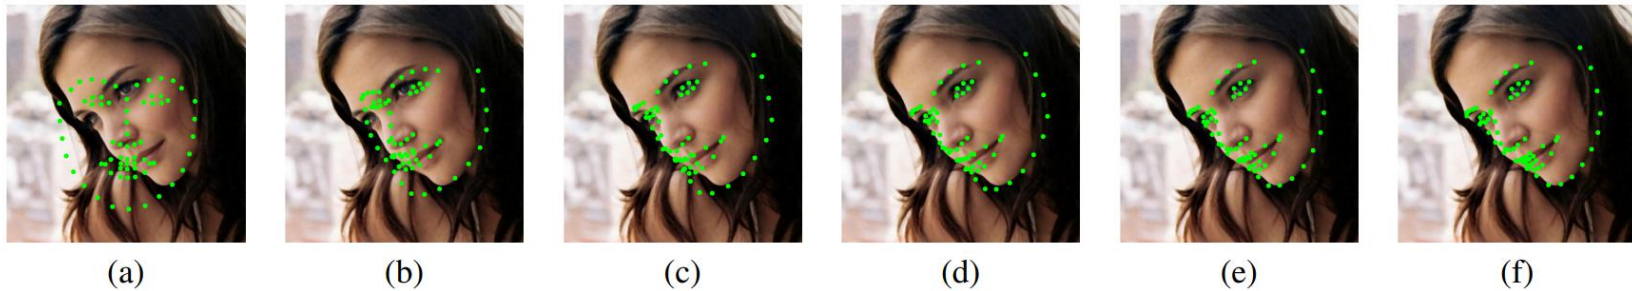


Figure 5: Illustration of face alignment results in different stages of cascaded regression (Fig. 1 in [51]). The shape estimate is initialized and iteratively updated through a cascade of regression trees: (a) initial shape estimate, (b)-(f) shape estimates at different stages.

In OpenCV, we can use **Facemark API and pre-trained model**

**C++ Documentation:**

[https://docs.opencv.org/4.5.5/db/dd8/classcv\\_1\\_1face\\_1\\_1Facemark.html](https://docs.opencv.org/4.5.5/db/dd8/classcv_1_1face_1_1Facemark.html)

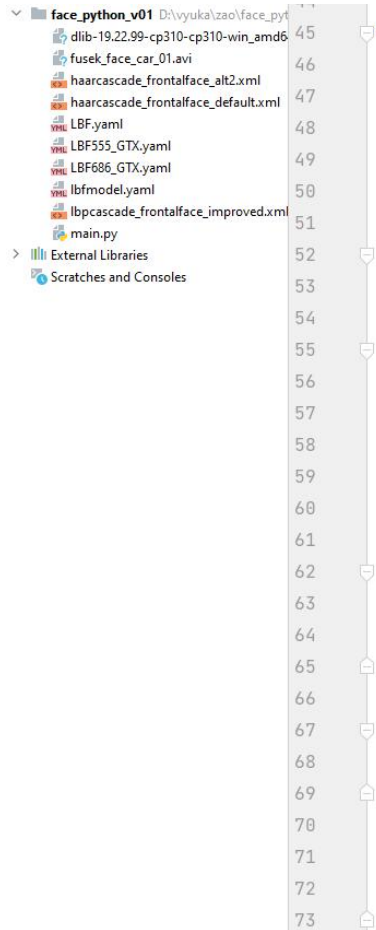
**Python interface:**

<https://gist.github.com/saiteja-talluri/1d0e4fc4c75774b936b99c7c52b65fe6>

<https://github.com/saiteja-talluri/GSoC-OpenCV>

**More about models and methods:**

<https://towardsdatascience.com/faster-smoother-smaller-more-accurate-and-more-robust-face-alignment-models-d8cc867efc5>



```

45 def facial_landmark():
46     cv2.namedWindow("face_detect", 0)
47     video_cap = cv2.VideoCapture("fusek_face_car_01.avi")
48     face_cascade = cv2.CascadeClassifier("lbpcascade_frontalface_improved.xml")
49     face_mark = cv2.face.createFacemarkLBF()
50     face_mark.loadModel("LBF555_GTX.yaml")
51
52 while True:
53     ret, frame = video_cap.read()
54     paint_frame = frame.copy()
55     if ret is True:
56         faces = face_cascade.detectMultiScale(frame,
57                                             scaleFactor=1.1,
58                                             minNeighbors=2,
59                                             minSize=(100, 100),
60                                             maxSize=(500, 500))
61
62         if len(faces) > 0:
63             status, landmarks = face_mark.fit(frame, faces)
64             for f in range(len(landmarks)):
65                 cv2.face.drawFacemarks(paint_frame, landmarks[f], (255, 255, 255))
66
67         for one_face in faces:
68             cv2.rectangle(paint_frame, one_face, (0, 0, 255), 12)
69             cv2.rectangle(paint_frame, one_face, (255, 255, 255), 4)
70
71     cv2.imshow("face_detect", paint_frame)
72     if cv2.waitKey(2) == ord("q"):
73         break

```



Alternatively, we can use <http://dlib.net/>

**Dlib** [27] is an open-source machine learning library. Among others, it has **Ensemble of Regression Trees (ERT)** [26] facial landmark detection algorithm, which is a cascade, based on gradient boosting. The authors use a “mean” face template as an initial approximation, then the template is refined over several iterations. The algorithm requires the face to be first detected in the frame (Viola-Jones [28] face detector is used). Note, that most facial landmark detection algorithms require face to be first detected. High speed is the main advantage of ERT (according to the authors, around 1 millisecond per face). The library contains ERT implementation, trained on 300W dataset. The algorithm is still actively used in the modern research thanks to an open implementation and speed. However, not so long ago it has been shown that neural networks are preferred in terms of quality for faces with large pose [29]. Mobile-friendly implementations of ERT are available.

The screenshot shows the Dlib C++ Library website. At the top, it says "Dlib C++ Library" and "ENHANCED BY Google". Below this is a search bar. The main content is a "Machine Learning Machine Learning Guide" diagram, which is a complex flowchart showing various machine learning tasks and their relationships. The diagram is divided into several sections: Classification, Data Transformations, Structured Prediction, Clustering, Regression, and Matrix Random Fields. A "START" icon is placed in the center of the diagram. To the left of the diagram is a sidebar with "The Library" and "Help/Info" sections. To the right is a "Primary Algorithms" section.

**The Library**

- Algorithms
- API Wrappers
- Bayesian Nets
- Compression
- Containers
- Graph Tools
- Image Processing
- Linear Algebra
- Machine Learning
- Metaprogramming
- Miscellaneous
- Networking
- Optimization
- Parsing

**Help/Info**

- Dlib Blog
- Examples: C++
- Examples: Python
- FAQ
- Home
- How to compile
- How to contribute
- Index
- Introduction
- License
- Python API
- Suggested Books

**Machine Learning Machine Learning Guide**

Dlib contains a wide range of machine learning algorithms. All designed to be highly modular, quick to execute, and simple to use via a clean and modern C++ API. It is used in a wide range of applications including robotics, embedded devices, mobile phones, and large high performance computing environments. If you use dlib in your research please cite:

**Primary Algorithms**

**Binary Classification**

- auto\_train\_rf\_classifier
- rvm\_trainer
- svm\_c\_ekm\_trainer
- svm\_c\_linear\_dcd\_trainer
- svm\_c\_linear\_trainer
- svm\_c\_trainer
- svm\_nu\_trainer
- svm\_pegasos
- train\_probabilistic\_decision\_function

**Multiclass Classification**

- one\_vs\_all\_trainer
- one\_vs\_one\_trainer
- svm\_multiclass\_linear\_trainer

**Regression**

- kris
- krr\_trainer
- mip
- random\_forest\_regression\_trainer
- rbf\_network\_trainer
- rls
- rr\_trainer
- rvm\_regression\_trainer
- svr\_linear\_trainer

## DLIB (to build in Windows, we need VS C++ and Python)

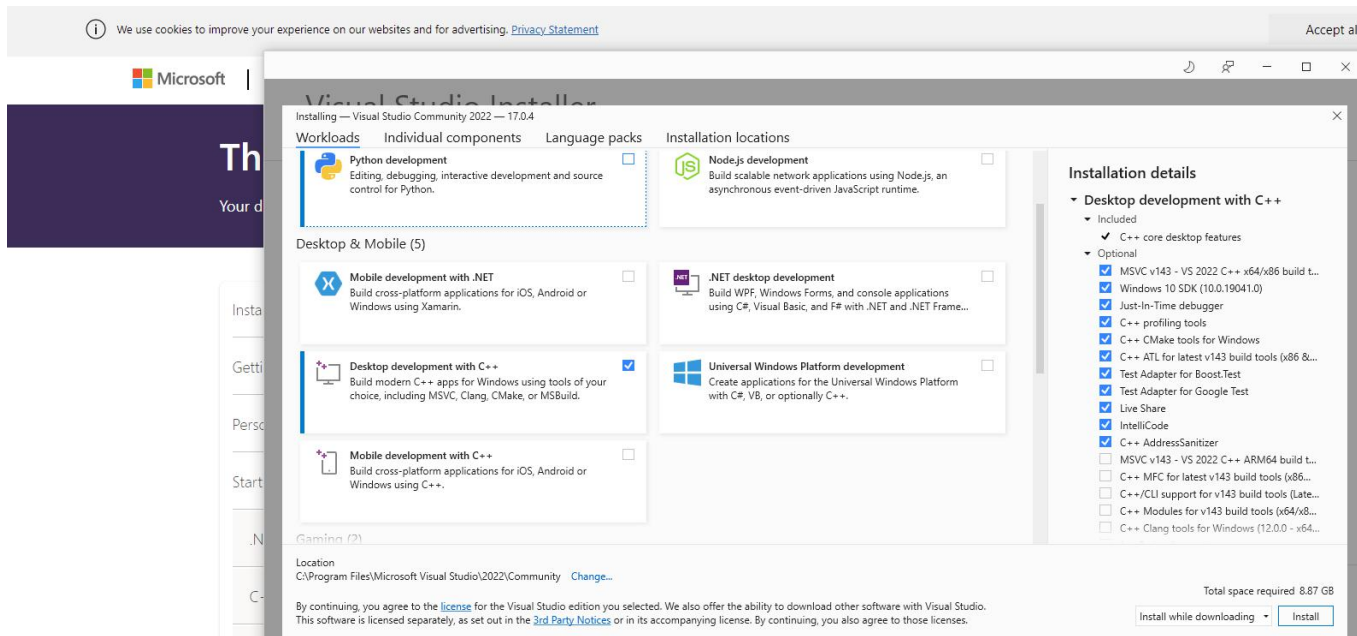
In my case, Dlib installation using pip in PyCharm works with **Python 3.7.9**

It means that you need to create a new python virtual venv with this python versio + pip install cmake

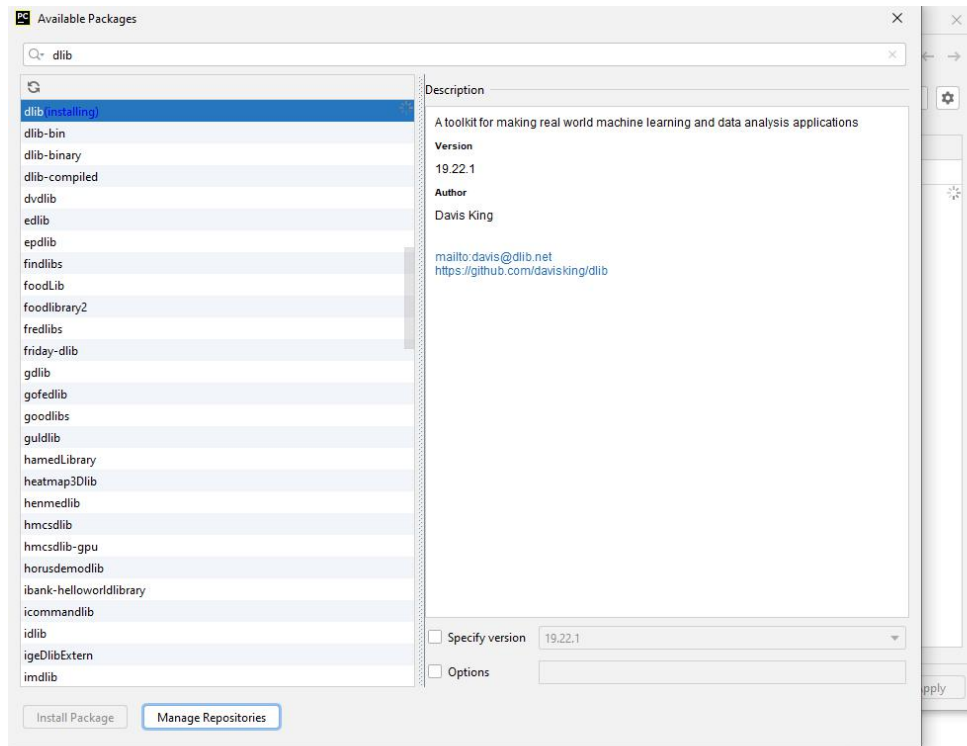
Python 3.7.9 - Aug. 17, 2020

Note that Python 3.7.9 *cannot* be used on Windows XP or earlier.

- Download [Windows help file](#)
- Download [Windows x86-64 embeddable zip file](#)
- Download [Windows x86-64 executable installer](#)
- Download [Windows x86-64 web-based installer](#)
- Download [Windows x86 embeddable zip file](#)
- Download [Windows x86 executable installer](#)
- Download [Windows x86 web-based installer](#)



## DLIB



```
> while True > if ret is True
hon Console [x] Terminal
e-built Python packages shared indexes // Always download // Download once // Don't show again // Configure... (8 minutes ago)
Installing package 'dlib'...
```

The screenshot shows the dlib.net website. On the left is a navigation menu with various categories. The 'Face Landmark Detection' link is highlighted with a green box. The main content area on the right contains text about open source licensing, instructions on how to contribute, and a list of major features.

Not secure | dlib.net

Compression  
Containers  
Graph Tools  
Image Processing  
Linear Algebra  
Machine Learning  
Metaprogramming  
Miscellaneous  
Networking  
Optimization  
Parsing

**Help/Info**  
Dlib Blog  
Examples: C++  
Examples: Python  
Binary Classification  
CNN Face Detector  
Face Alignment  
Face Clustering  
Face Detector  
Face Jittering/Augmentation  
**Face Landmark Detection**  
Face Recognition  
Find Candidate Object Locations  
Global Optimization  
Linear Assignment Problems  
Sequence Segmenter  
Structural Support Vector Machines  
SVM-Rank  
Train Object Detector  
Train Shape Predictor  
Video Object Tracking  
FAQ

environments. Dlib's [open source licensing](#) allows you to use it in any application, free of charge.

To follow or participate in the development of dlib subscribe to [dlib on github](#). Also be sure to read the [how to contribute](#) page if you intend to submit code to the project.

To quickly get started using dlib, follow these instructions to [build dlib](#).

## Major Features

- **Documentation**
  - Unlike a lot of open source projects, this one provides complete and precise documentation for every class and function. There are also debugging modes that check the documented preconditions for functions. When this is enabled it will catch the vast majority of bugs caused by calling functions incorrectly or using objects in an incorrect manner.
  - Lots of example programs are provided
  - *I consider the documentation to be the most important part of the library.* So if you find anything that isn't documented, isn't clear, or has out of date documentation, tell me and I will fix it.
- **High Quality Portable Code**
  - Good unit test coverage. The ratio of unit test lines of code to library lines of code is about 1 to 4.
  - The library is tested regularly on MS Windows, Linux, and Mac OS X systems. However, it should work on any POSIX system and has been used on Solaris, HPUX, and the BSDs.
  - No other packages are required to use the library. Only APIs that are provided by an out of the box OS are needed.
  - There is no installation or configure step needed before you can use the library. See the [How to compile](#) page for details.

```
5
6 def facial_landmark_dlib():
7     detector = dlib.get_frontal_face_detector()
8     landmark_predictor = dlib.shape_predictor("shape_predictor_68_face_landmarks.dat")
9     video_cap = cv2.VideoCapture("fusek_face_car_01.avi")
10
11 while video_cap.isOpened():
12     ret, frame = video_cap.read()
13     if ret is True:
14         paint_frame = frame.copy()
15         faces = detector(frame, 0)
16
17         for i, f in enumerate(faces):
18             pt1 = (f.left(), f.top())
19             pt2 = (f.right(), f.bottom())
20             print("id-top-bot {} - {} - {}".format(i, pt1, pt2))
21             cv2.rectangle(paint_frame, pt1, pt2, (0, 0, 255), 12)
22             cv2.rectangle(paint_frame, pt1, pt2, (255, 255, 255), 4)
23
24             shape = landmark_predictor(frame, f)
25             print(shape.parts())
26             for ip, p in enumerate(shape.parts()):
27                 if ip in [20, 25, 30]:
28                     point = (p.x, p.y)
29                     cv2.circle(paint_frame, point, 5, (255, 255, 255), -1)
30                     cv2.circle(paint_frame, point, 2, (0, 0, 0), -1)
31
32             cv2.imshow("opencv_frame", paint_frame)
33             if cv2.waitKey(2) == ord("q"):
34                 break
35     else:
36         break
```

# Facial Landmark Detection MediaPipe

MediaPipe

Solutions Framework

Search

English

## Solutions

Overview Guide Examples API

Filter

Studio

Vision tasks

- Object detection
- Image classification
- Image segmentation
- Interactive segmentation
- Gesture recognition
- Hand landmark detection
- Image embedding
- Face detection
- Face landmark detection
  - Overview
  - Android
  - Web
  - Python
- Pose landmark detection
- Face stylization
- Holistic landmark detection

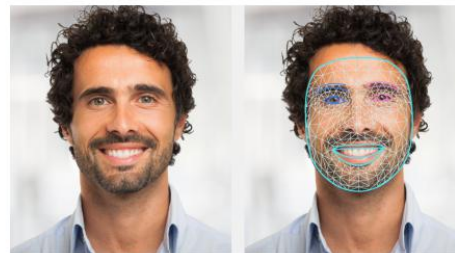
**Attention:** This MediaPipe Solutions Preview is an early release. [Learn more](#)

Home > MediaPipe > Solutions > Guide

Was this helpful?

## Face landmark detection guide

The MediaPipe Face Landmarker task lets you detect face landmarks and facial expressions in images and videos. You can use this task to identify human facial expressions, apply facial filters and effects, and create virtual avatars. This task uses machine learning (ML) models that can work with single images or a continuous stream of images. The task outputs 3-dimensional face landmarks, blendshape scores (coefficients representing facial expression) to infer detailed facial surfaces in real-time, and transformation matrices to perform the transformations required for effects rendering.



Try it! →

On this page

[Get Started](#)

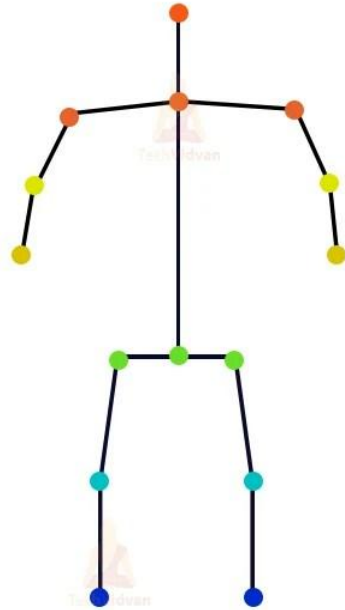
Task details

Features

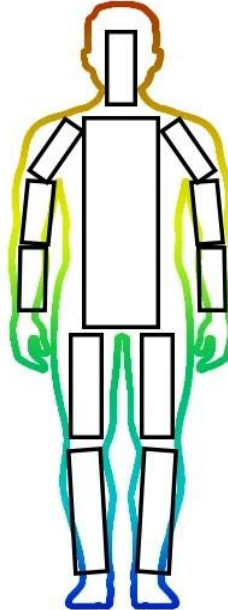
Configurations options

Models

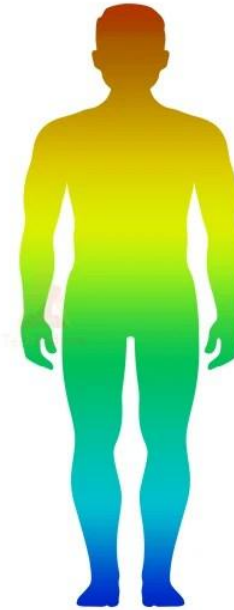
## Types of Human Pose Estimation Models



(a) Kinematic



(b) Planar

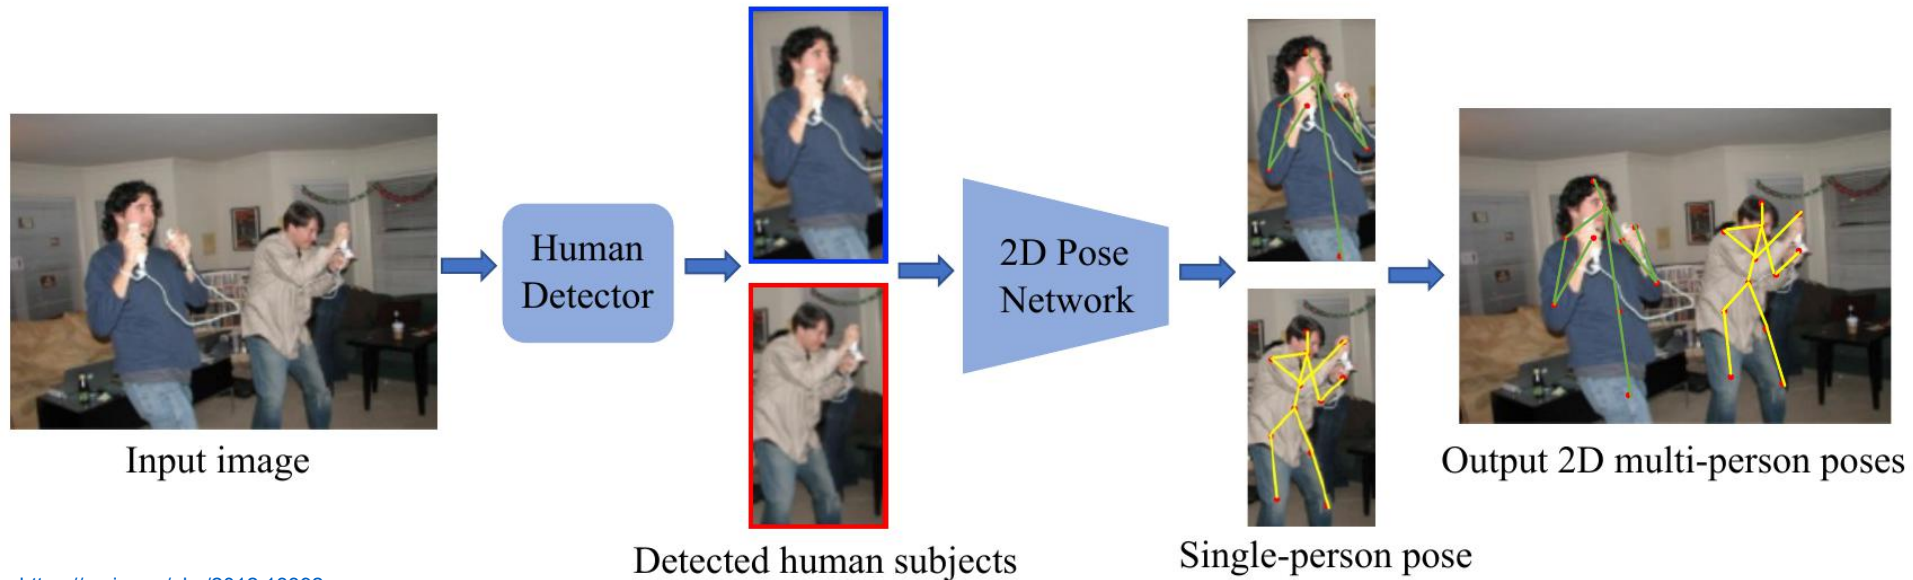


(c) Volumetric

**2D multi-person pose estimation: top-down and bottom-up methods.**

**Top-down** approaches have two sub-tasks:

- (1) human detection and
- (2) pose estimation in the region of a single human.



(a) Top-Down Approaches



**2D multi-person pose estimation: top-down** and **bottom-up** methods.

**Bottom-up** approaches also have two sub-tasks:

- (1) detect all keypoints candidates of body parts and
- (2) associate body parts in different human bodies and assemble them into individual pose representations.



(b) Bottom-Up Approaches