Face Recognition: from EigenFaces to DeepFace

Artem Chernodub

ZZ Photo

IMMSP NASU

AI Ukraine, Kharkiv, 12 September 2015.

Face Recognition

Standard Face Recognition Scenario

Step 1. Face Detection

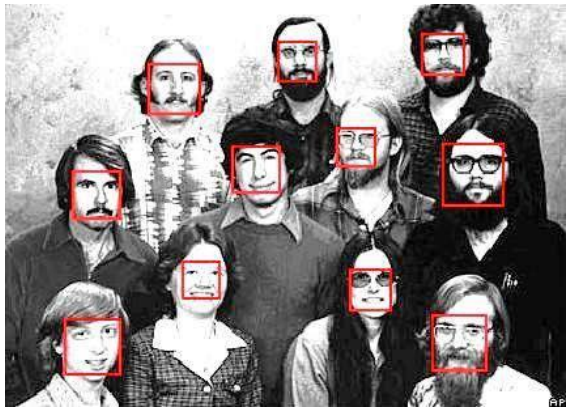

Step 2. Face Alignment

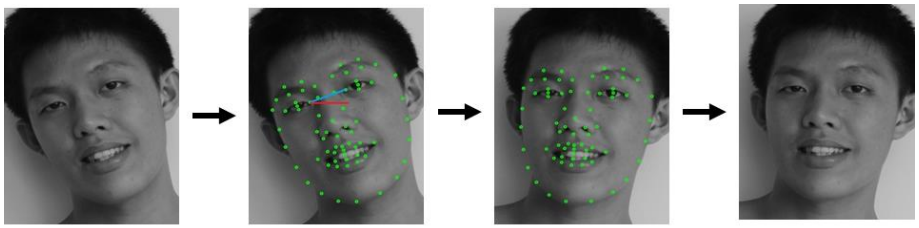

Step 3. Face Matching

Face Matching (Face Recognition)

A) Training the classifier

Features

Classifier

Answer

B) Matching two face images

Recognition of Face Attributes

- man;
- white;
- 45-50 years;
- glasses – no;
- moustache – yes;
- beard – yes.

Beauty Recognition

http://www.linkface.cn/face/beauty

Чернодуб А.Н., Пащенко Ю.А., Головченко К.А. Нейросетевая система определения привлекательности лица человека // «Нейроинформатика-2013», Москва, 21-25 января 2013, с. 254 — 259.

Face Matching

EigenFaces, Principal Component Analysis (PCA) for Face Matching, 1991

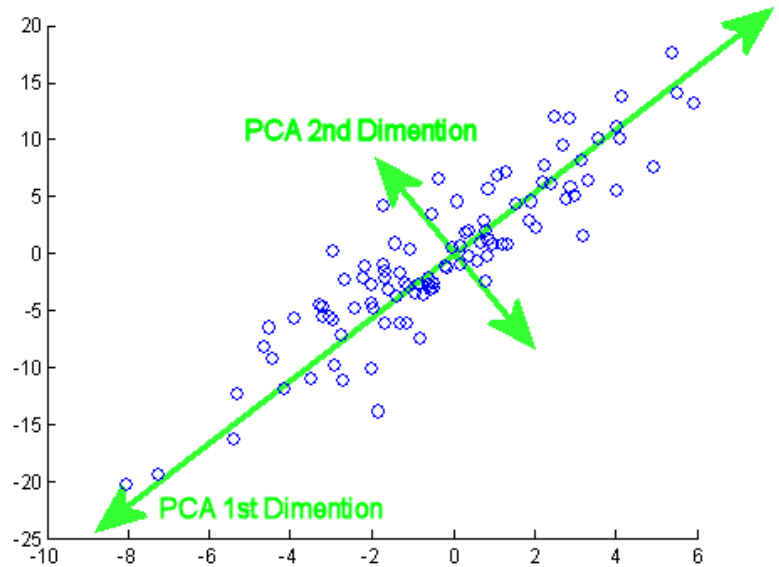

- Well-known method for dimensionality reduction.
- Building features from the training data.
- Impressive visualization of the basis vectors.

M. Turk, A. Pentland. Face Recognition using EigenFaces // Journal of cognitive neuroscience 3 (1), p. 71-86.

Local Binary Pattern Histograms (LBPH), 2006

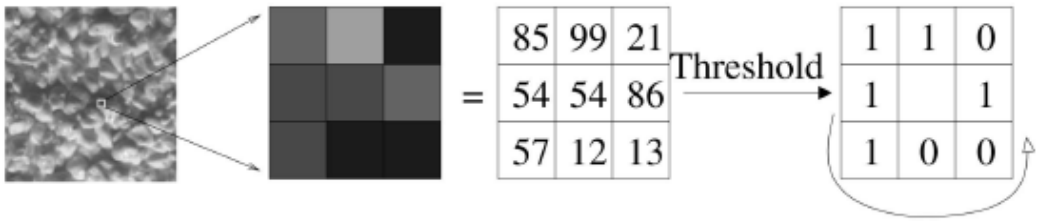

- Fast & robust features for texture descriptions applied to face recognition.
- χ^2 distance for comparison of histograms.
- Matrix of importance regions (weighted χ^2).

T. Ahonen, A. Hadid. Face Description with Local Binary Patterns: Application to Face Recognition // IEEE Transactions on Pattern Analysis and Machine Intelligence, Vol. 28, No. 12, p. 2037-2041.

Deep Learning = Learning of Representations (Features)

<u>The traditional model of pattern recognition (since the late 50's):</u>
fixed/engineered features + trainable classifier

→ Hand-crafted Feature Extractor → Trainable Classifier

<u>End-to-end learning / Feature learning / Deep learning:</u>
trainable features + trainable classifier

→ Trainable Feature Extractor → Trainable Classifier

Deep Feedforward Neural Networks

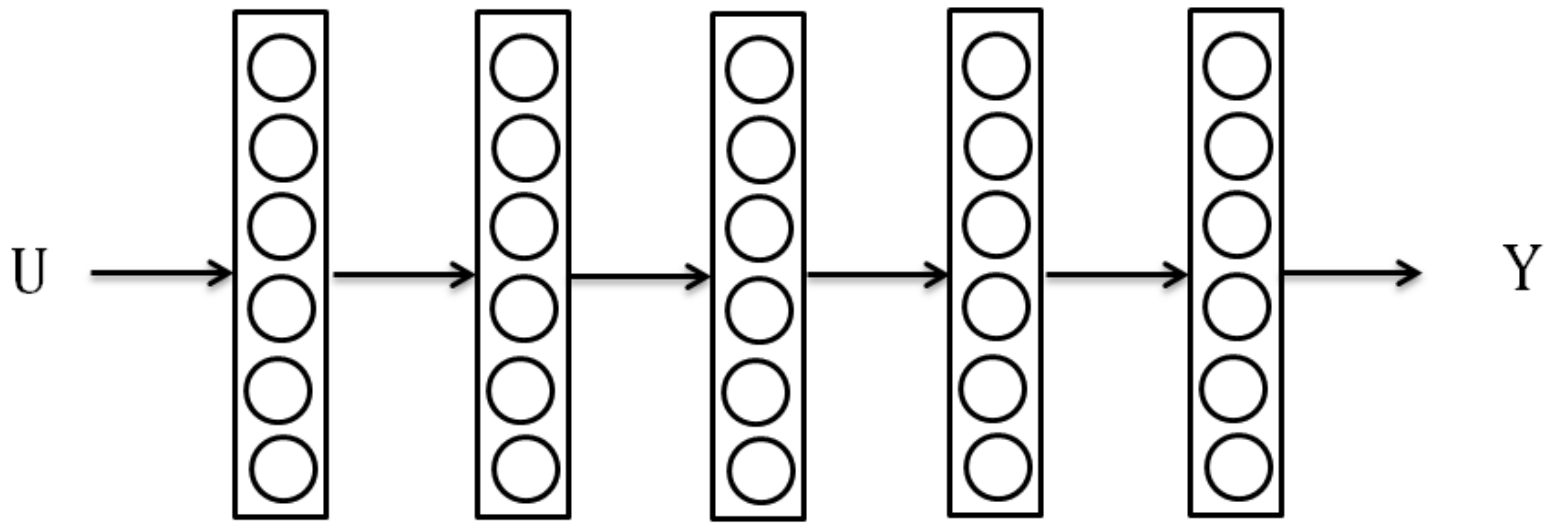

- 2-stage training process: i) unsupervised pre-training; ii) fine tuning (vanishing gradients problem is beaten!).
- Number of hidden layers > 1 (usually 6-9).
- 100K – 100M free parameters.
- No (or less) feature preprocessing stage.

FLOPS comparison

Type	Name	Flops	Cost
Mobile	Raspberry Pi 1st Gen, 700 Mhz	0,04 Gflops	$35
Mobile	Apple A8	1,4 Gflops	$700 (in iPhone 6)
CPU	Intel Core i7-4930K (Ivy Bridge), 3.7 GHz	140 Gflops	$700
CPU	Intel Core i7-5960X (Haswell), 3.0 GHz	350 Gflops	$1300
GPU	NVidia GTX 980	4612 Gflops (single precision), 144 Gflops (double precision)	$600 + cost of PC (~$1000)
GPU	NVidia Tesla K80	8740 Gflops (single precision), 2910 Gflops (double precision)	$4500 + cost of PC (~1500)

https://ru.wikipedia.org/wiki/FLOPS

Deep Face (Facebook), 2014

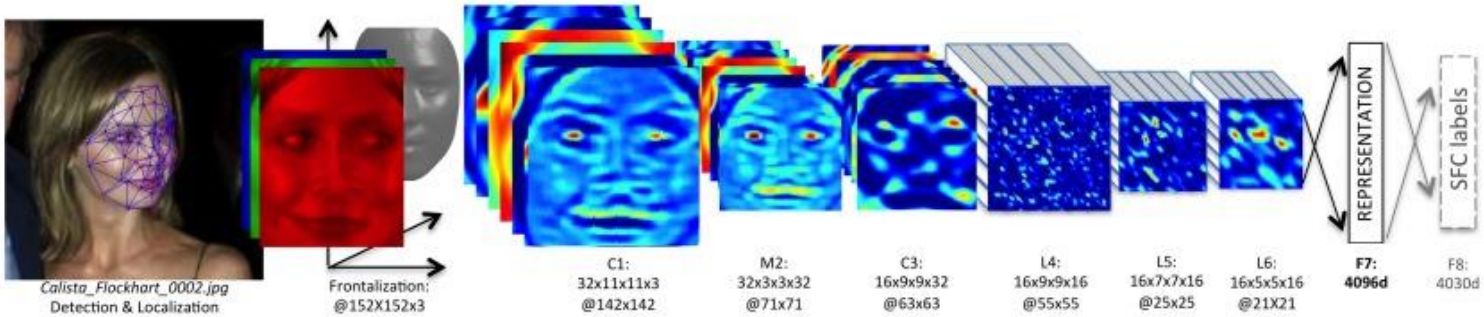

Model	# of parameters	Accuracy, %, LFW
Deep Face Net	128M	97.25
Human level	N/A	~ 97.64

Training data: 4M facial images

Y. Taigman, M. Yang, M.A. Ranzato, L. Wolf. DeepFace: Closing the Gap to Human-Level Performance in Face Verification // CVPR 2014.

Convolutional Neural Networks: Return of Jedi

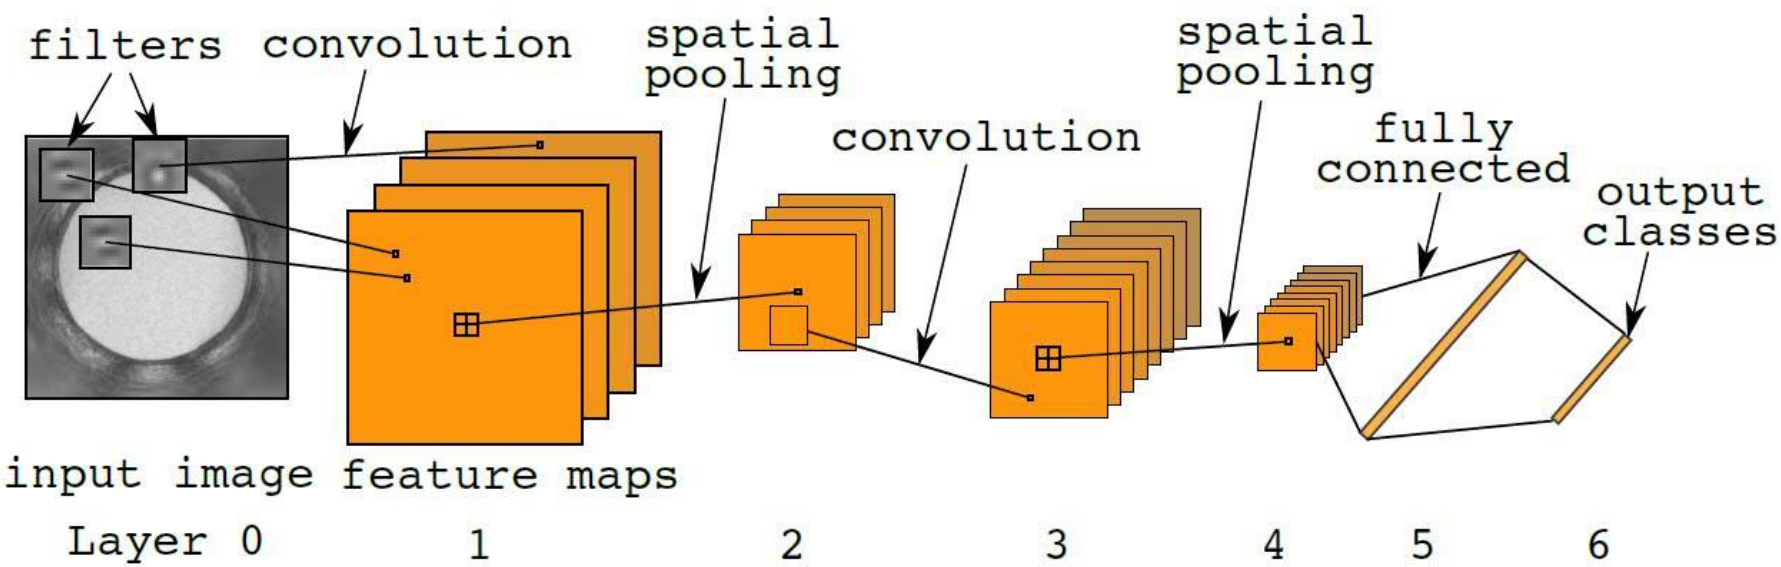

Andrej Karpathy and Fei-Fei. *CS231n: Convolutional Neural Networks for Visual Recognition* http://cs231n.github.io/convolutional-networks

Yoshua Bengio, Ian Goodfellow and Aaron Courville. Deep Learning // An MIT Press book in preparation
http://www-labs.iro.umontreal.ca/~bengioy/DLbook

DeepFace Training Framework

Step 1. Supervised training for identification

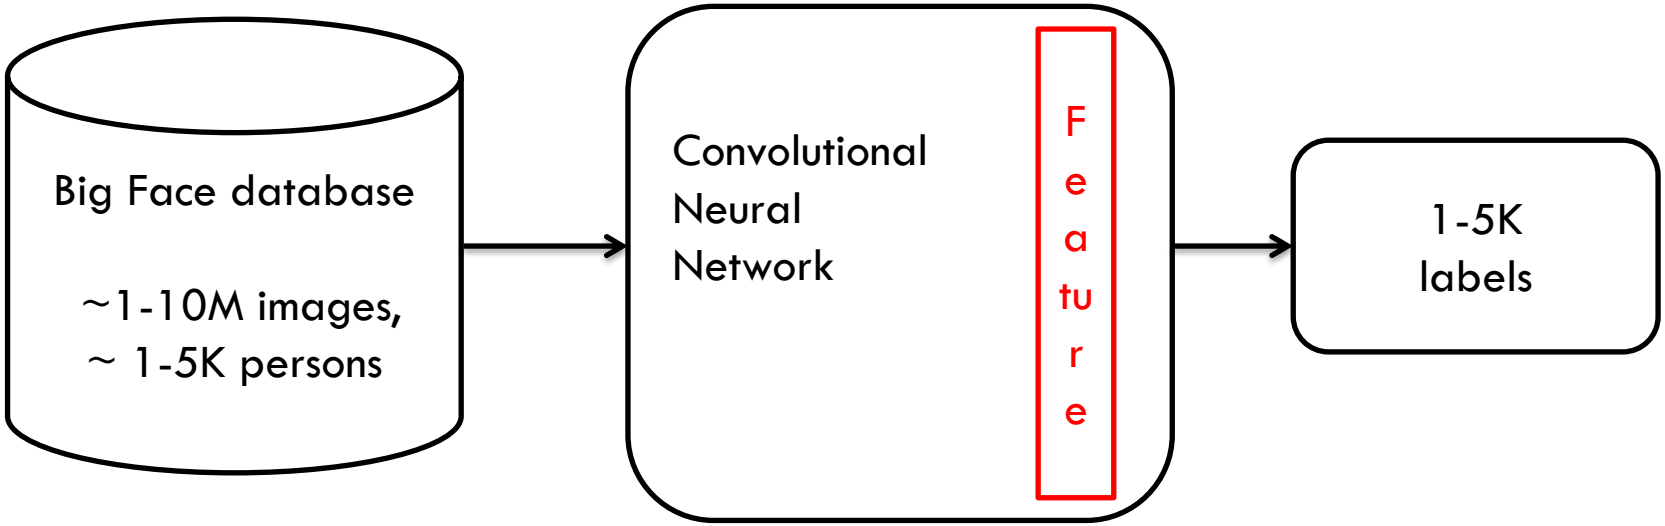

Step 2. Use produced features for face matching

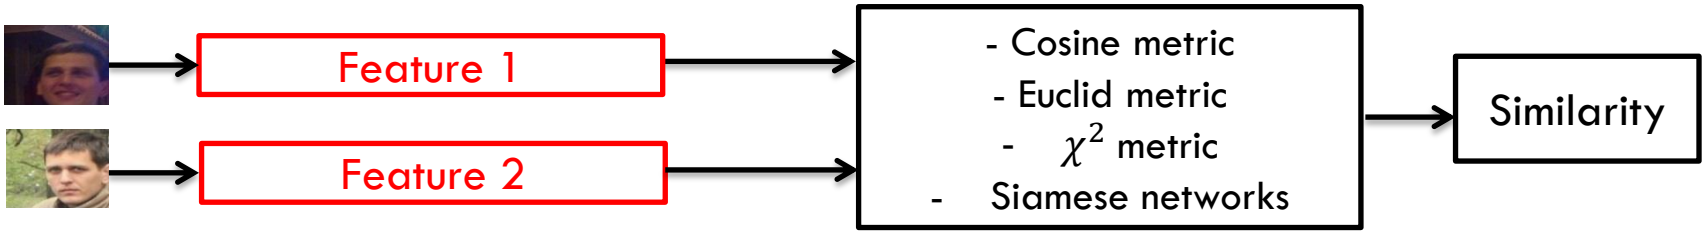

Databases for Deep Face Training

- Facebook's Social Face Classification (SCF) dataset, 2014. 4.4M labeled faces, 4030 people with 800 to 1200 faces each. Private dataset.

- Visual Geometry Group Dataset, Oxford, 2015. 2622 people with 1000 faces each. Private dataset.

- CASIA WebFace Database, 2014. 10575 people, 500K faces. Public dataset. http://www.cbsr.ia.ac.cn/english/CASIA-WebFace-Database.html

AT&T Database (Olivetti, ORL): 1992-1994

- Grayscale images, 92x112.
- 40 subjects, 10 images per subject.
- Different facial expressions, poses, glasses.
- Studio, single session.

F. Samaria, A. Harter. Parameterisation of a Stochastic Model for Human Face Identification // Proceedings of 2nd IEEE Workshop on Applications of Computer Vision, Sarasota FL, December 1994.

http://www.uk.research.att.com/

Face Databases (benchmarks)

FERET dataset, 1994-1996

- 2413 still facial images, representing 856 individuals, studio, two sessions.
- Different facial expressions, poses, light.

P.J. Phillips, H. Moon, P. Rauss, S. A. Rizvi. The FERET September 1996 Database and Evaluation Procedure // First International Conference, AVBPA'97 Crans-Montana, Switzerland, March 12–14, 1997, pp. 395-402

http://www.nist.gov/srd/

Labeled Faces in the Wild dataset, 2007

- Images with faces collected from the web.
- Pairs comparison, restricted mode.
- test: 10-fold cross-validation, 6000 face pairs.

G. B. Huang, M. Ramesh, T. Berg, E. Learned-Miller. Labeled Faces in the Wild: A Database for Studying Face Recognition in Unconstrained Environments // University of Massachusetts, Amherst, Technical Report 07-49, October, 2007.
http://vis-www.cs.umass.edu/lfw/

Results for Face Matching, LWF (option "Labeled outside data")

	Method	Accuracy
1	EigenFaces	60,20%
2	LBP-classic (3K features)	72,43%
3	High-dim LBP (100K features)	95,17%
4	DeepFace	97,25%
5	Human level	97,64%
6	DeepID3	99,53%
7	Baidu	99,77%

http://vis-www.cs.umass.edu/lfw/results.html

Face Detection

Viola-Jones Object Detector

- Very popular for Face Detection.
- Different features except HAAR: LBP, SURF/SIFT, etc.
- Available free in OpenCV library (http://opencv.org).
- OpenCV implementation supports Windows, Linux, Mac OS, iOS (not officially) and Android.
- May be trained with MATLAB tool *traincascade*.

Images pyramid for Viola-Jones

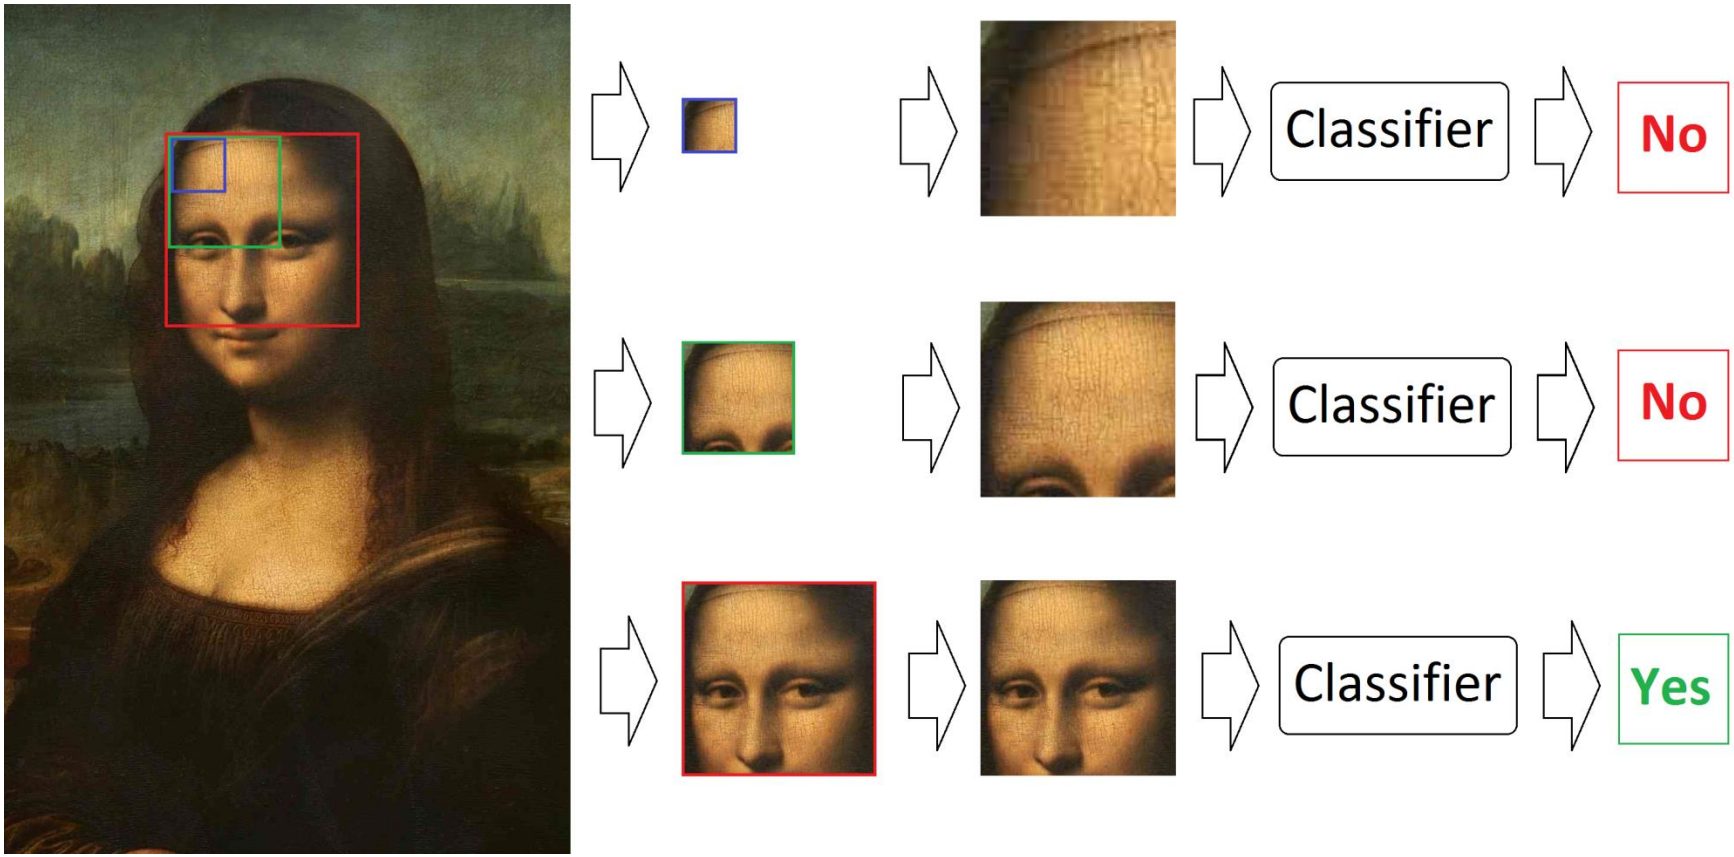

FDDB: Face Detection Data Set and Benchmark

- 5171 faces in a set of 2845 LWF images.
- Special software tool for testing.

Vidit Jain and Erik Learned-Miller. FDDB: A Benchmark for Face Detection in Unconstrained Settings // Technical Report UM-CS-2010-009, Dept. of Computer Science, University of Massachusetts, Amherst. 2010.

http://vis-www.cs.umass.edu/fddb

Viola-Jones Object Detector Classifier Structure

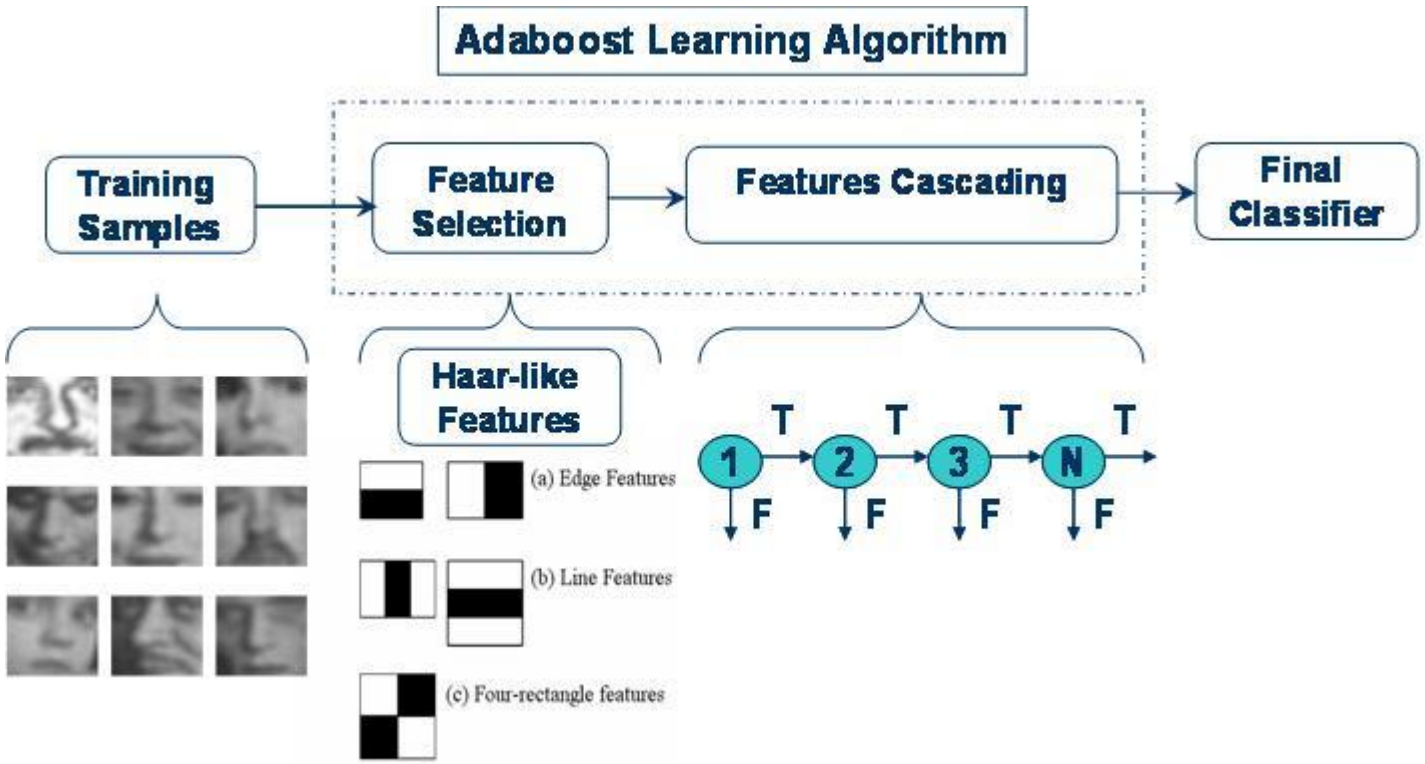

P. Viola, M. Jones. Rapid object detection using a boosted cascade of simple features // Proceedings of the 2001 IEEE Computer Society Conference on Computer Vision and Pattern Recognition, CVPR 2001.

FDDB: state-of-the-art

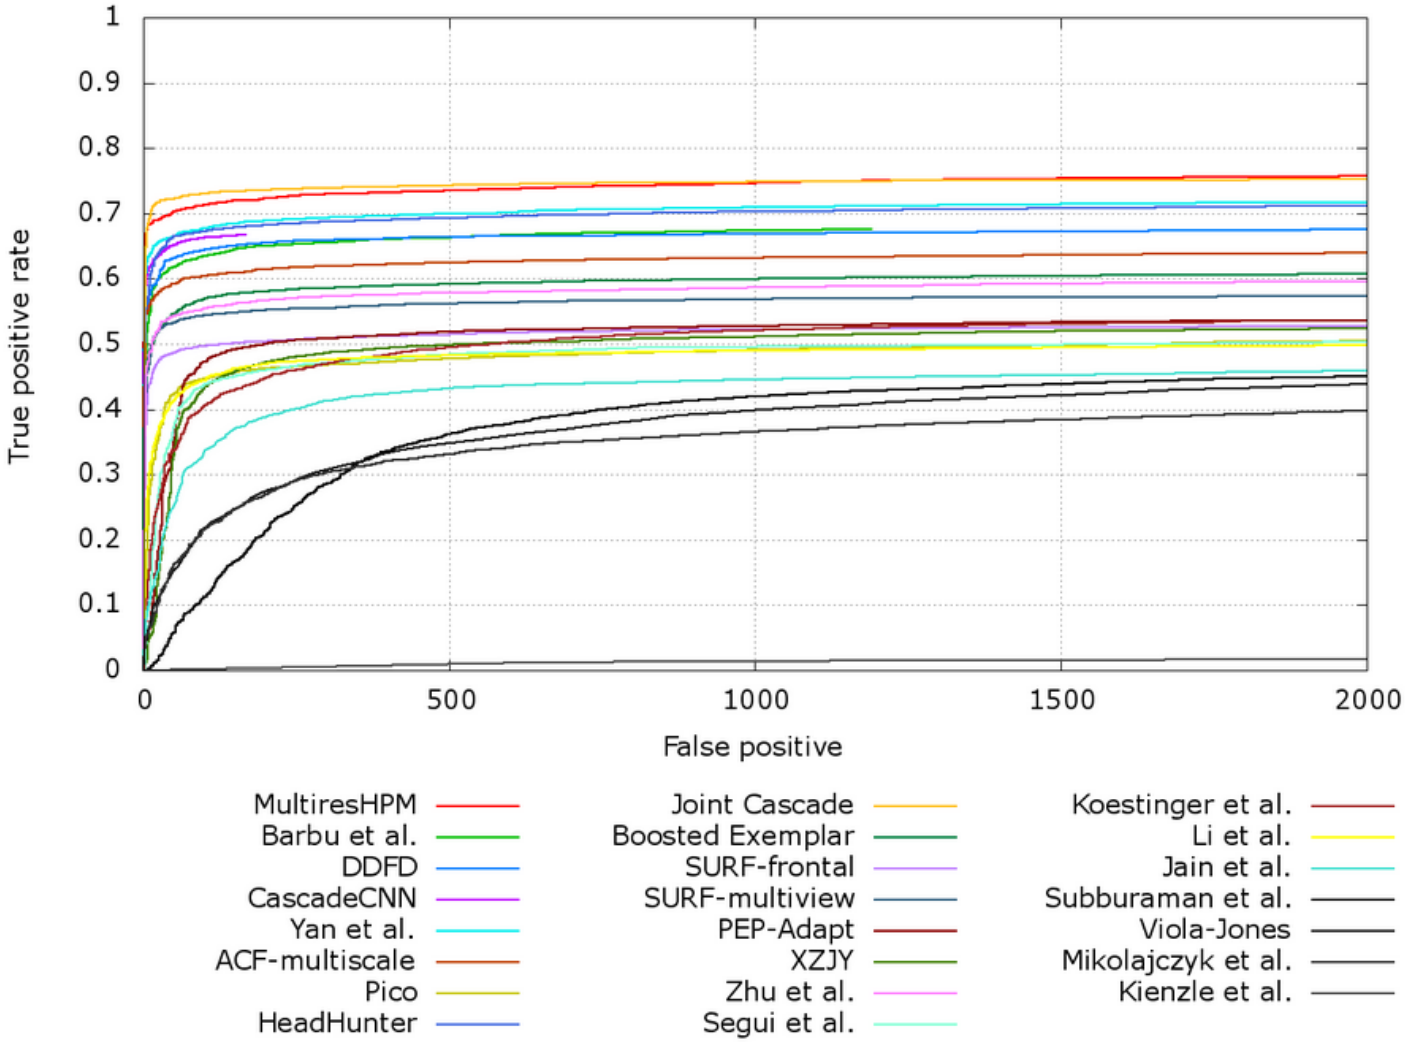

http://vis-www.cs.umass.edu/fddb/results.html

Face Alignment

3D Face Alignment

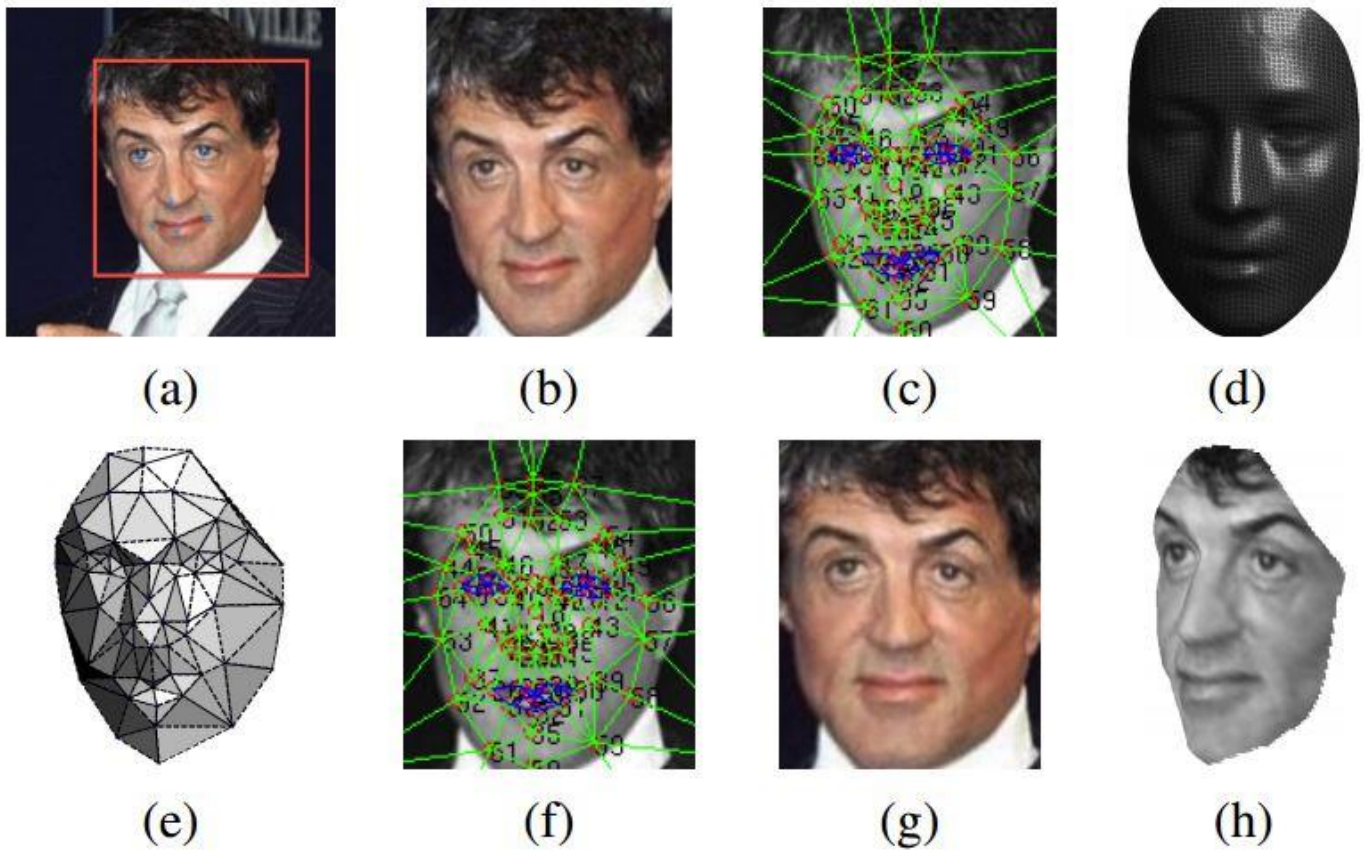

Y. Taigman, M. Yang, M.A. Ranzato, L. Wolf. DeepFace: Closing the Gap to Human-Level Performance in Face Verification // CVPR 2014.

2D Face Alignment

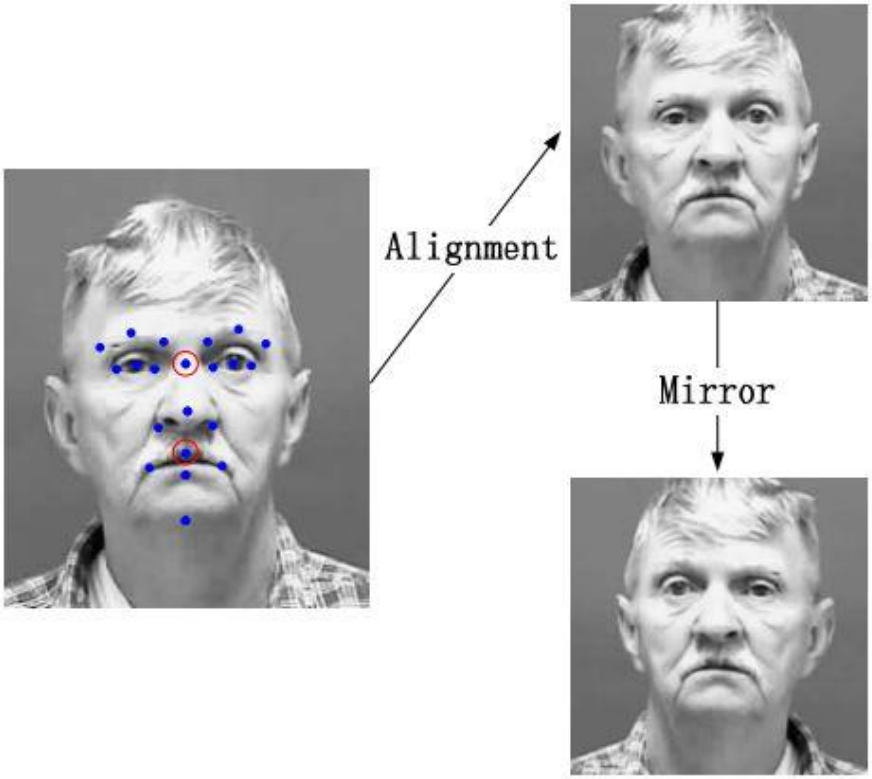

D. Yi, Z. Lei, S. Liao, S.Z. Li. Learning Face Representation from Scratch // CoRR abs/1411.7923 (2014). [CASIA]

Some Computer Vision / Machine Learning tools

- OpenCV, VLFeat – popular Computer Vision libraries.
- mexopencv – MATLAB interface for OpenCV.
- pythonXY – all-in-one Machine Learning Package.
- libSVM – reliable cross-platform Support Vector Machine library.
- NetLab – Shallow Neural Networks & other ML models, from Christopher Bishop's book.
- Caffe, Theano, Torch, matconvnet – frameworks for training Deep Neural Networks.

Thanks!

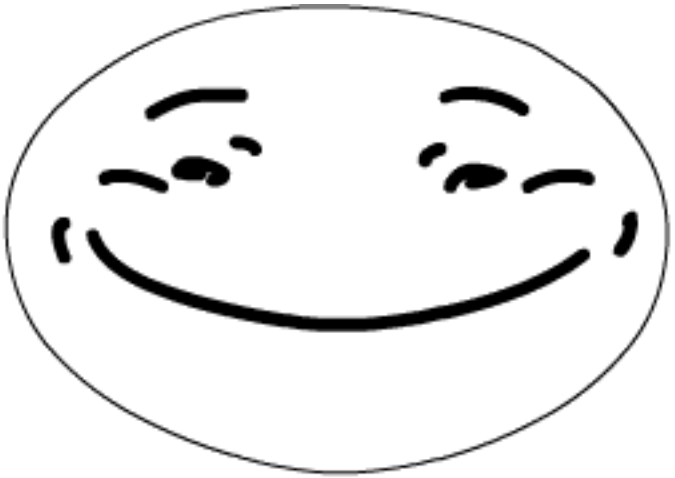

contact: a.chernodub@gmail.com