Face Search at Scale

Anil Jain
Michigan State University

July 10, 2017

Information Content in a Face

Identity: John

Demographics:

Age: ~ 40; gender: male;

ethnicity: white

Attributes:

Hair: Short, Brown

Moustache: Yes

Beard: Yes

Mole: Yes

Scar: Yes

Social cues:

Expression, emotion,…

Outline

- Face Recognition
- Applications
- Challenges
- State of the Art
- Summary

Face Verification

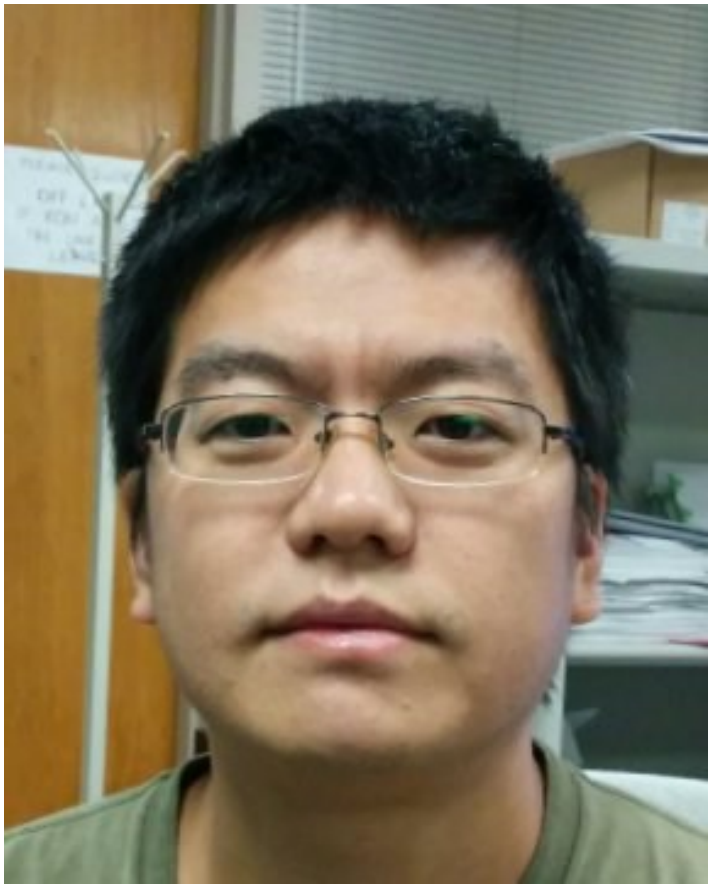

Same Person?

Face Search

Probe *Gallery*

Closed-set v. Open-set Search

Forensic Face Search

Gallery (120M gallery)

Who is this?

Law Enforcement

Search System

One of them?

Forensic examiner

Face Recognition Milestones

1964
Woodrow Bledsoe
Automated face recognition (AFR)

1973
Takeo Kanade
First AFR thesis

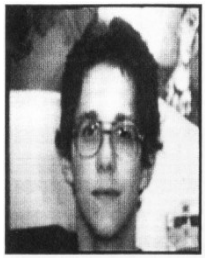

1991
Turk & Pentland
Eigenface

1996
Penev & Atick
Local Feature Analysis

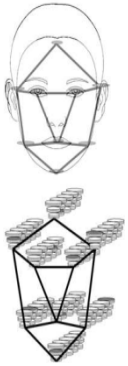

1997
Wiskott et al.
Elastic Bunch Graph Matching

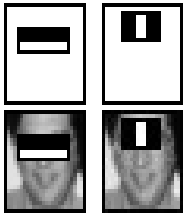

2001
Viola & Jones
Face detector

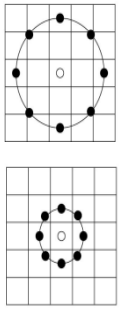

2006
Ahonen et al.
Local Binary Pattern (LBP)

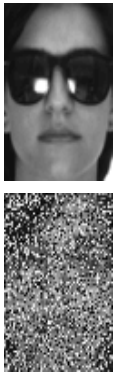

2009
Wright et al.
Sparse representation

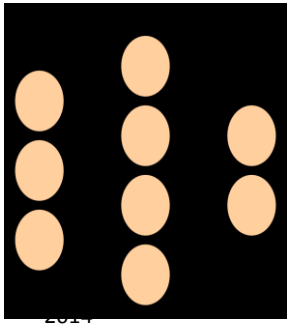

2014
Jia et. al.
Deep Network Library
Caffe

1915
35mm still camera

1991
Kodak
Digital camera
1024p

1990s
Surveillance camera
480p @ 30fps

2000
Sharp
First camera phone
320p

2010
RGB-D camera
Microsoft Kinect
480p @ 30 fps
Depth accuracy:
~ 2 mm @ 1 m distance

2013-2014
Wearable camera
Google Glass
720p @30fps

Nov. 2011
Samsung
Galaxy Nexus
Face Unlock

2015
Google& Intel
Smartphone
RGB-D Camera

2015+
Body Camera
Used by police officers

Growing Interest in Face Recognition

- Technology Drivers
  - Security (covert acquisition, IR, thermal,..)
  - Prevalence of surveillance cameras
  - Mobile phones
  - Social media
- Technology Enablers
  - Processors (2M comparisons/sec/core)
  - Deep networks
  - Large training sets
  - Benchmark datasets of increasing complexity
  - Legacy database s

Applications

Surveillance Cameras Everywhere!

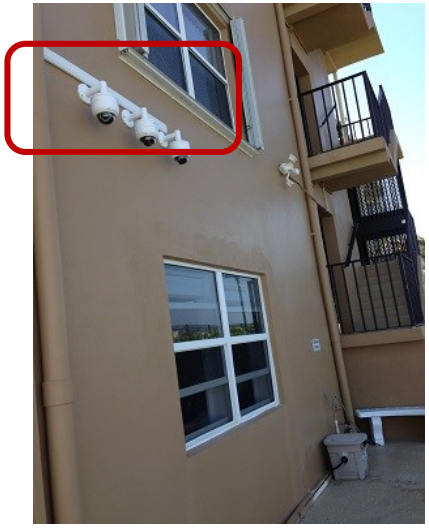

6:00 AM, Home

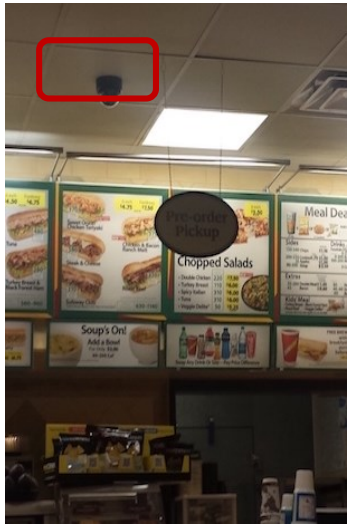

6:15 AM, Fast Food

6:35 AM, ATM

6:45 AM, Gas Station

7:00 AM, Parking Lot

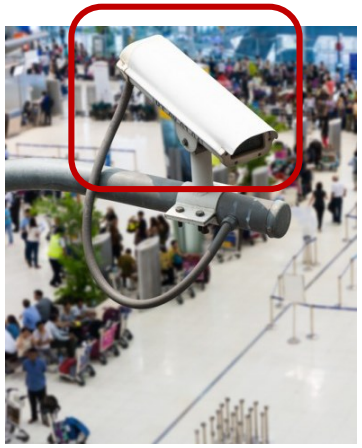

7:10 AM, Airport

7:30 AM, Security

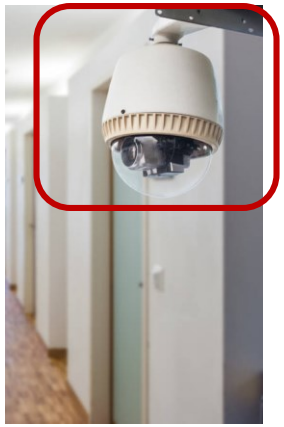

3:00 PM, Hotel

~200 million surveillance cameras; billion of hours of videos/week!

Face Recognition in Video

Widespread looting and rioting

Face recognition lead to many arrests

Extensive CCTV Network

Many suspects could not be identified

2011 London riots

https://www.newscientist.com/article/mg21128266-000-face-recognition-technology-fails-to-find-uk-rioters/

Boston Marathon Bombing

(April 2013)

Tamerlan Tsarnaev

Dzhokhar Tsarnaev

International Border Crossing

SmartGate, Australia & NZ

ePassports from eligible countries

HK-Schenzen border

Fusion of face & fingerprint

Passenger Verification

Matching face image to photo on ID card

川大智胜
WISESOFT

De-duplication: Driver License Database

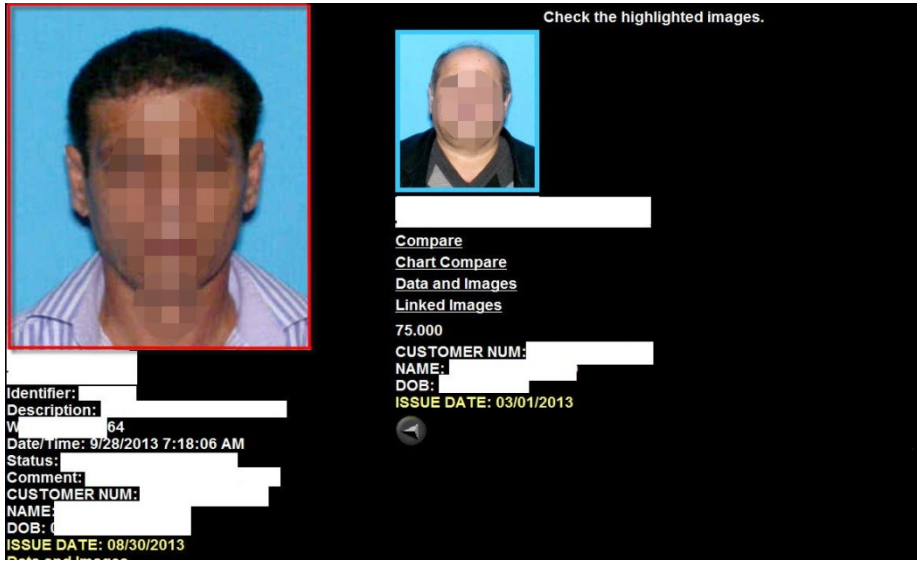

Face-based scrubbing of **13.5M records (~30M photos)** in Michigan DMV database; photos of different subjects in the same record!

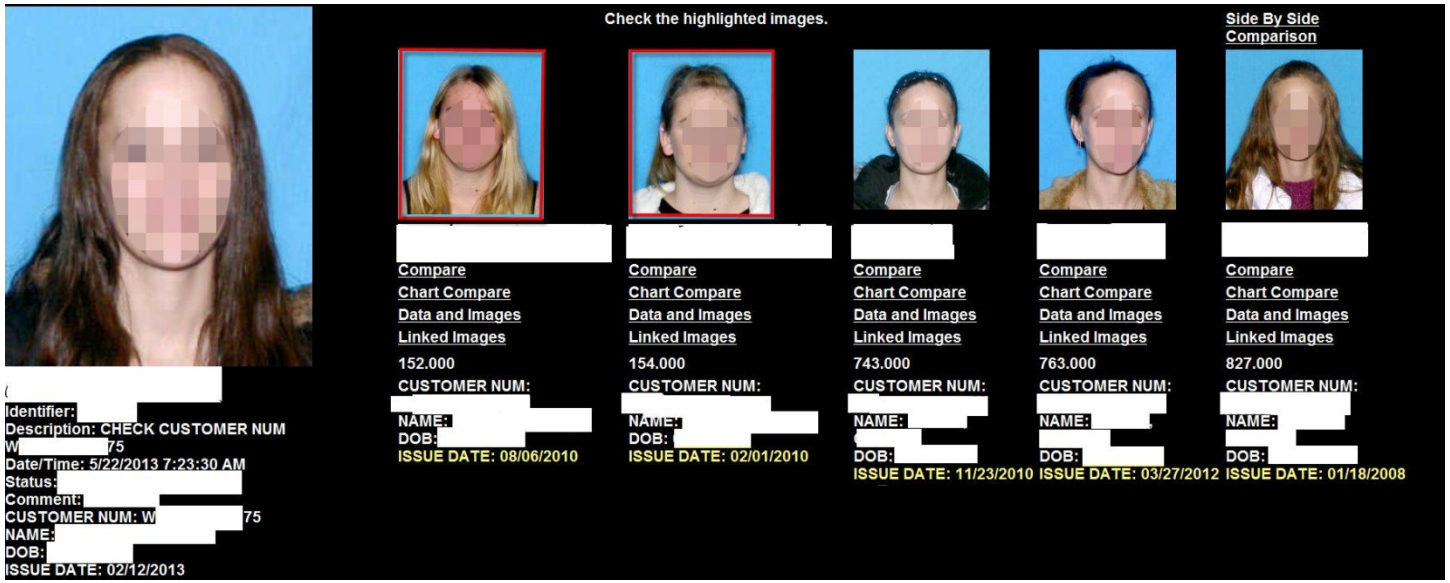

Courtesy: Pete Langenfeld, Michigan State Police

De-duplication

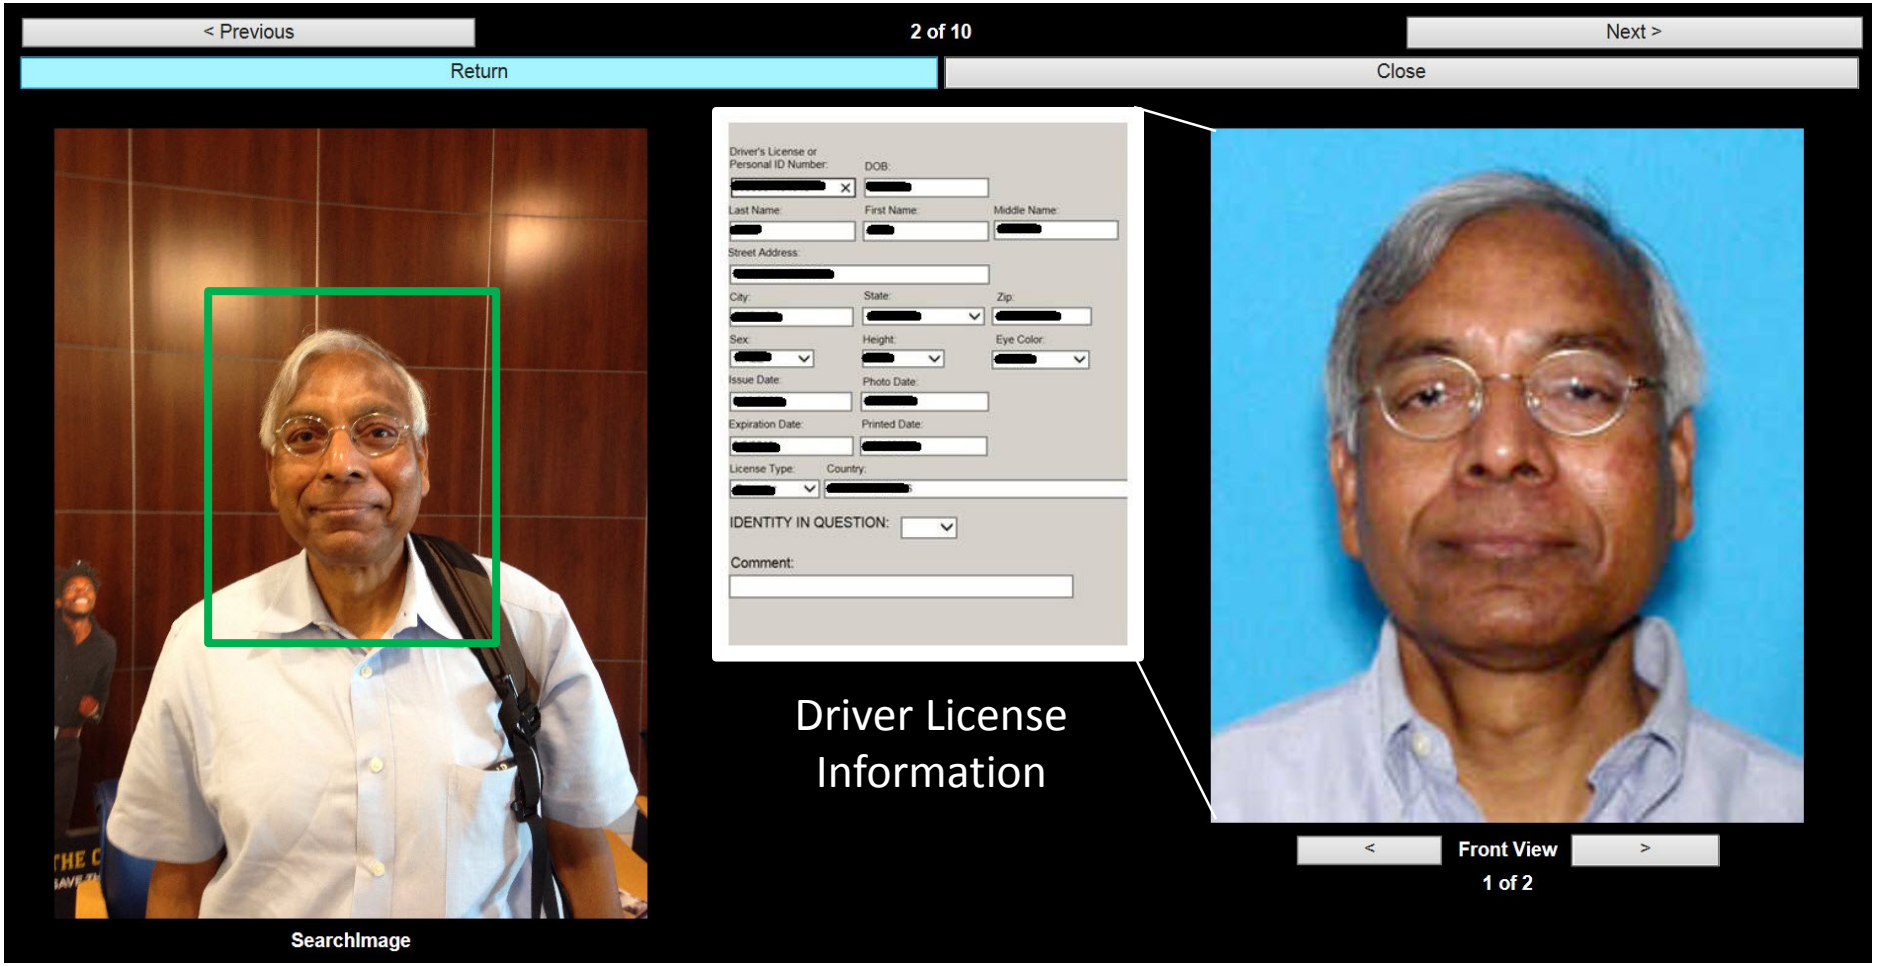

Gallery: 34 million (30M DMV photos, 4M mugshots) 2009 driver license

Smile Makes a difference!

Top-10 retrievals

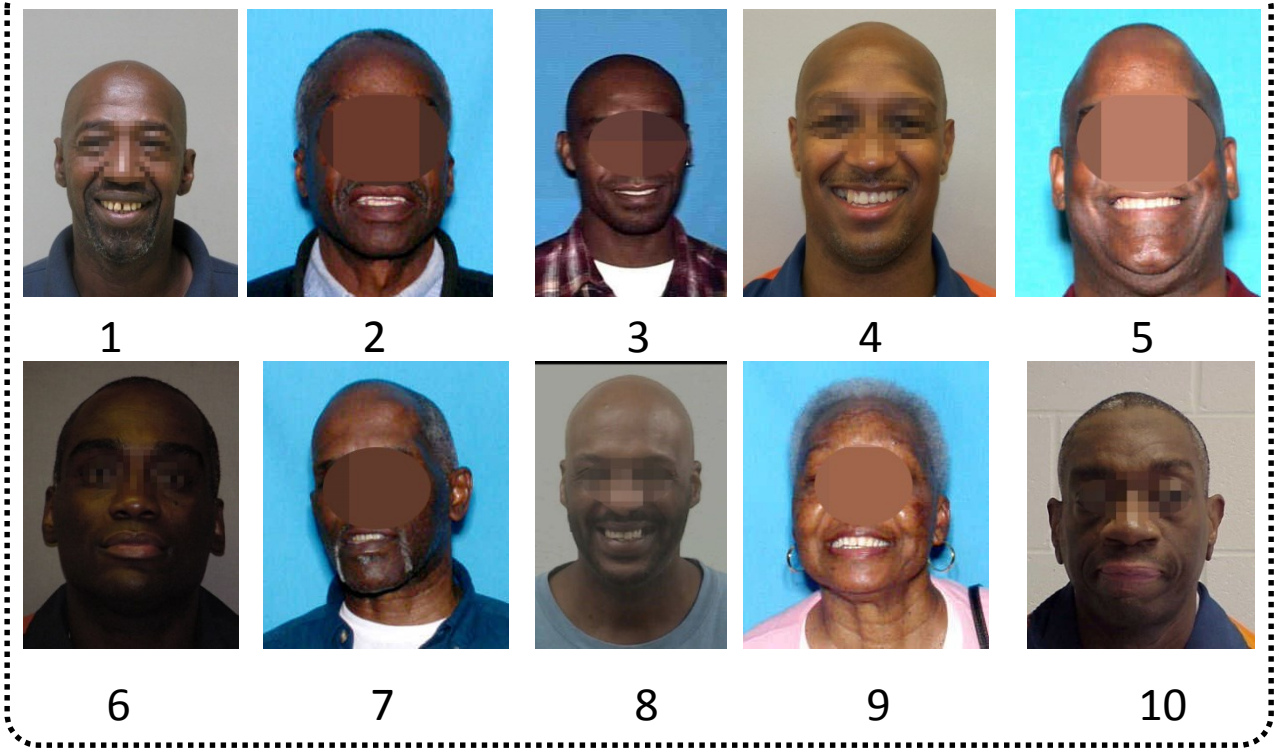

Gallery: 34 million (30M DMV photos, 4M mugshots)

Mobile Phones

Joseph Van Os / Getty Images

More cell phone accounts than world's population; $1 Trillion in mobile payments

Mobile Face Unlock

Uploaded: Dec 6, 2011 YouTube

Photo Tagging


Townhall
Katie Pavlich
La Opinión
Maria Peña
New York Times
Jonathan Mahler
Fox Business
Blake Burman
Raghubir Goyal
Newsmax
John Gizzi
Weekly Standard
Michael Warren
Daily Signal
Fred Lucas
Fox News
Howard Kurtz
BBC
Tara McKelvey
Yahoo News
Hunter Walker
Talk Media News
Jon-Christopher Bua
Ami Magazine
Jake Turx
Al Jazeera
Kimberly Halkett
New York Post
Daniel Halper
Breitbart
Charlie Spiering
Washington Examiner
Sarah Westwood
National Journal
George Condon
ABC News
Cecilia Vega
Politico
Annie Karni
Fox Radio
Jon Decker
One America News Network
Trey Yingst
UNF News
Kyle Mazza
Time
Zeke Miller
USA Today
David Jackson
Wall Street Journal
Carol Lee
Associated Press
Mark Smith
Washington Times
Dave Boyer
NPR
Mara Liasson
Bloomberg
Justin Sink
Washington Post
John Wagner
ABC News
Jonathan Karl
CBS News
Margaret Brennan
Fox News
John Roberts
NBC News
Hallie Jackson
Reuters
Emily Stephenson
Associated Press
Catherine Lucey

Social Media

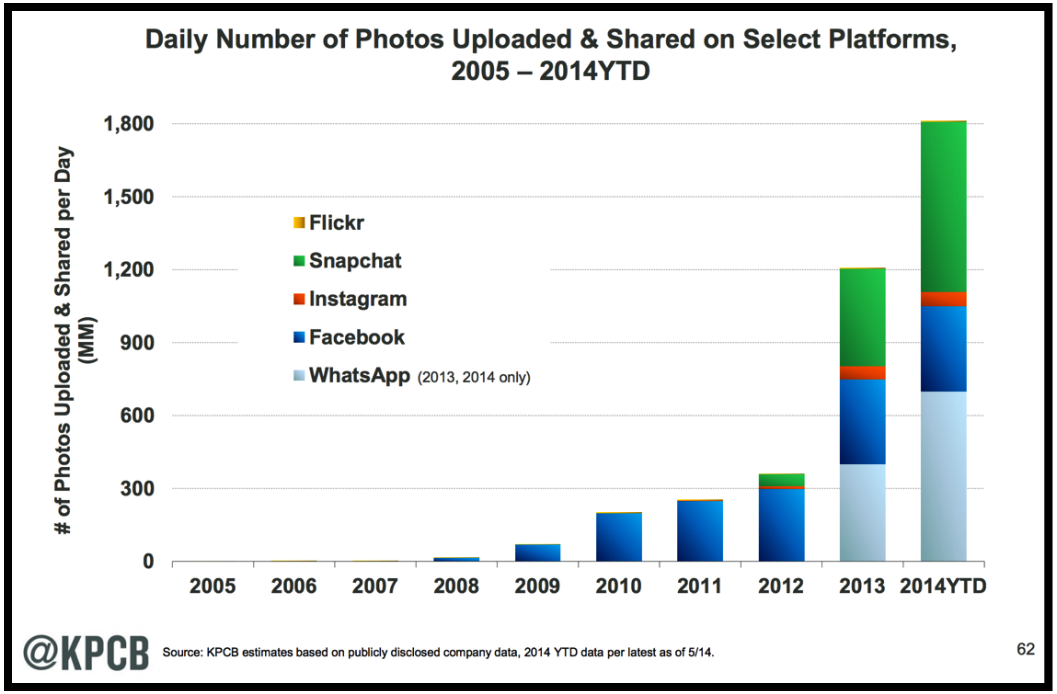

- ***~trillion image*** shares per year, and increasing
- Challenges: accuracy and efficiency

Constrained Face Recognition

- Cooperative subjects: Small intra-subject variations (**FERET**)
- Operational face data (**mugshots, visa images**)
  - Limited user cooperation
  - Effect of aging

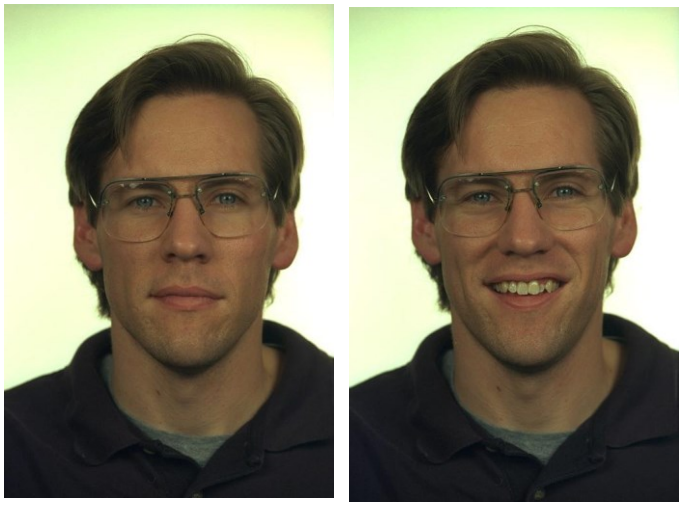

FERET Images

PCSO Mugshot Images

Unconstrained Face Recognition

- **LFW**
  - Images of celebrities and public figures
  - Faces detect by Viola-Jones detector
- **IJB-A**
  - Semi-automatic data collection
  - Manually selected identities & annotation


LFW Images


IJB-A Images

State of the Art: Verification

FRGC v2.0 (2006) MBGC (2010) LFW (2007) IJB-A (2015)

LFW Standard Protocol
99.77% (Accuracy)
3,000 genuine & 3,000 imposter pairs;
10-fold CV

LFW BLUFR Protocol
88% TAR @ 0.1% FAR
156,915 genuine, ~46M imposter pairs;
10-fold CV

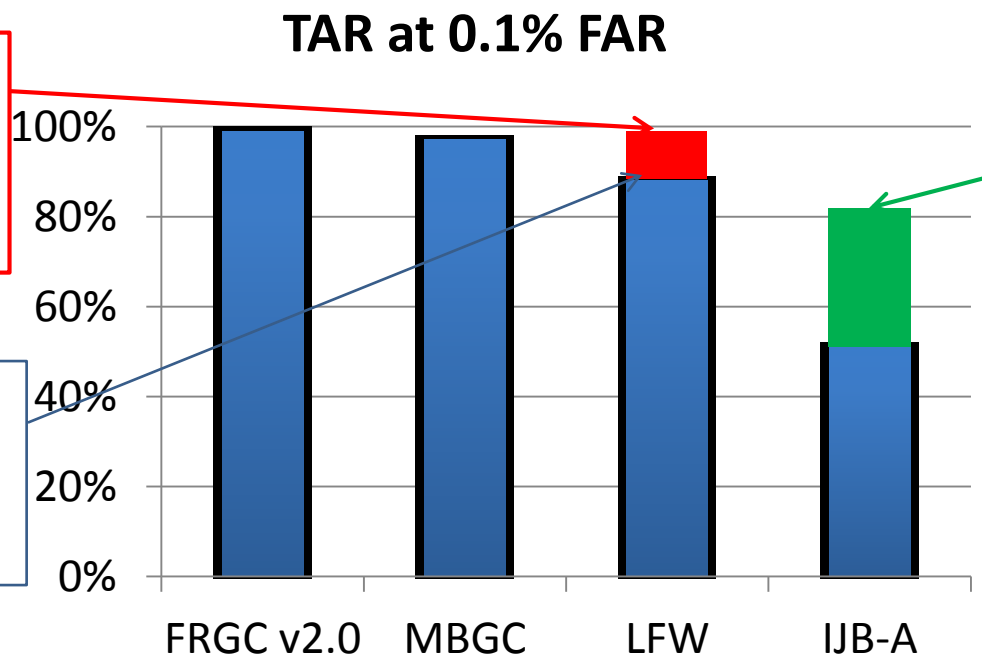

IJB-A
80%TAR @ 1%FAR
10-fold CV

State of the Art: Verification

NIST IJB-B database: TAR @0.01% FAR = 70%

Automated Face Recognition

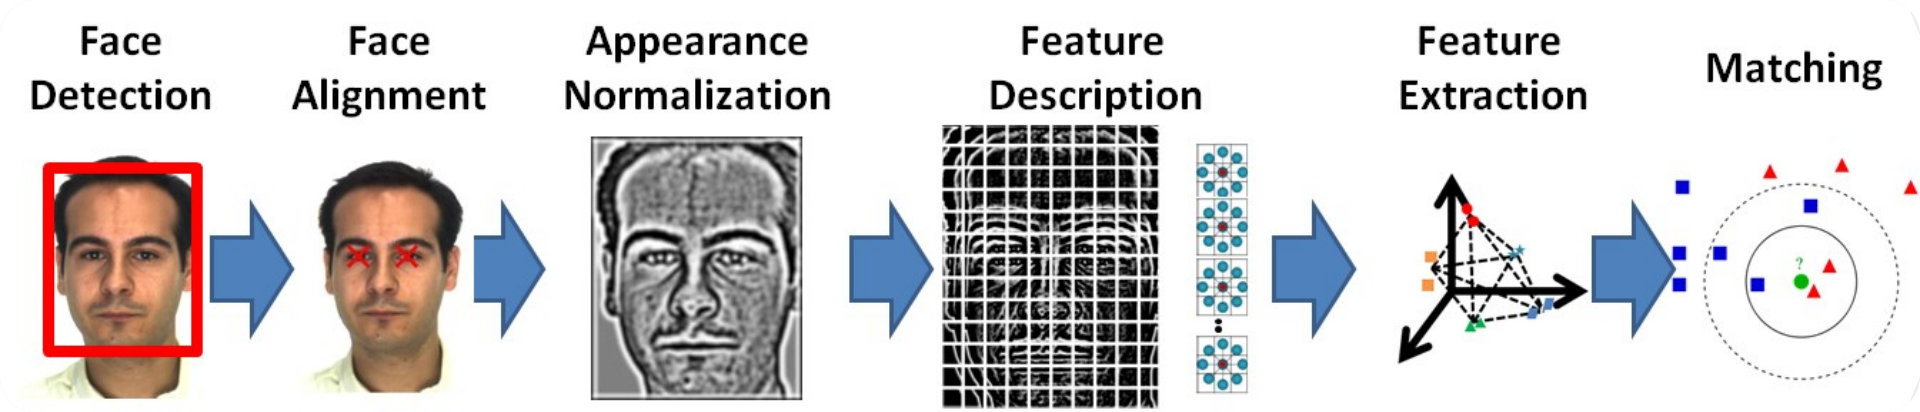

- Most face recognition algorithms follow this pipeline

Face Detection

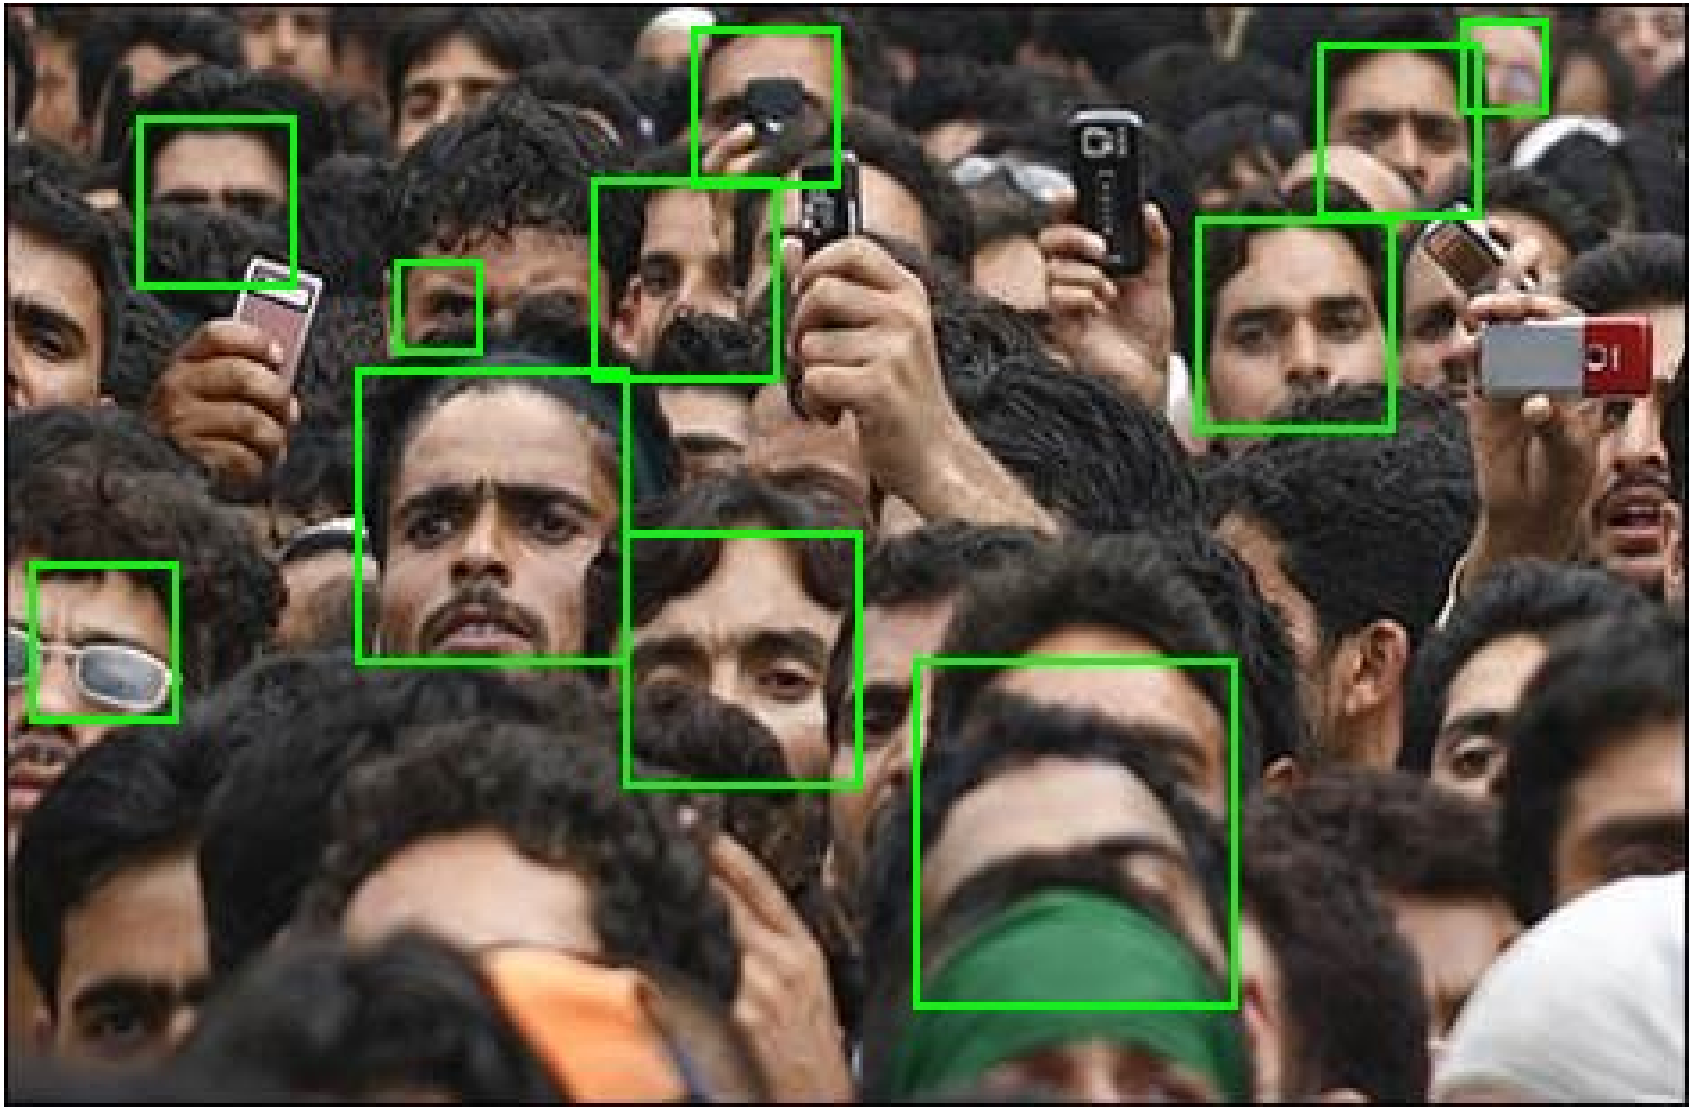

Number of mobile phone users worldwide in 2016 is estimated to be about 4.8 billion

Learning Face Representation

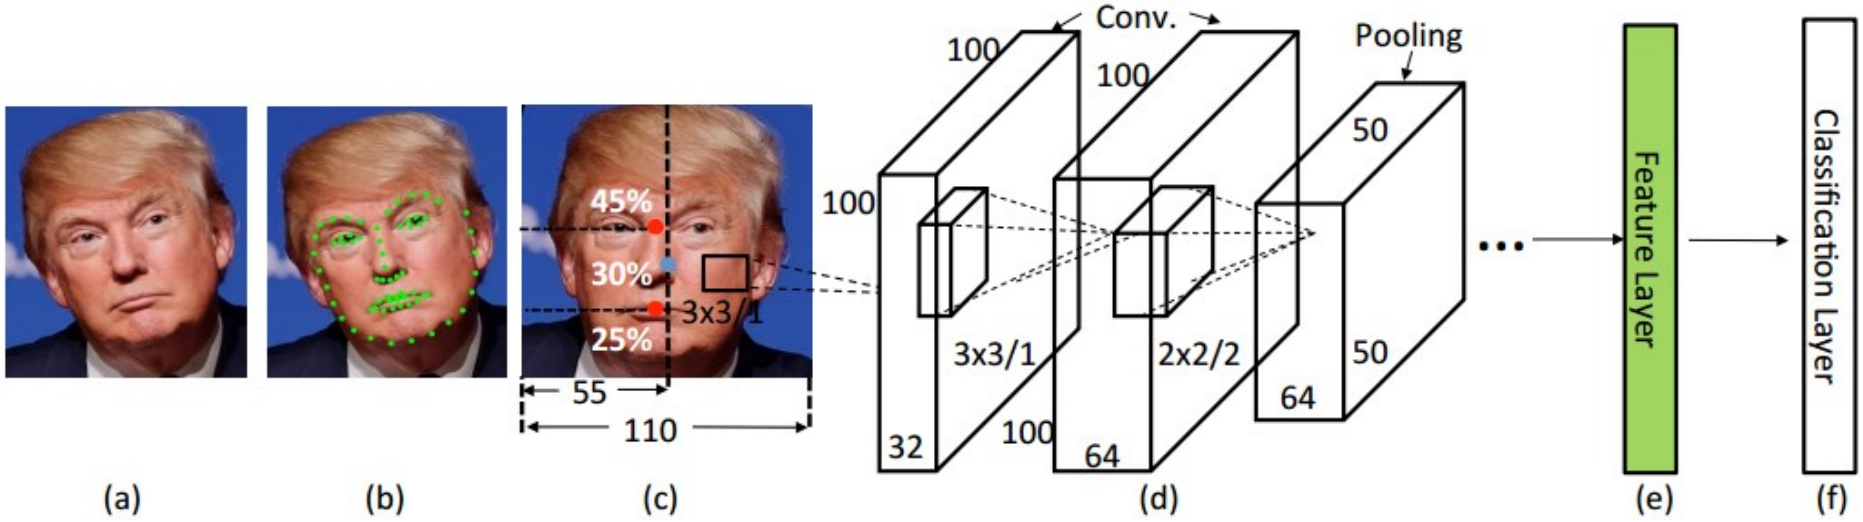

(a) Input RGB image, (b) detected keypoints, (c) normalized face image, (d) a convolutional neural network, (e) 320-dimensional feature vector and (f) softmax classification layer for training only

Network Training

- ConvNet is trained with CASIA-Webface
  - 494,414 images of 10,575 subjects (**training bias?)**
- Preprocessing: face and landmark detection
  - Align face images using the centers of eyes & mouth

#subjects = 10,575

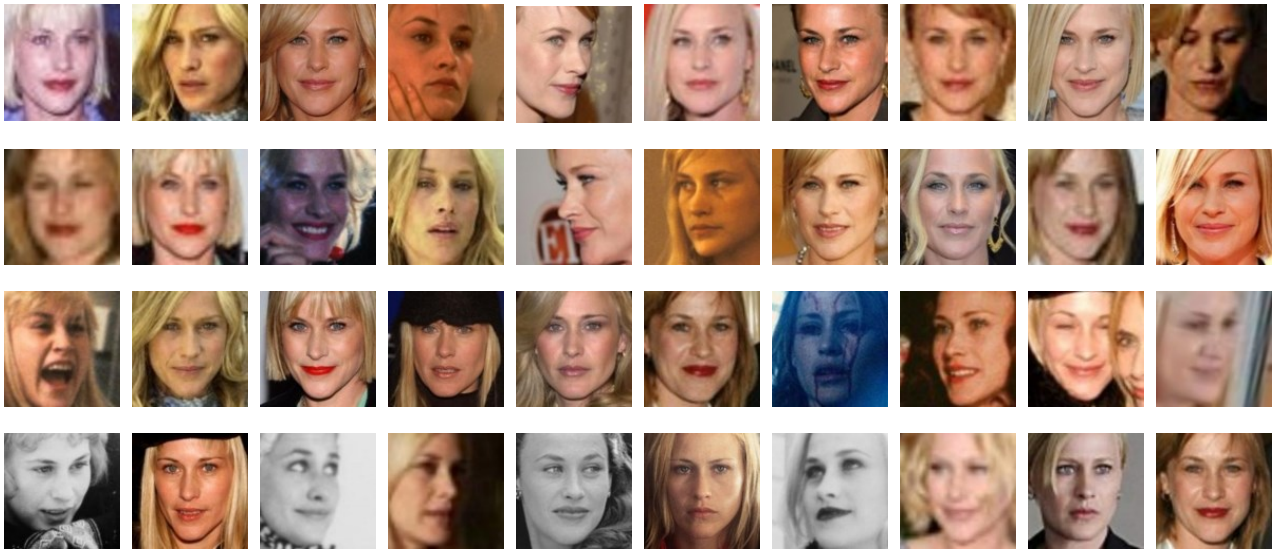

Total # images with landmarks = 435,689

Face v. Non-Face

- In 120M faces, we estimate ~2.3% non-faces
- Some “non-faces” are of statues, toys, etc., but some are completely wrong

False-Positive Face Detections

Search Example

- 3 Images of TV Anchor Tammy Leitner, added to 120M background set
- Top-10 retrieval results for one query:

Search for Sherry Jones (120M Gallery)





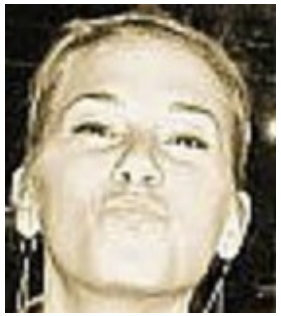








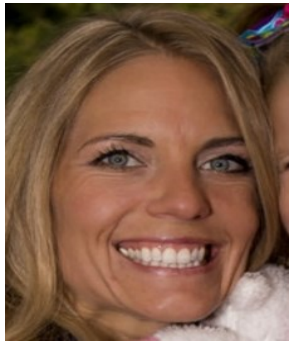

Search for Dzhokhar Tsarnaev (120M Gallery)

Network Architecture & Training

	Training Set	Network	VR@FAR=0.1%
(1)	CASIA-Webface	From Scratch [50]	84.41%
(2)	CASIA-Webface	From Scratch [50], fusion of 9 models	88.00%
(3)	CASIA-Webface	18-layer ResNet	82.06%
(4)	VGG-CASIA	50-layer ResNet	88.67%
(5)	VGG-CASIA	50-layer ResNet, 10-crop	89.74%
(6)	CASIA-Webface	50-Layer Pre-ResNet	88.36%
(7)	CASIA-Webface	50-Layer Pre-ResNet, 10-crop	89.64%
(8)	VGG-Face	50-Layer Resnet	81.40%
(9)	VGG-Deduplicated	50-Layer Pre-ResNet	86.98%
(10)	VGG-Deduplicated+CASIA-Webface	50-Layer Pre-ResNet	91.04%
(11)	VGG-Deduplicated+CASIA-Webface	50-Layer Pre-ResNet, 10-crop	92.22%
(12)	VGG-Deduplicated+CASIA-Webface	101-Layer Pre-ResNet	91.18%
(13)	VGG-Deduplicated+CASIA-Webface	101-Layer Pre-ResNet, 10-crop	92.10%

LFW Under BLUFR Protocol

IARPA Janus Program

IARPA's Janus program aims to dramatically improve the current performance of face recognition tools by fusing the rich spatial, temporal, and contextual information available from the multiple views captured by today's "media in the wild".

https://www.iarpa.gov/index.php/research-programs/janus

Some Challenges

Pose, Illumination, Expression

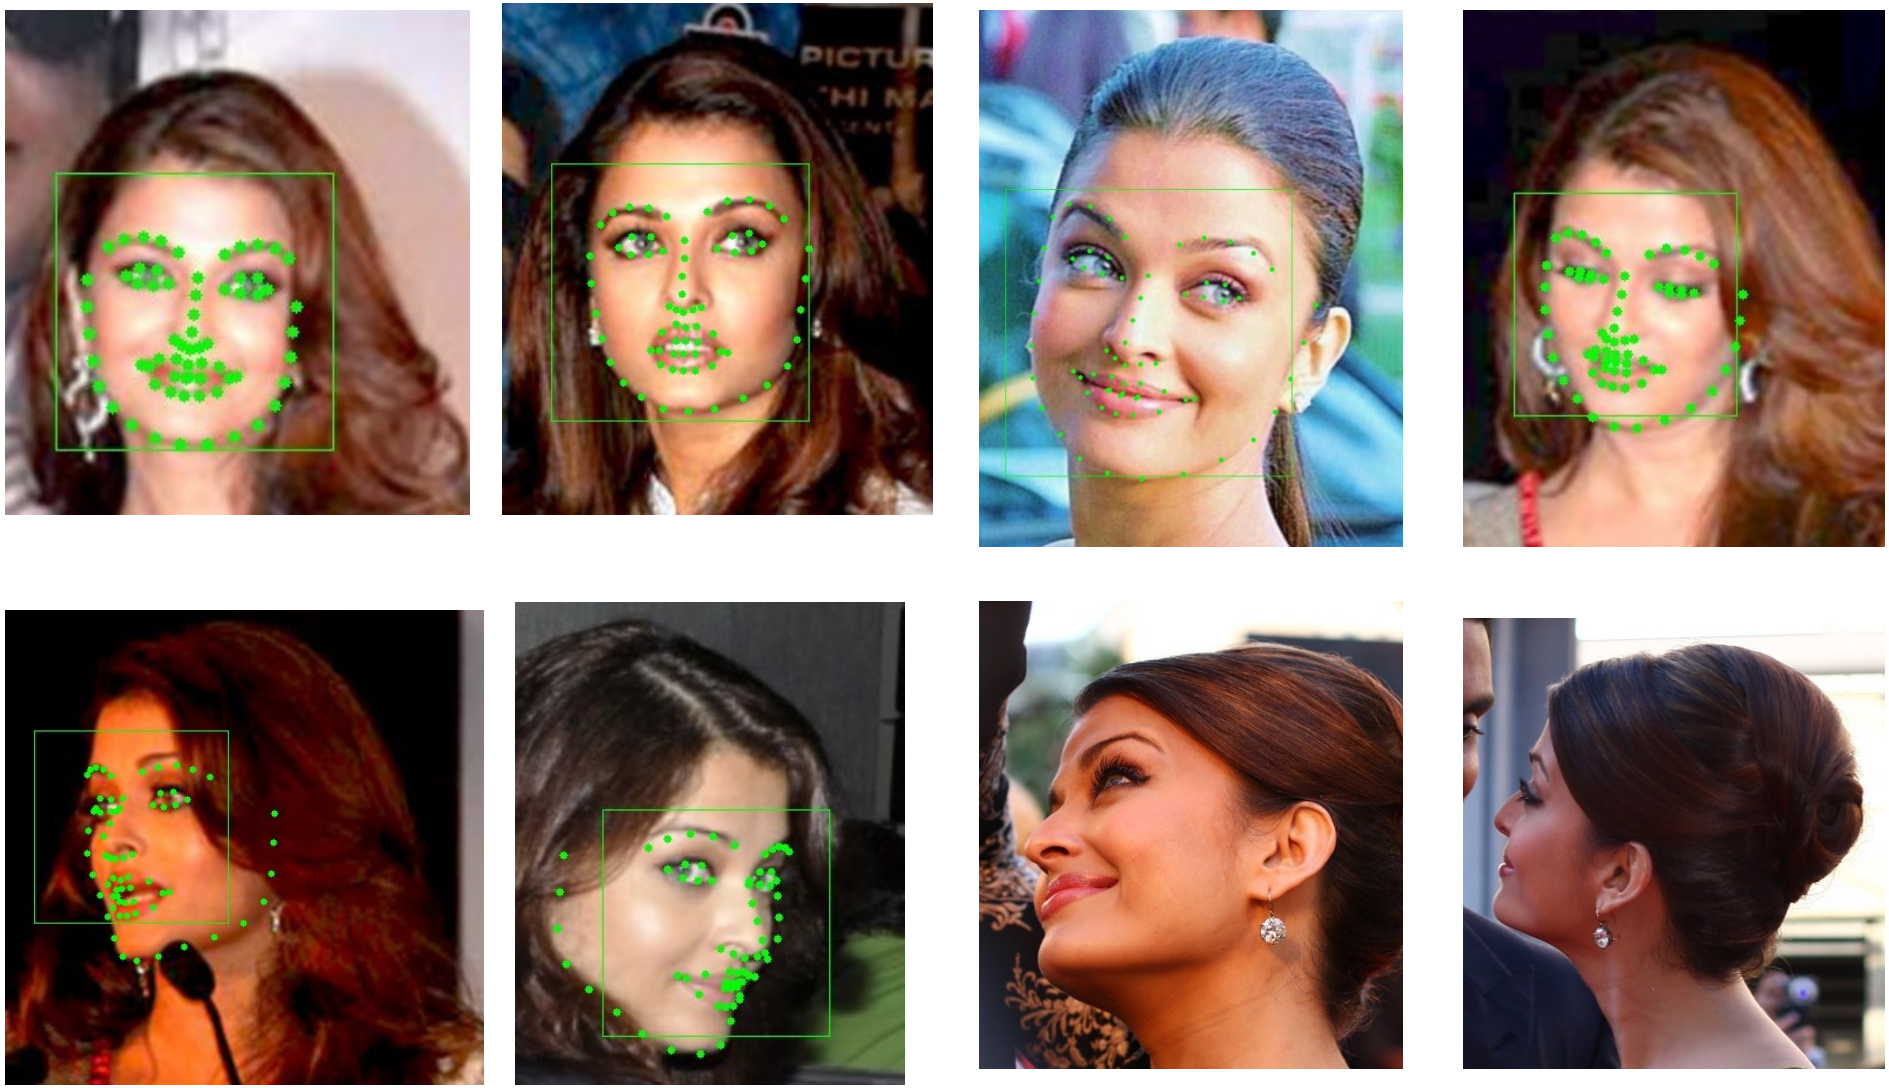

Images of one subject in NIST IJB-A data, overlaid with V-J detector & dlib landmarks

Facial Aging and Doppelgangers

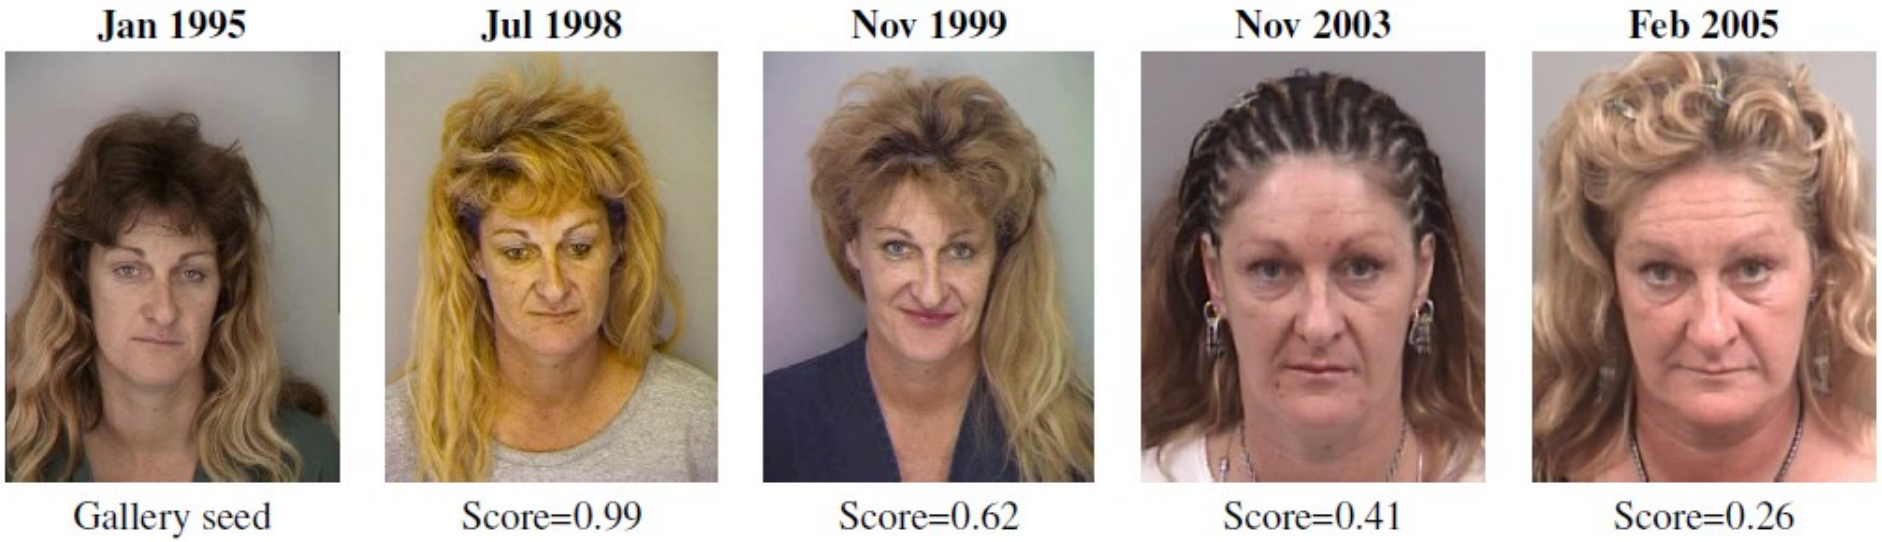

http://www.theguardian.com/theguardian/2010/dec/05/barack-obama-doppelganger-ilham-anas

Scars, Marks & Tattoos

Detroit police linked at least six armed robberies at an ATM after matching a tipster's description of the suspect's distinctive tattoos

www.DetroitisScrap.com/2009/09/567/

Which Ones Are Real?

Which Ones Are Real?

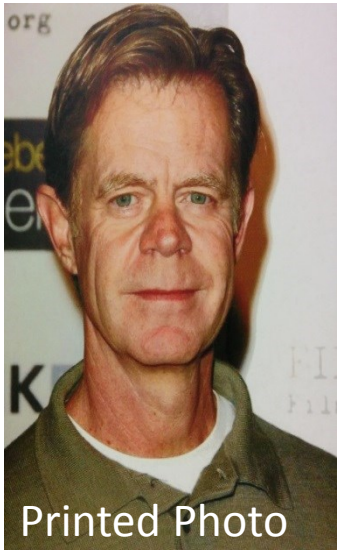

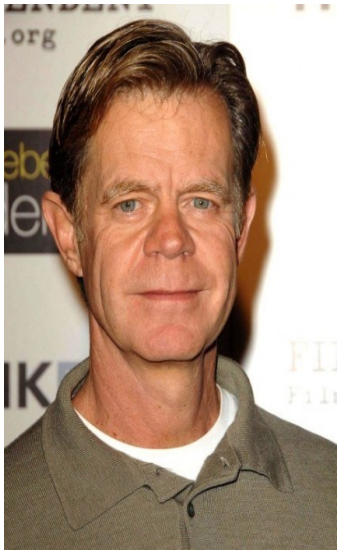

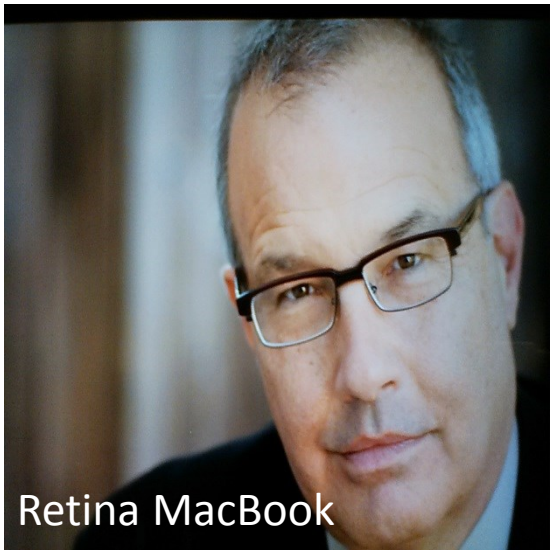

Face Image Recovery from Templates

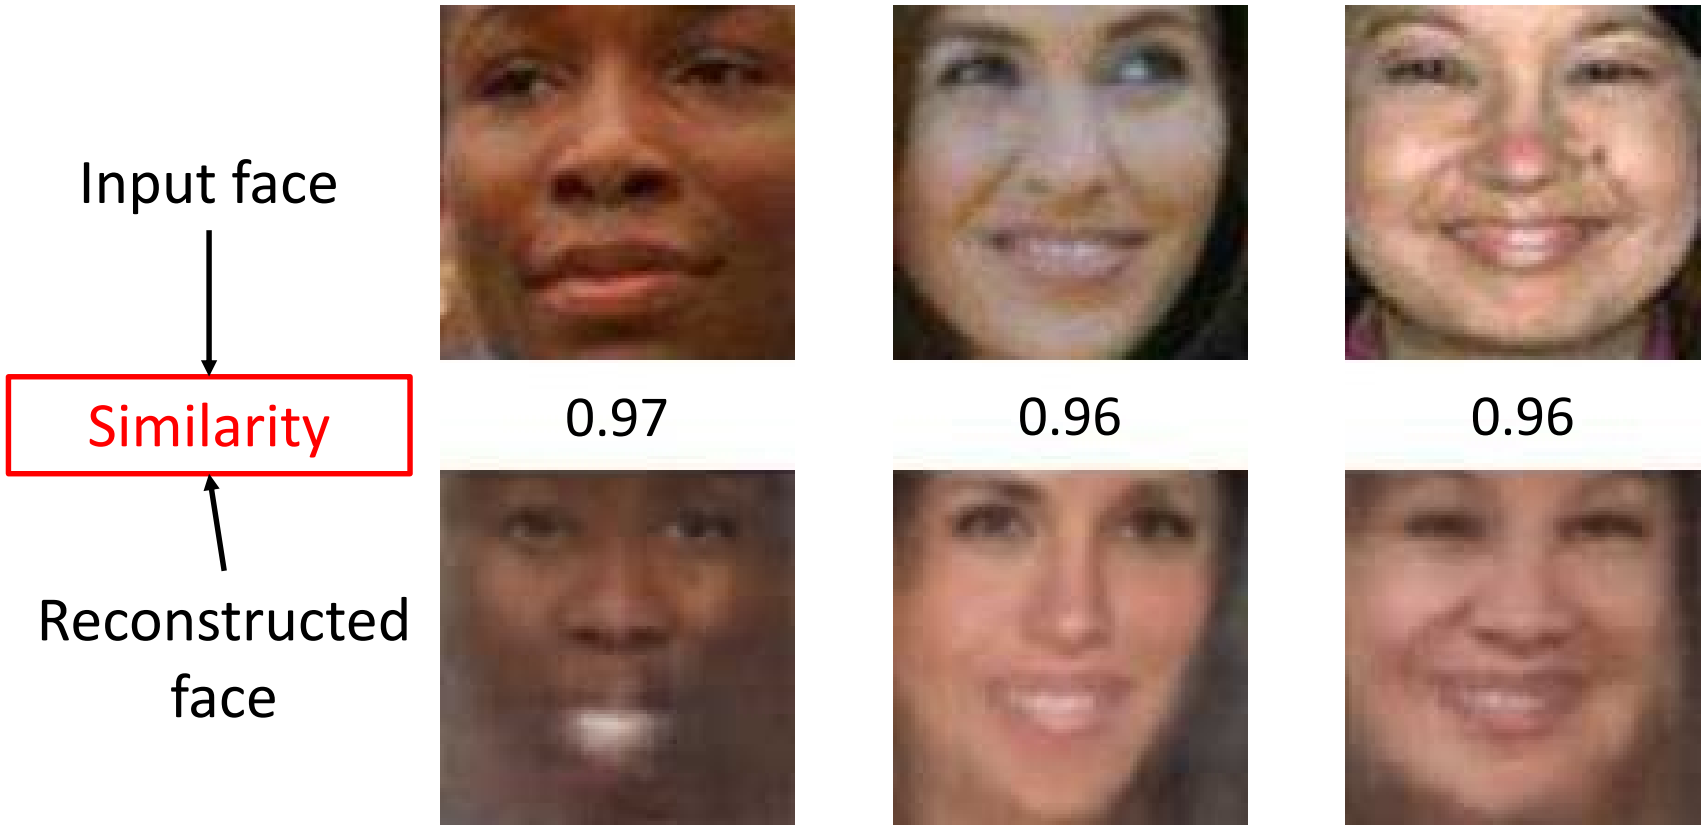

Threshold@ FAR = 0.1% is 0.78

Capacity of Face Recognition?

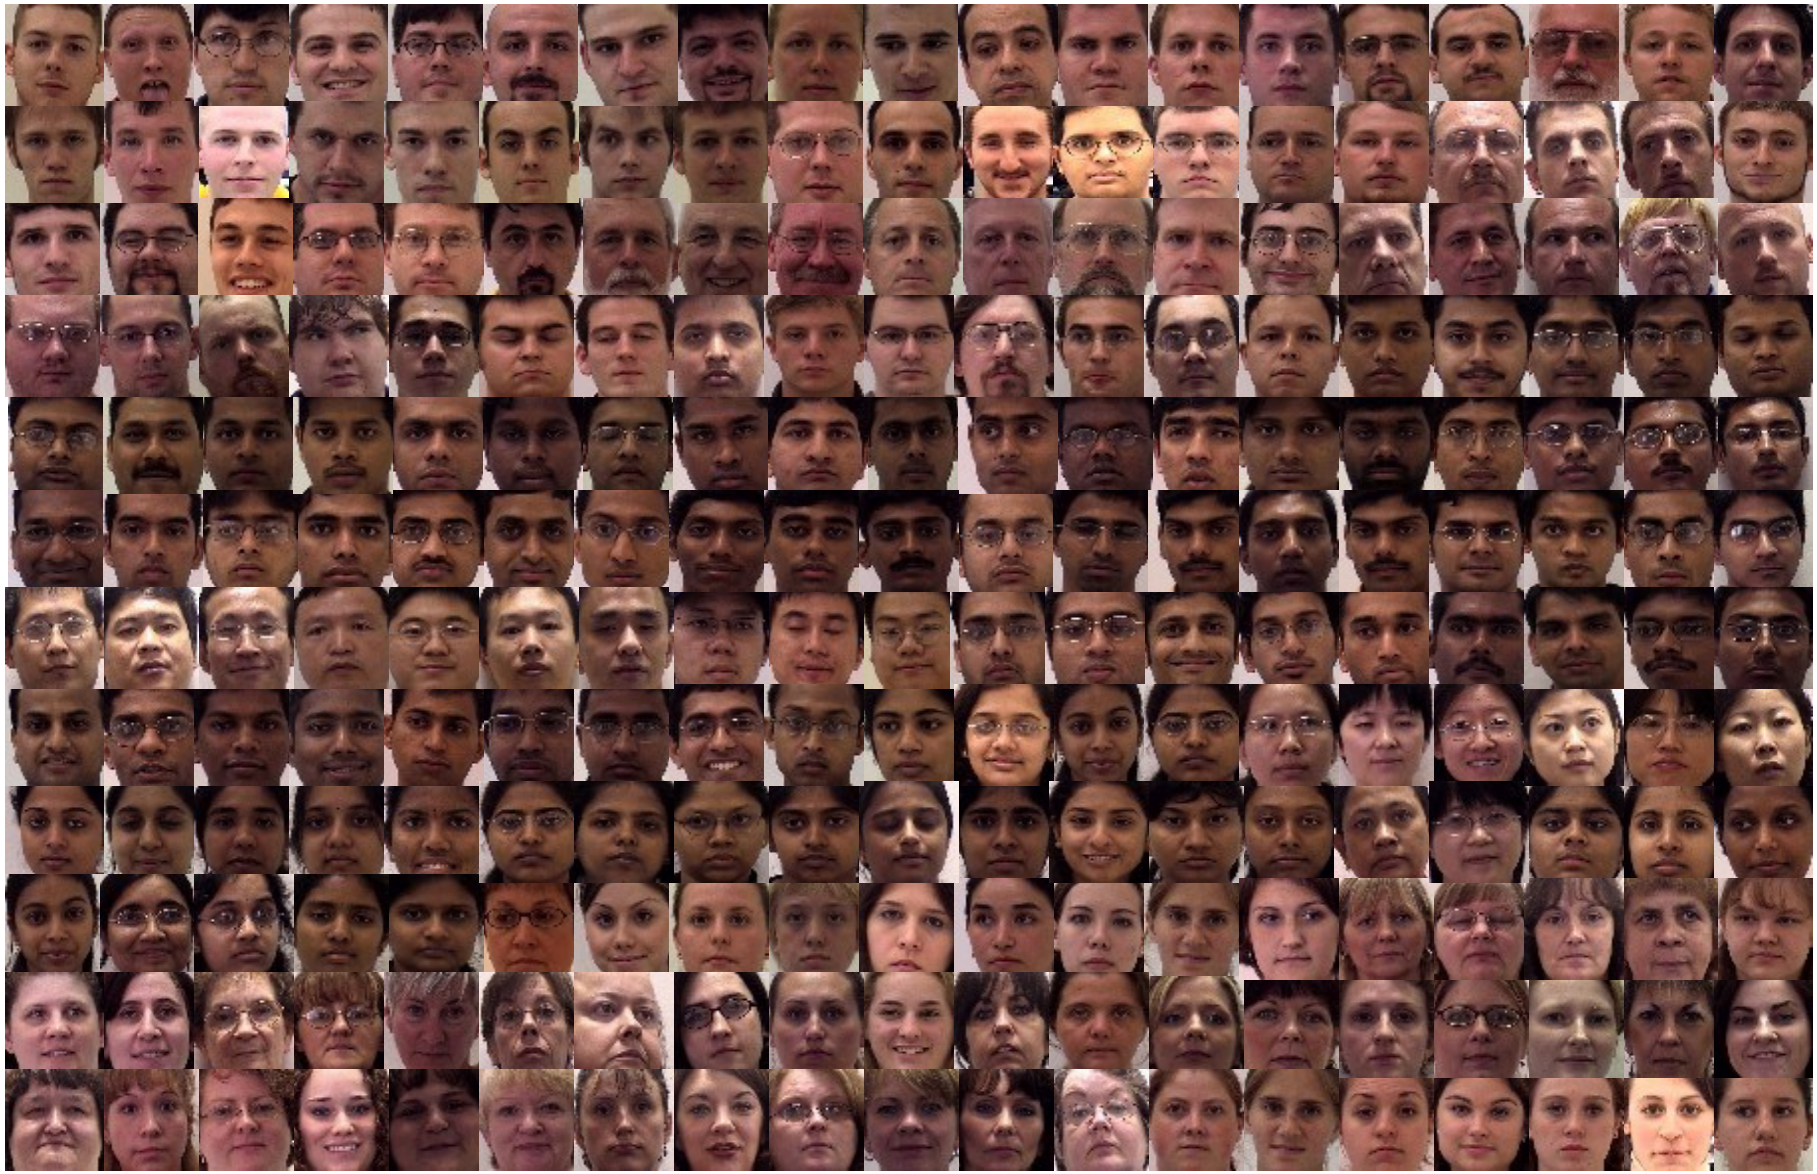

Summary

- Face recognition is now a major topic of research; growing no. of FR systems deployed
- State-of-the-art: High accuracy for constrained & cooperative subjects; low accuracy for unconstrained face recognition of non-cooperative subjects
- Need recognition systems robust to pose, illumination & expression, aging, and low resolution video
- User concerns: Data security & privacy

References

- L. Best-Rowden and A. K. Jain, "Automatic Face Image Quality Prediction", arXiv preprint arXiv:1706.09887, 2017
- G. Mai, K. Cao, P.C. Yuen and A.K. Jain, "Face Image Reconstruction from Deep Templates", arXiv preprint arXiv:1703.00832, 2017
- L. Best-Rowden and A.K. Jain, "Longitudinal Study of Automatic Face Recognition", *IEEE Trans. Pattern Analysis & Machine Intelligence*, 2017 DOI:10.1109/TPAMI.2017.2652466
- C. Otto, D. Wang and A. K. Jain, "Clustering Millions of Faces by Identity", *IEEE Trans. Pattern Analysis & Machine Intelligence*, 2017 https://arxiv.org/abs/1604.00989
- D. Wang, C. Otto and A. K. Jain, "Face Search at Scale", *IEEE Transactions on Pattern Analysis and Machine Intelligence*, DOI 10.1109/TPAMI.2016.2582166, June 2016