Partial Face Recognition

Shengcai Liao

NLPR, CASIA

April 29, 2015

Background

- Cooperated face recognition
  - People are asked to stand in front of a camera with good illumination conditions
    - Border pass, access control, attendance, etc.
  - Mostly solved

Background

- Unconstrained face recognition
  - Images are captured arbitrarily without or with little user cooperation
    - Video surveillance, hand held system, etc.
  - Difficult task

Background

- Partial face recognition in unconstrained environments

Background

- Partial faces in unconstrained environments

Scenario	External occlusion	Self occlusion	Facial accessories	Limited field of view (FOV)	Extreme illumination	Sensor saturation
Examples	occlusion by other objects	non-frontal pose	hat, sunglasses, scarf, mask	partially out of camera's FOV	gloomy or highlighted facial area	underexposure or overexposure
Image						

Face Recognition and the London Riots
Summer 2011

Widespread looting and rioting:

Extensive CCTV Camera Network:

FR leads to many arrests:

Yet, many suspects still unable to be identified by COTS FRS:

Partial Face Recognition (PFR)

- Problem
  - Recognize an arbitrary partial face image captured in uncontrolled environment
- Importance
  - Recognize a suspect in crowd
  - Identify a face from its partial image
- Difference from traditional face recognition
  - Alignment?
  - Feature representation?
  - Classification?

Alignment Free Partial Face Recognition: Application Scenario

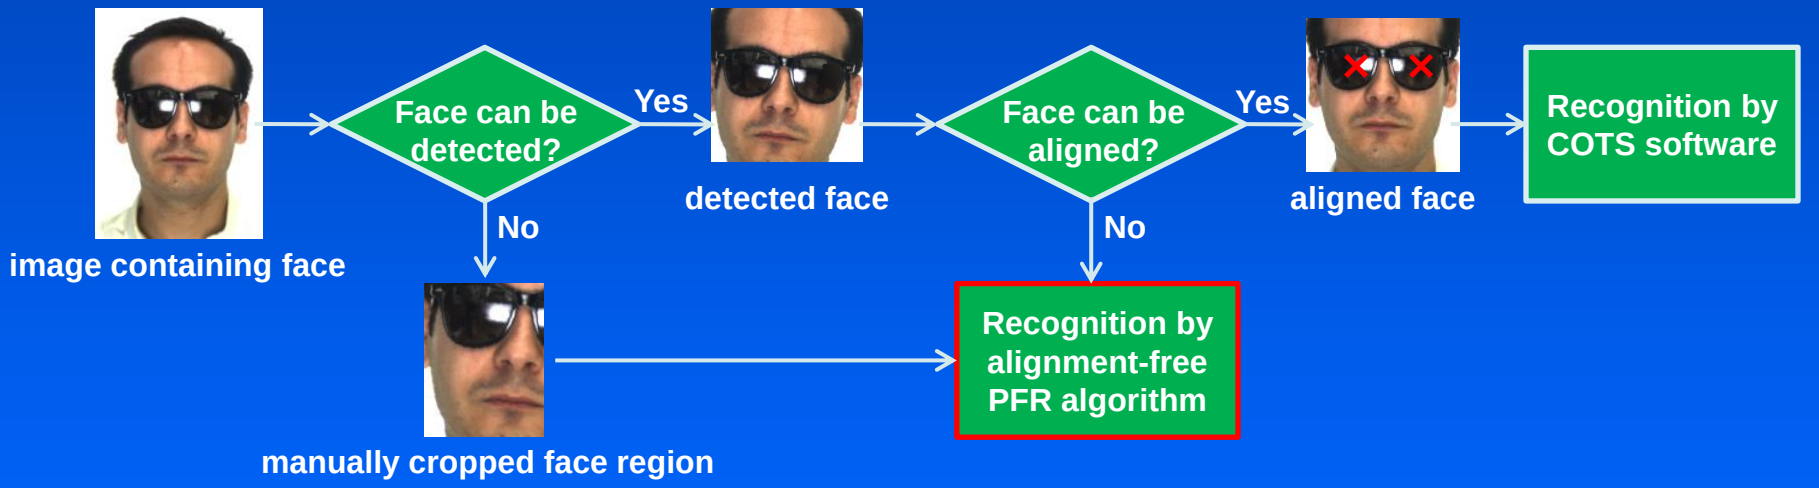

Alignment Free Partial Face Recognition: Overview

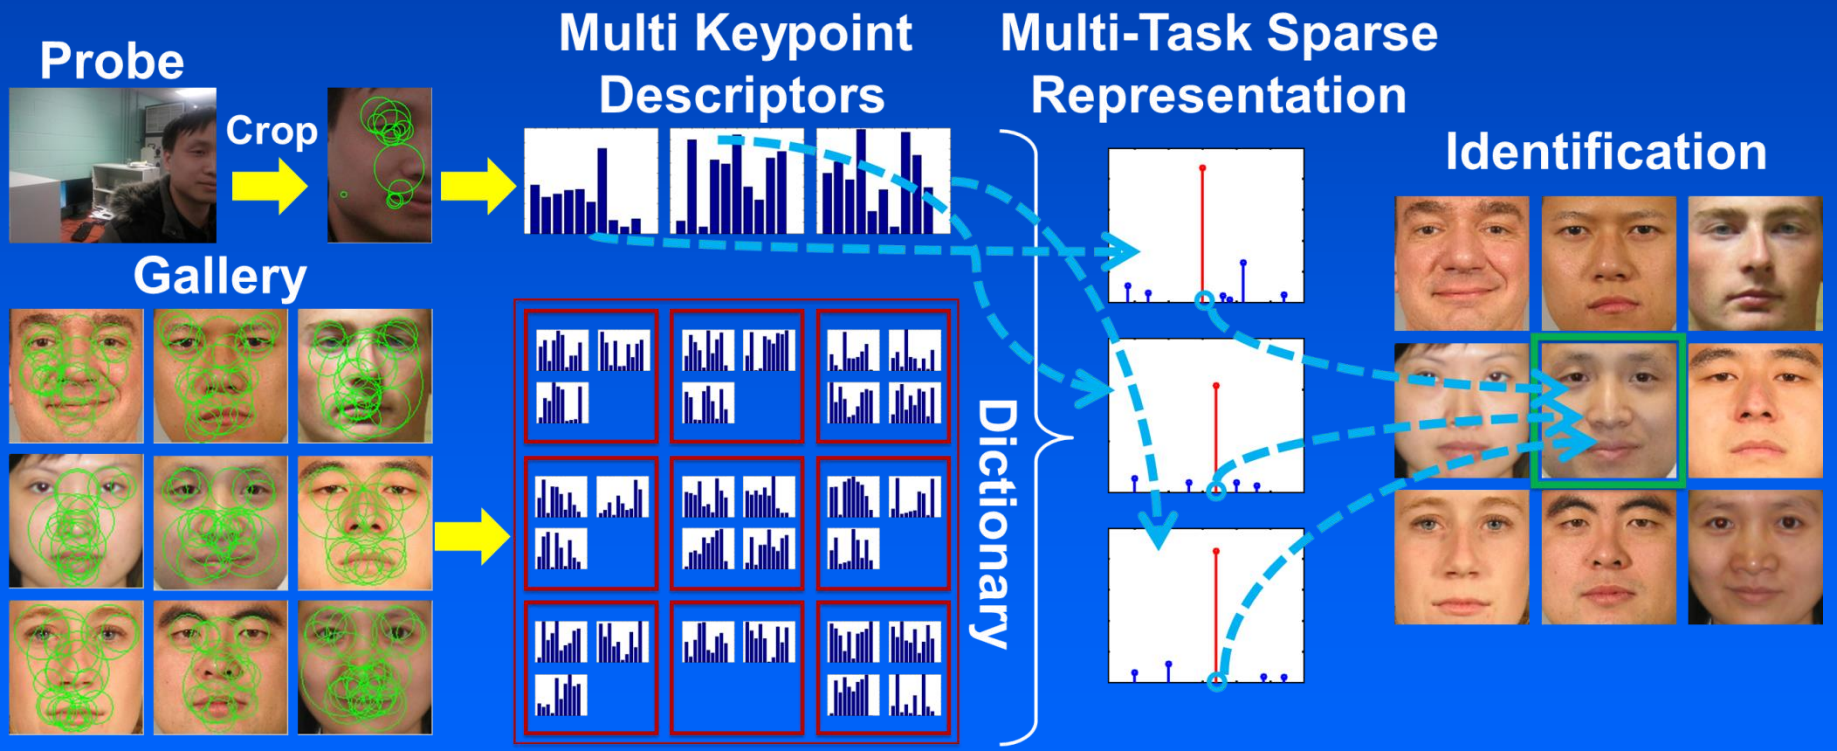

Interest Point Based Local Descriptor

keypoint detection

interest region description

matching

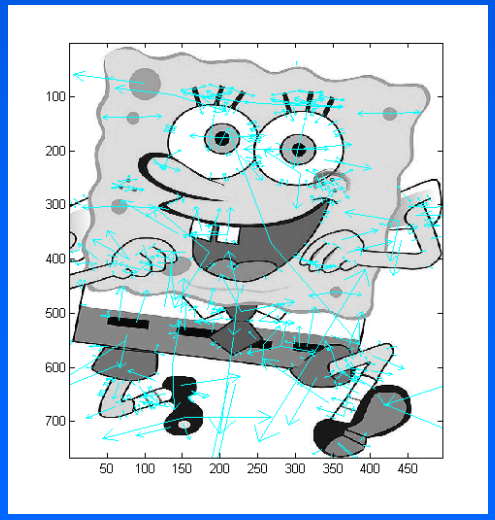

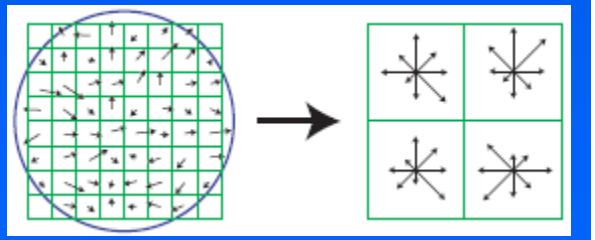

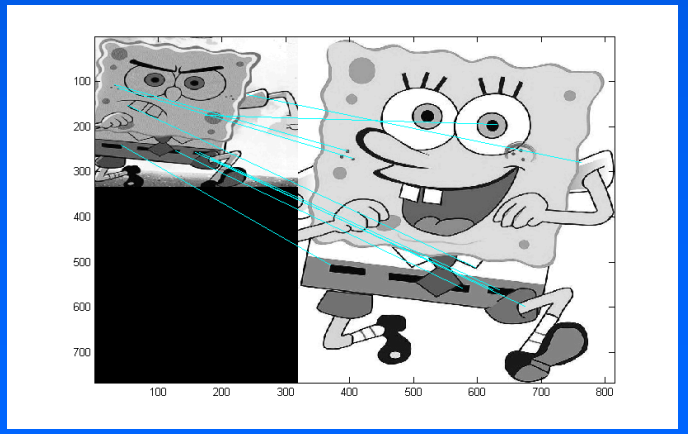

Interest Point Based Local Descriptor

- Image retrieval
- Image matching
- Object recognition
- Texture recognition
- Robot localization
- …

Interest Point Based Local Descriptor

- Intensity histogram
- SIFT
- HOG
- GLOH
- PCA-SIFT
- SURF

Face Description with Interest Points

- SIFT detector
  - Detects blobs
  - Limited keypoints
- CanAff detector
  - Canny edge based
  - Plenty keypoints

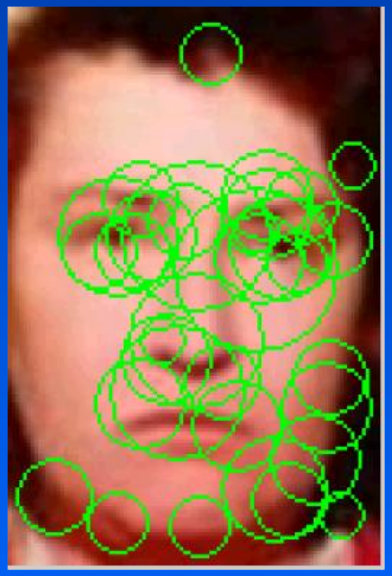

SIFT
(37 keypoints)

CanAff
(571 keypoints)

K. Mikolajczyk, A. Zisserman, and C. Schmid, “Shape recognition with edge-based features,” in *Proceedings of the British Machine Vision Conference*, 2003.

Face Description with Interest Points

- Advantages of interest point detectors
  - Detections of local structures, not pre-defined components
    - Good for partial faces
  - Affine invariance
    - Good for pose/viewpoint changes
  - High repeatability
    - Good for partial face matching

Face Description with Interest Points

- Gabor Ternary Pattern (GTP) based descriptor

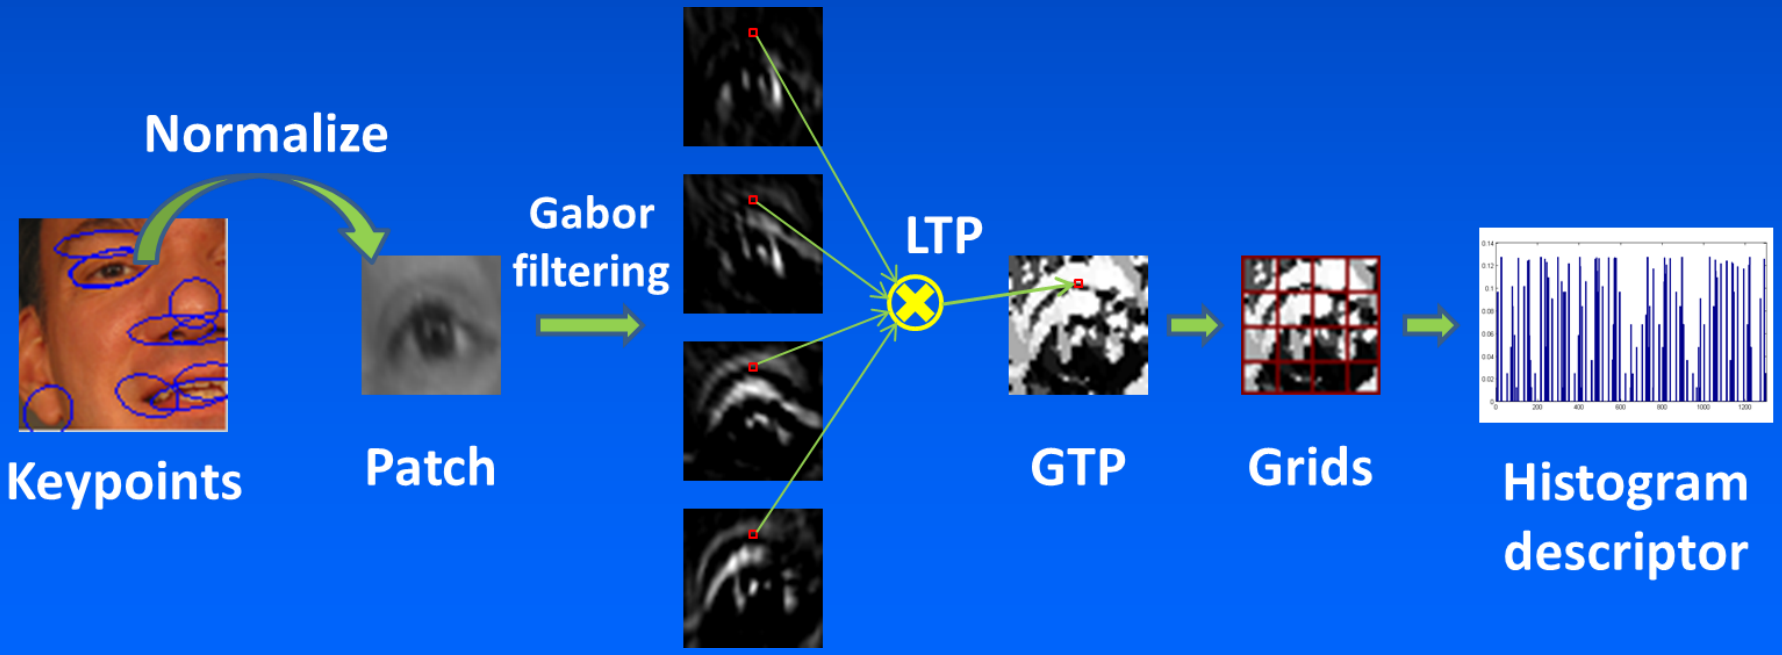

Multi Keypoint Descriptors (MKD)

- Each image is described by a set of keypoints and descriptors:
  - Keypoints: $p_1, p_2, \ldots, p_k$
  - Descriptors: $d_1, d_2, \ldots, d_k$
- The number of descriptors, k, may be different from image to image

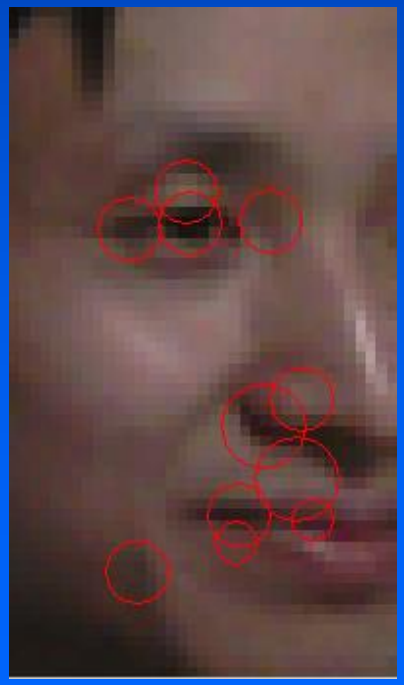

MKD based Sparse Representation Classification (MKD-SRC)

- Descriptors of the same class c can be viewed as a sub-dictionary: $\mathbf{D}_c = (\mathbf{d}_{c_1}, \mathbf{d}_{c_2}, \cdots, \mathbf{d}_{c_{k_c}})$
- A gallery dictionary is built: $\mathbf{D} = (\mathbf{D}_1, \mathbf{D}_2, \cdots, \mathbf{D}_C)$
- For each descriptor y_i in a test sample $\mathbf{Y} = (y_1, y_2, \cdots, y_k)$, solve

$$\hat{\mathbf{x}}_i = arg\min \|\mathbf{x}_i\|_1, \ s.t. \ \mathbf{y}_i = \mathbf{D}\mathbf{x}_i$$

- Determine the identity of the test sample by SRC:

$$\min_c r_c(\mathbf{Y}) = \frac{1}{k}\sum_{i=1}^{k} \|\mathbf{y}_i - \mathbf{D}_c\delta_c(\hat{\mathbf{x}}_i)\|_2^2$$

MKD based Sparse Representation Classification (MKD-SRC)

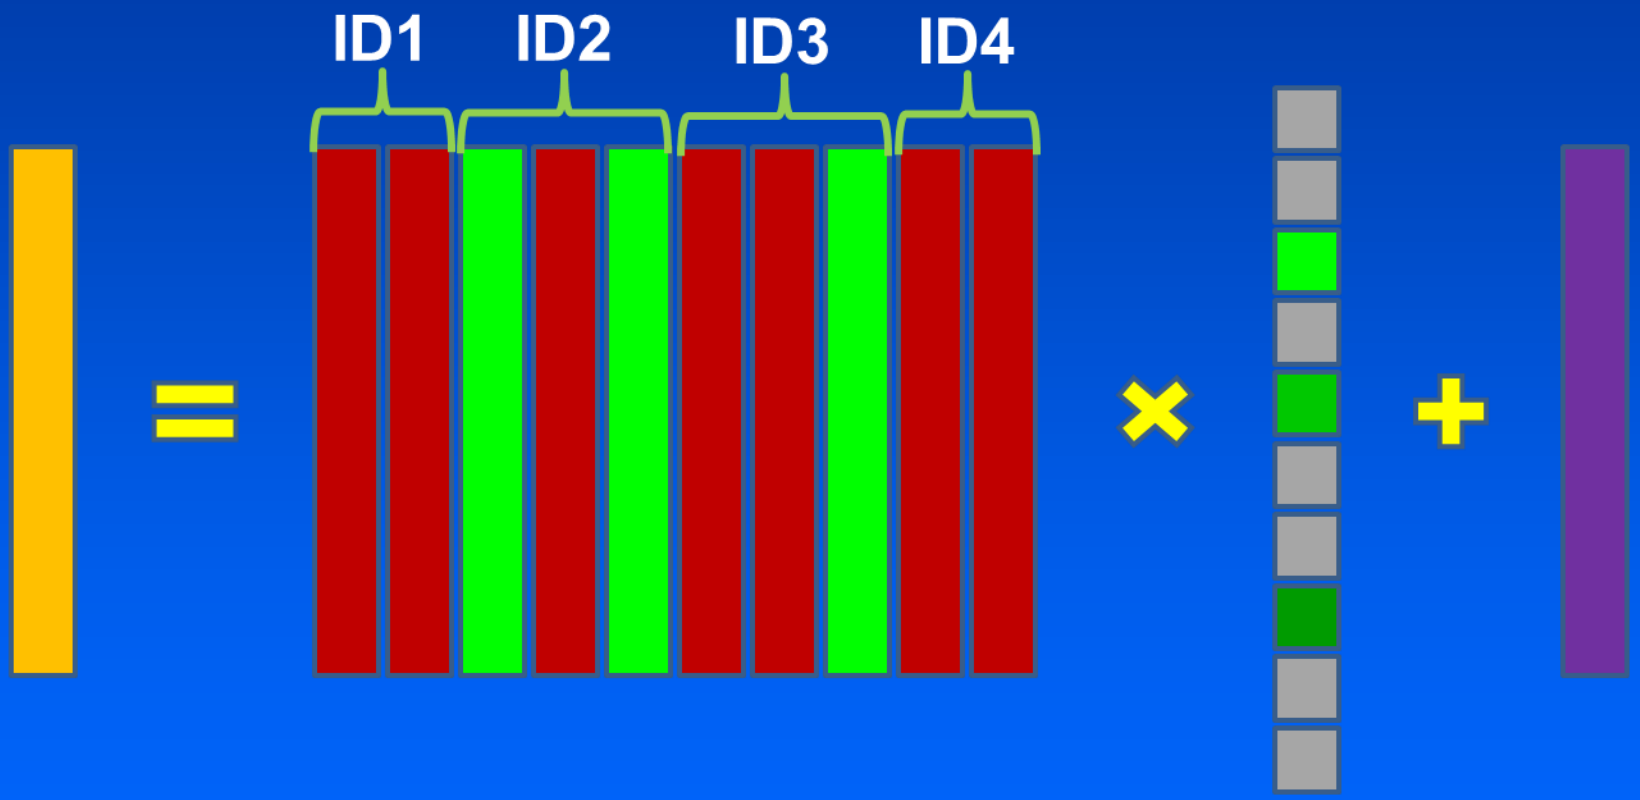

$$y_{M\times 1} = D_{M\times K} * x_{K\times 1} + \varepsilon_{M\times 1}$$

An Example of MKD-SRC Solution

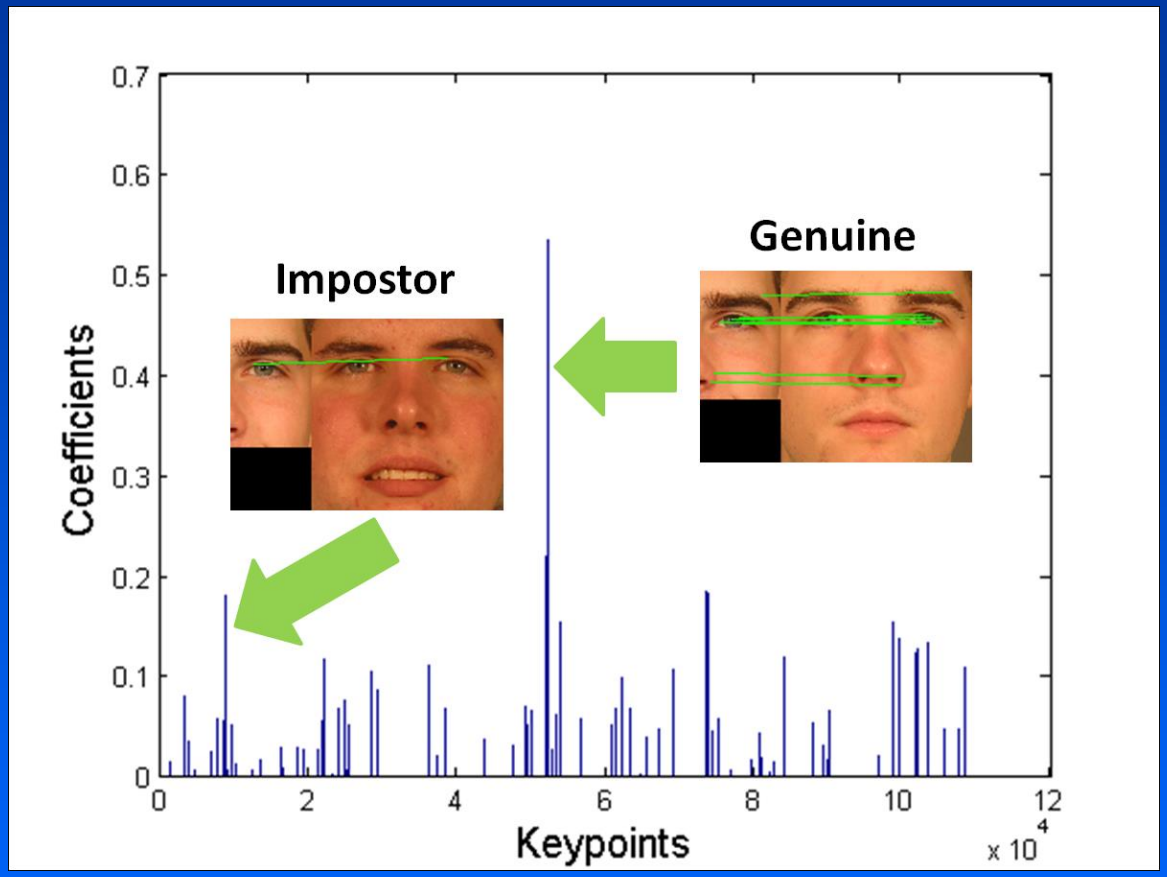

Genuine

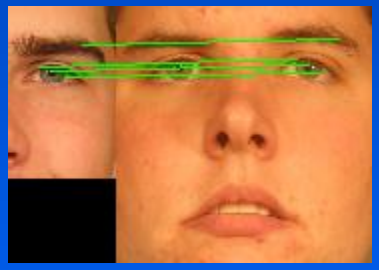

Impostor

Lowe's SIFT

- MKD-SRC is more discriminant in recognizing partial faces

Fast Atom Filtering

- In the dictionary, the number of atoms, K, can be of the order of millions
- Fast atom filtering

$$\mathbf{c}_i = \mathbf{D}^T \mathbf{y}_i,\ i = 1, 2, \cdots, k.$$

  For each y_i, filter out T (T<<K) atoms, i.e. T largest values of c_i, resulting in a small sub-dictionary
- The filtering scales linearly w.r.t. K, while the remaining MKD-SRC task takes a constant time

Effects of the Fast Atom Filtering

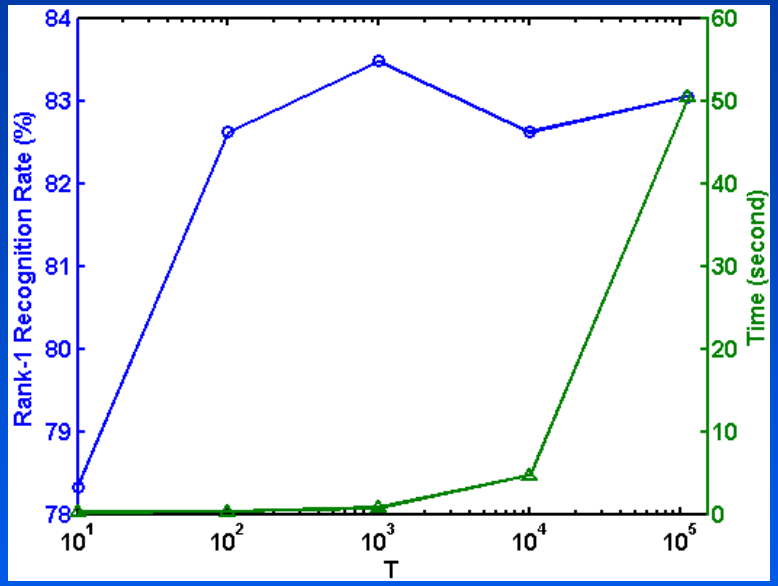

T	Accuracy	Time
10^1	78.33%	0.2s
10^2	82.62%	0.3s
10^3	83.48%	0.8s
10^4	82.62%	4.7s
$= K$	83.05%	50.4s

- A subset of FRGCv2, with 1,398 gallery images and 466 probe images, resulting in K=111,643 for the dictionary

Extension to Partial Face Verification

Differences with Previous Methods

	Lowe's SIFT	Wright's SRC	MKD-SRC
Size of descriptor per image	variable	fixed	variable
Face image requirement	alignment-free	aligned and cropped	alignment-free
Collaborative Representation	×	√	√
Holistic face	√	√	√
Arbitrary partial face	√	×	√

Differences with Previous Methods

Approach	Scenario	Image requirement	Database used	#Subjects
Subspace [13], [14], [15]; SRC [7], [12]; SVM [16], [17]; Part-based fusion[36], [26], [32], [33], [34], [35]; Single component[29], [30], [31]	Occlusion	Aligned & cropped frontal faces	AT&T, ORL, AR, Yale B, FERET, FRGC, Multi PIE	$\leq 1,196$
Multi-view [19], [20]; Cross-view [21], [22], [23], [24], [25], [26]	Arbitrary pose	Alignment via facial landmarks	FERET, PIE	≤ 250
Skin texture [50]	Limited FOV	Frontal face	MBGC	114
Proposed method (MKD-SRC)	Occlusion, arbitrary pose, limited FOV	Alignment free	FRGC, AR, LFW, PubFig	$> 20,000$

Experimental Settings

- Open-set face identification: FRGC 2.0+, AR+, PubFig+
- Face verification: LFW

Database	FRGCv2.0+	AR+	PubFig+	LFW
Scenario	Partial patch	Occlusion	Pose & occlusion	Pose & occlusion
#Subjects	20,466	20,135	5,140	5,749
#Gallery	10,466	10,135	5,083	6,000
#Probe	25,562	11,530	8,027	6,000
Methods	MKD-SRC	MKD-SRC, PittPatt, FaceVACS, PCA+LDA, LBP		MKD-SRC, [70], PittPatt, [71], FaceVACS, PCA+LDA, LBP

Open-set Face Identification

- Task: determine the identity of the probe, or reject the probe
- Practical scenarios: watch-list surveillance, attendance, forensic search, SNS photo tagging, etc.

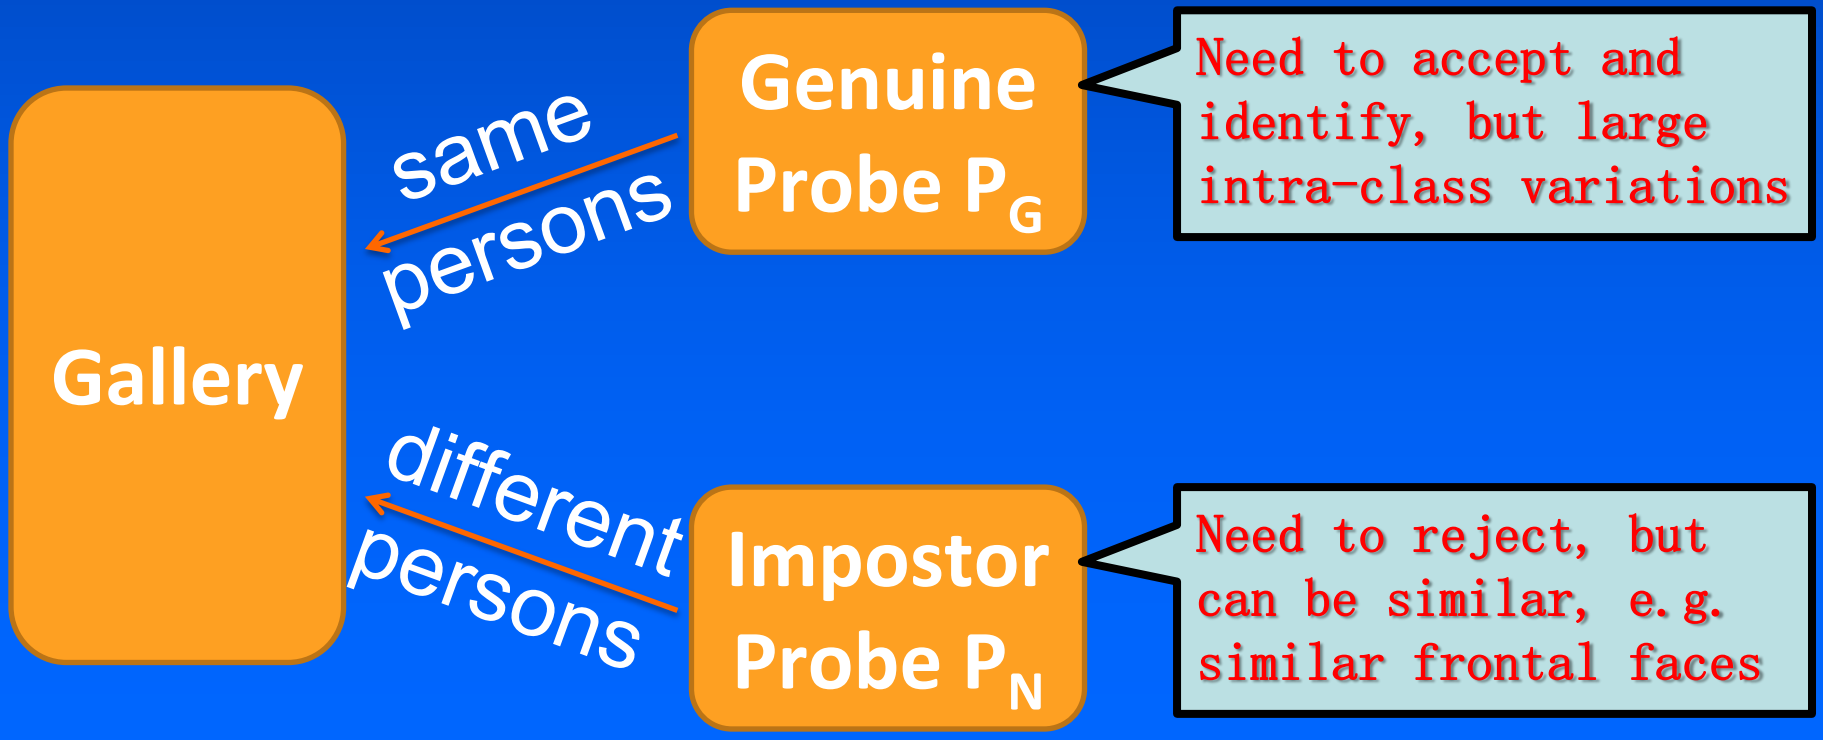

Open-set Face Identification

- Performance measures:
  - Detection and identification rate: percentage of images in P_G that correctly accepted and identified
  - False accept rate: percentage of images in P_N that falsely accepted
  - Detail and a recent benchmark can be seen in http://www.cbsr.ia.ac.cn/users/scliao/projects/blufr/

Experiments with Partial Faces

- FRGCv2.0+
  - Gallery: 466 FRGC + 10,000 background
  - Probe: 15,562 P_G (partial face) + 10,000 P_N
  - 1 image per subject in gallery

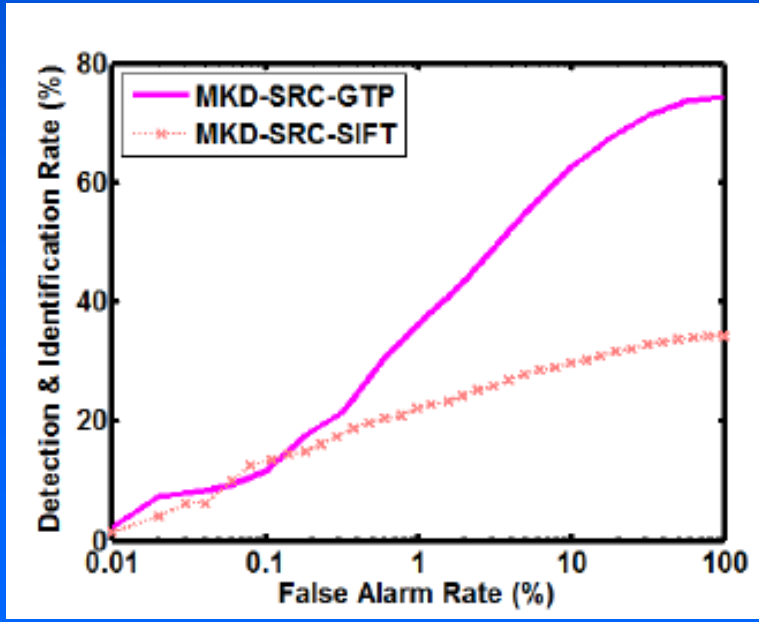

Experiments on Holistic Occluded Faces

- AR+
  - Gallery: 135 AR + 10,000 background
  - Probe: 1530 P_G (occluded) + 10,000 P_N
  - 1 image per subject in gallery
  - each subject has 6 (one session) or 12 (two sessions) images
  - All images in P_G are with sunglasses or scarf, and illumination variations

Experiments on Holistic Occluded Faces

- Gallery




- Probe






  - It can be seen that faces are not aligned very well

Experiments on Holistic Occluded Faces

- Closed-set identification
  - Wright's SRC is not robust with only one training sample per class, though manually aligned faces were used

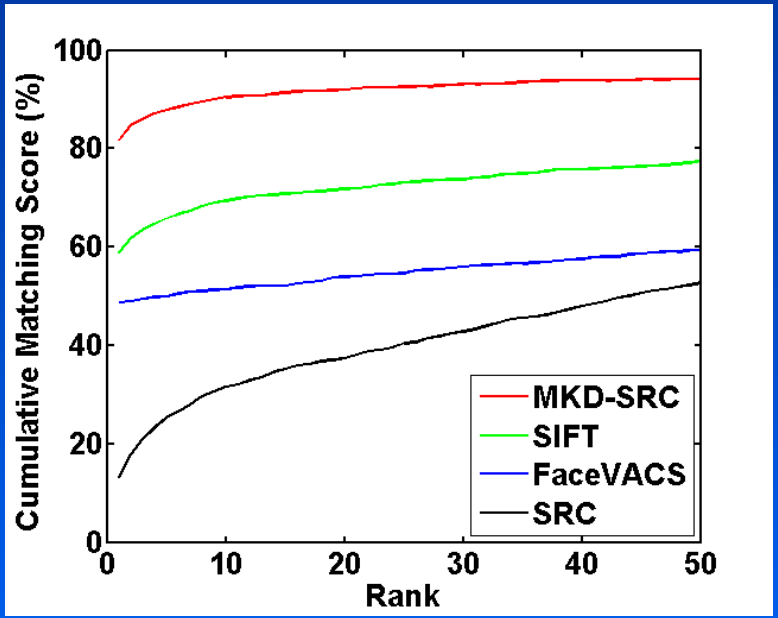

Methods	Recognition Rate (Rank-1 Rate)
MKD-SRC	81.70%
SIFT Keypoint Match by Lowe	58.89%
FaceVACS	48.76%
SRC by Wright et al.	13.20%

Open-set Identification on AR+

- Challenging task:
  - Gallery: frontal, no occlusion, 1 image / class
  - P_G: sunglasses or scarf, illumination
  - P_N: frontal, no occlusion
- MKD-SRC is able to reject 99% impostors (FAR=1%) while accepting >55% genuine samples at rank-1

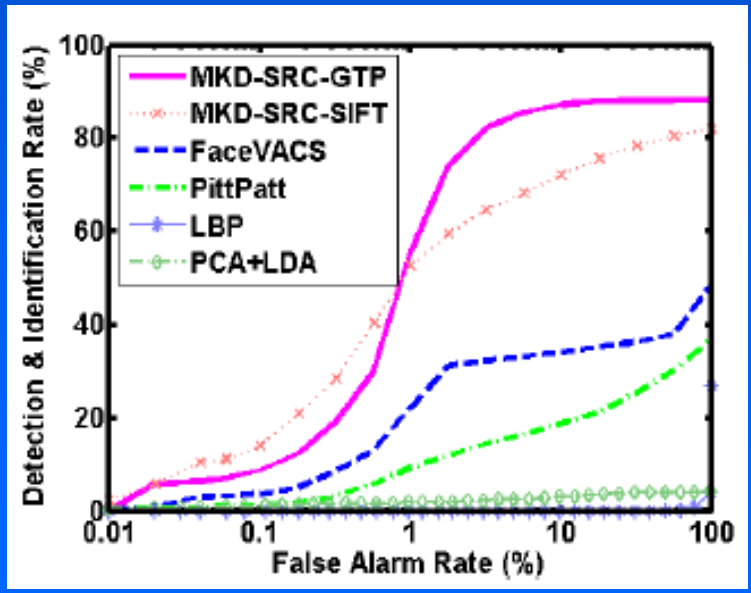

Experiments on Labeled Faces in the Wild (LFW)

- LFW: real faces from the internet, with possible non-frontal view or occlusion
- 13,233 images of 5,749 subjects
- Verification scenario; images restricted protocol
- 10 random subsets for test, with each subset having 300 genuine and 300 impostor pairs

Experiments on LFW

- MKD-SRC outperforms FaceVACS and the best image-restricted method V1-like
- Fusion of MKD-SRC and PittPatt outperforms the best method A.P.
- MKD-SRC-GTP is much better than MKD-SRC-SIFT

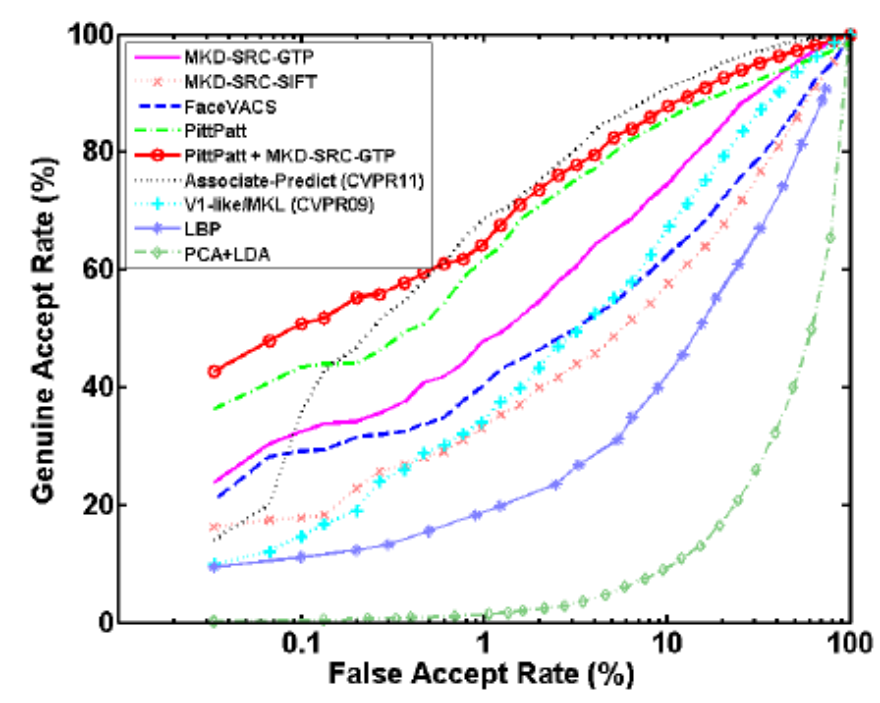

Experiments on LFW

- Correctly (top row) and incorrectly (bottom row) recognized face images from the LFW database by MKD-SRC

Experiments on LFW

- A subset of partial faces from LFW
  - Sunglasses, hats, occlusions by hand or other objects, large pose variations (>45˚)

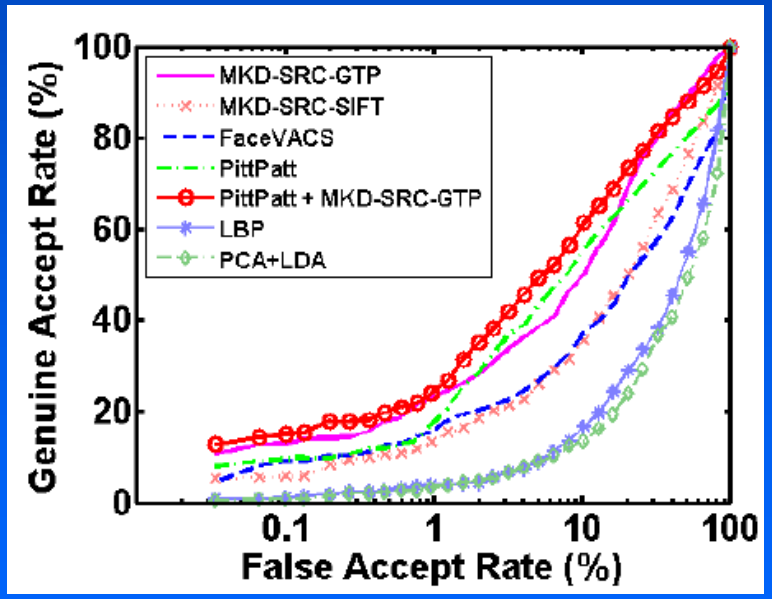

- Note: limitations of LFW and a new benchmark are discussed in http://www.cbsr.ia.ac.cn/users/scliao/projects/blufr/

Experiments on PubFig+

- Gallery: 83 PubFig + 5,000 LFW
- Probe: 817 P_G (occluded) + 7,210 P_N
- 1 image per subject in gallery

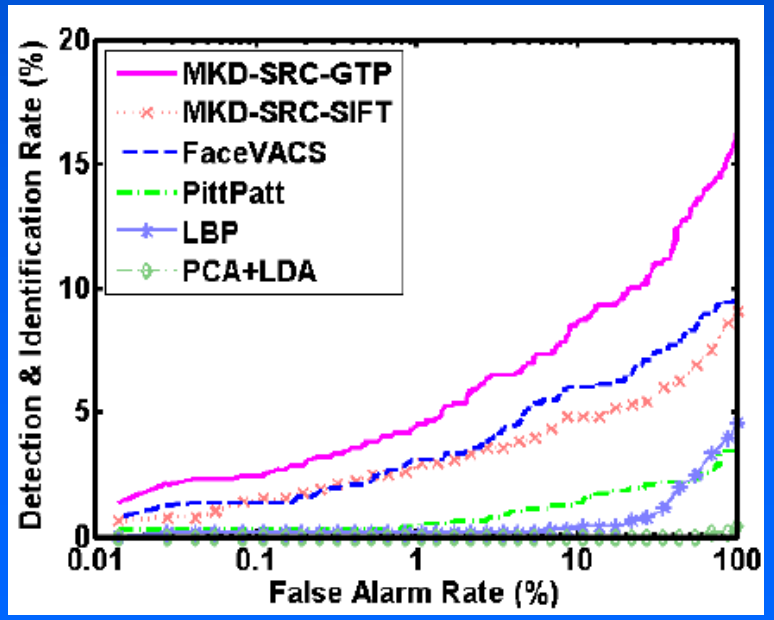

Summary

- Addressing the general partial face recognition problem without alignment
- A unified face recognition framework for both holistic and partial faces
- Improves SRC for the one-sample-per-class problem
- Multi keypoint descriptors enables variable-length face description

Suggestions for Future Work

- PFR is important but difficult. The proposed matching framework is not the only way to recognize partial faces. There are other possibilities, e.g. Weng et al. Robust feature set matching for partial face recognition, ICCV 2013
- There may be other elegant partial face description methods
- Automatic PFR is even more difficult. Partial face detection is still missing

Thank you!