

Discriminative Invariant Kernel Features: A Bells-and-Whistles-Free Approach to Unsupervised Face Recognition and Pose Estimation

Dipan K. Pal, Felix Juefei Xu and Marios Savvides*

CVPR 2016

Carnegie Mellon University
CyLab Biometrics Center

The problem

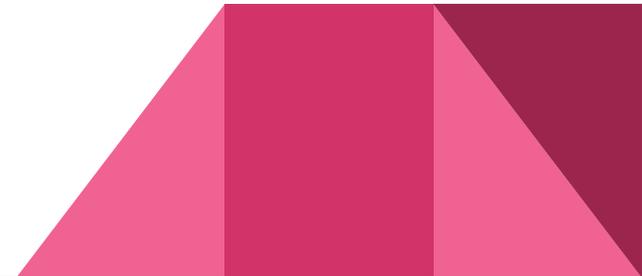
Two complementary tasks

To perform two complementary tasks **simultaneously** using a **single unsupervised** feature extractor



Who is this subject?

What is the subject's pose?



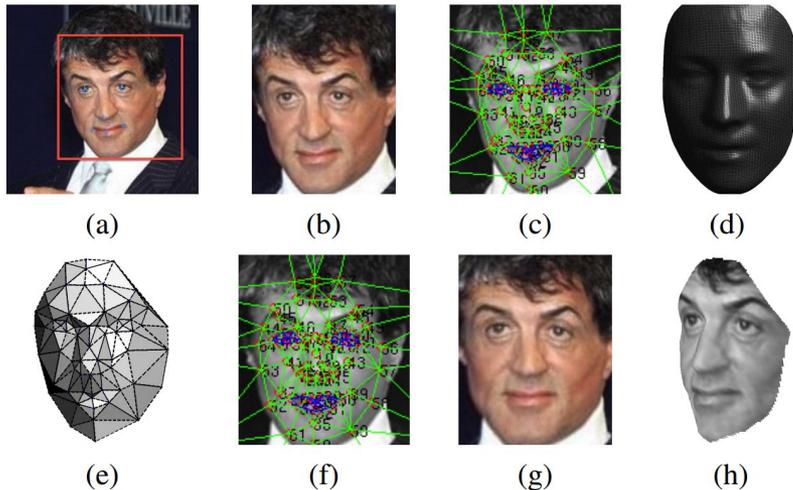
The problem

Landmark-free

The paper focuses on **dense landmark-free** (only two eye center locations) face recognition and pose estimation

Extends to a completely landmark-free approach which is also **alignment free**

This work:



Taigman, Yaniv, et al. "Deepface: Closing the gap to human-level performance in face verification." *CVPR* 2014.

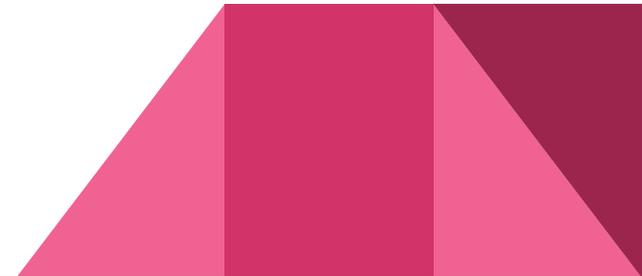
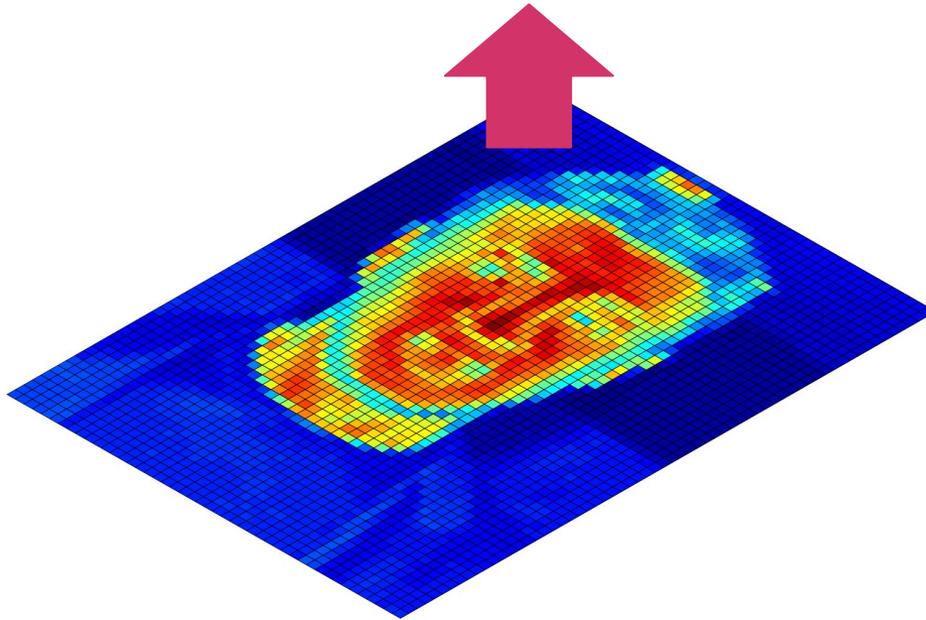


The problem

Discriminative Invariant Feature

We extract a **single** highly **discriminative** provably **group invariant** non-linear feature for both tasks from raw pixels

*Invariant feature extraction
from raw pixels.*

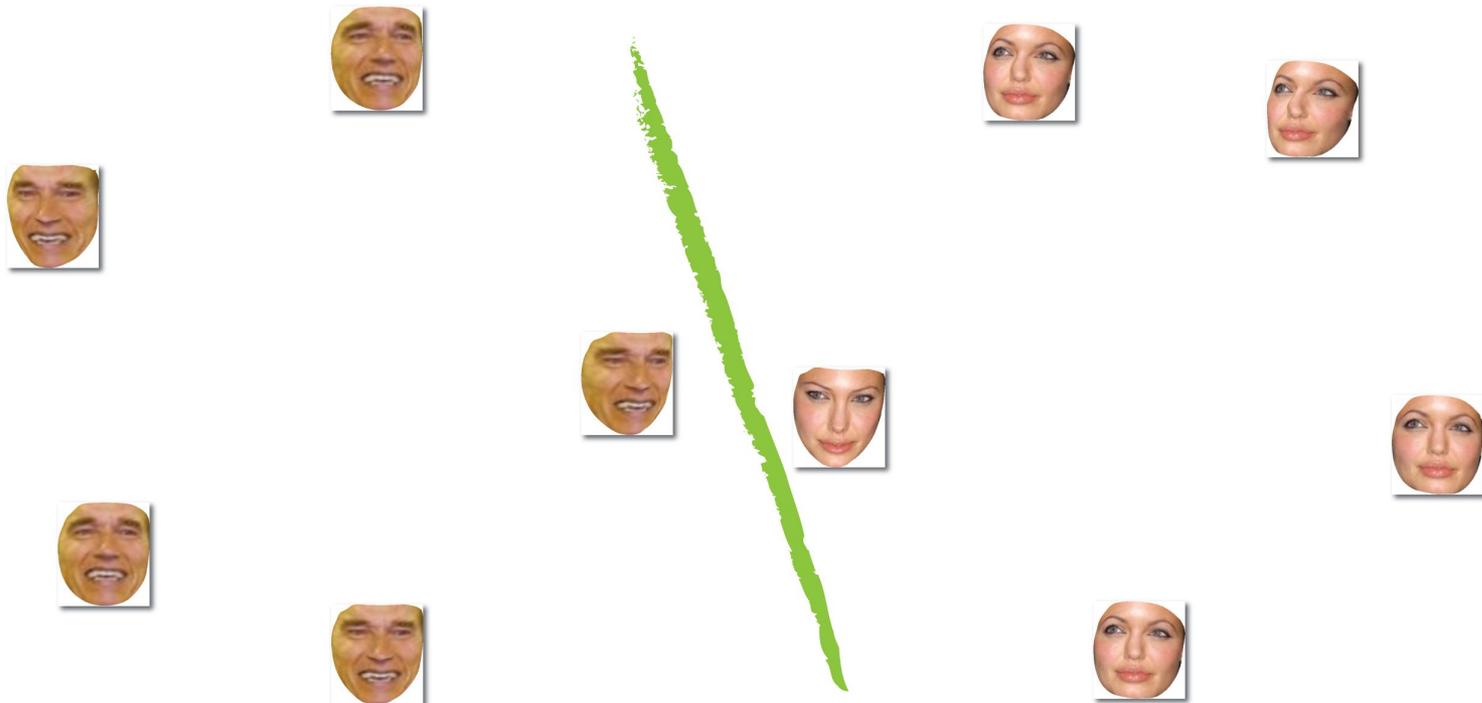


The approach

Invariance to Transformations

Nuisance transformations groups such as the translation, rotation group, **increase complexity** of the learning problem

Invariance to such transformations can drastically **reduce complexity**

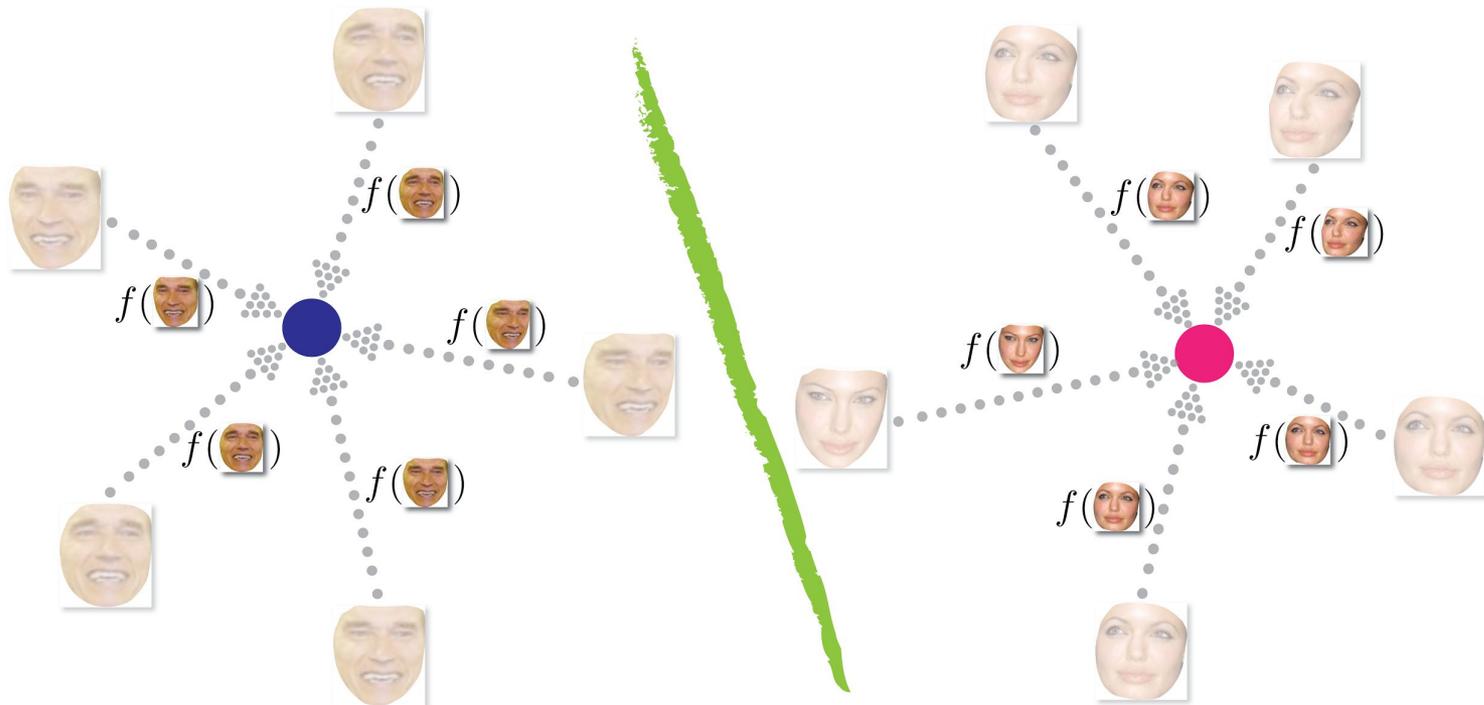


The approach

Invariance to Transformations

Nuisance transformations groups such as the translation, rotation group, **increase complexity** of the learning problem

Invariance to such transformations can drastically **reduce complexity**



The approach

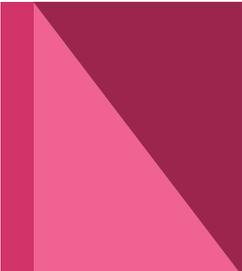
Invariance to Transformations

Nuisance transformations groups such as the translation, rotation group, **increase complexity** of the learning problem

Invariance to such transformations can drastically **reduce complexity**


$$f(\text{img}_1) = \dots = f(\text{img}_2)$$

$$f(\text{img}_1) = \dots = f(\text{img}_2)$$

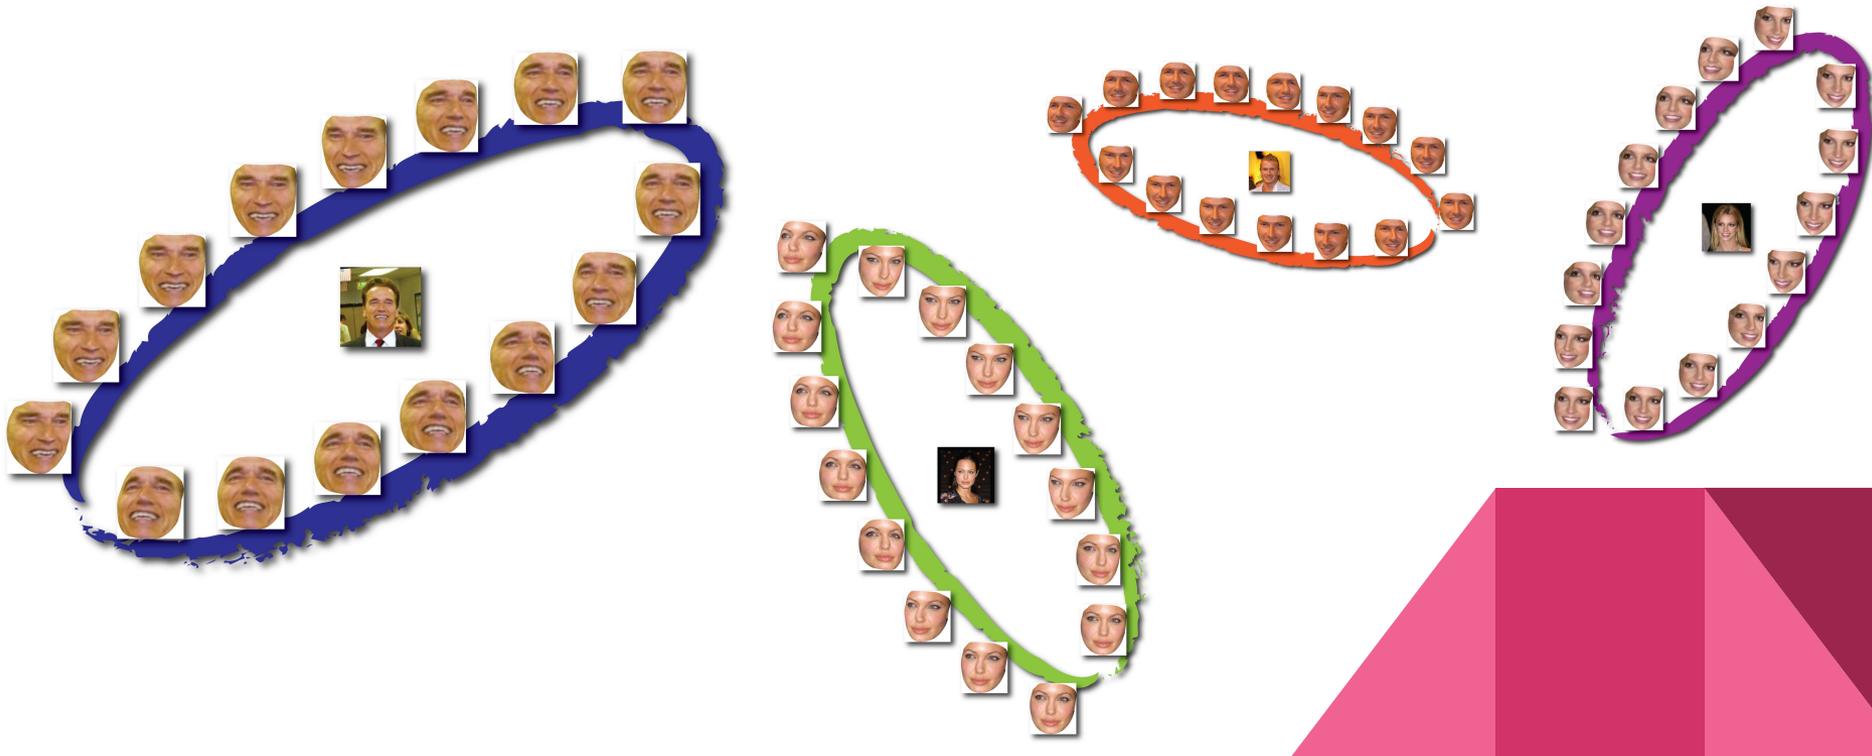


Previous work

Linear Invariant Features

Previous work builds linear invariant that are **implicitly** (but not explicitly) discriminative

When a group of transformations act on an object, they create an **orbit**



Previous work

Linear Invariant Features

Previous work builds linear invariant that are **implicitly** (but not explicitly) discriminative

The orbit is **unique** to the object, and is an **invariant** to the transformation group



is invariant

Previous work

Linear Invariant Features

Previous work builds linear invariant that are **implicitly** (but not explicitly) discriminative

Hence **any measure** of the orbit is an **invariant implicitly discriminative** feature

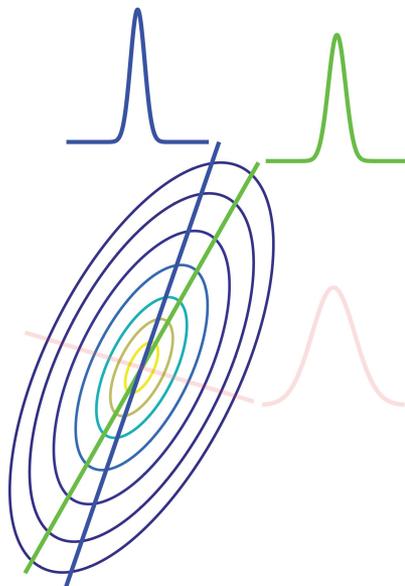


Our approach

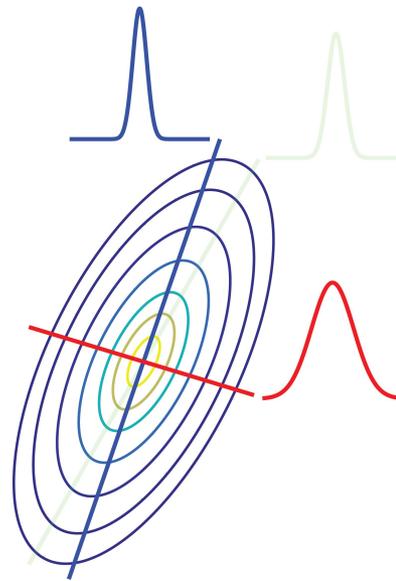
Linear Discriminative Invariance

To characterize the orbit, previously simply sampled templates were used

Explicit discrimination provides better matching. More over the **learnt templates** still **form a group** of transformed templates, hence invariance theory holds.



Sampled templates



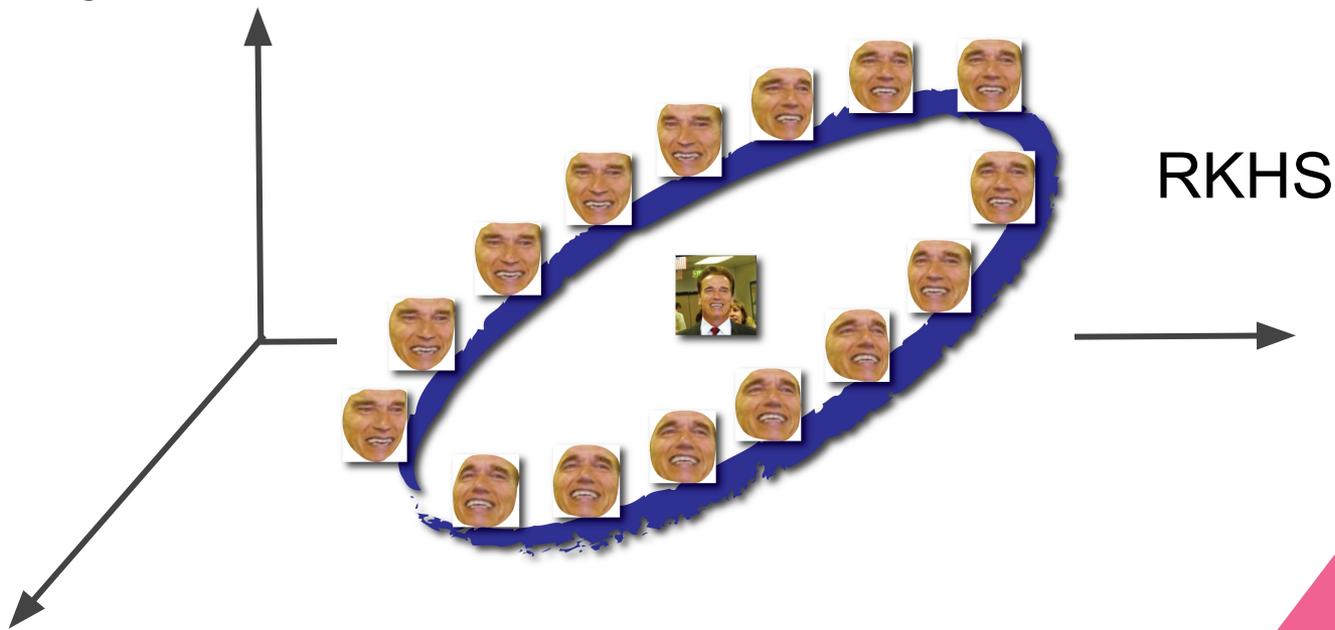
Discriminatively learned templates

Our approach

Non-linear Discriminative Invariance

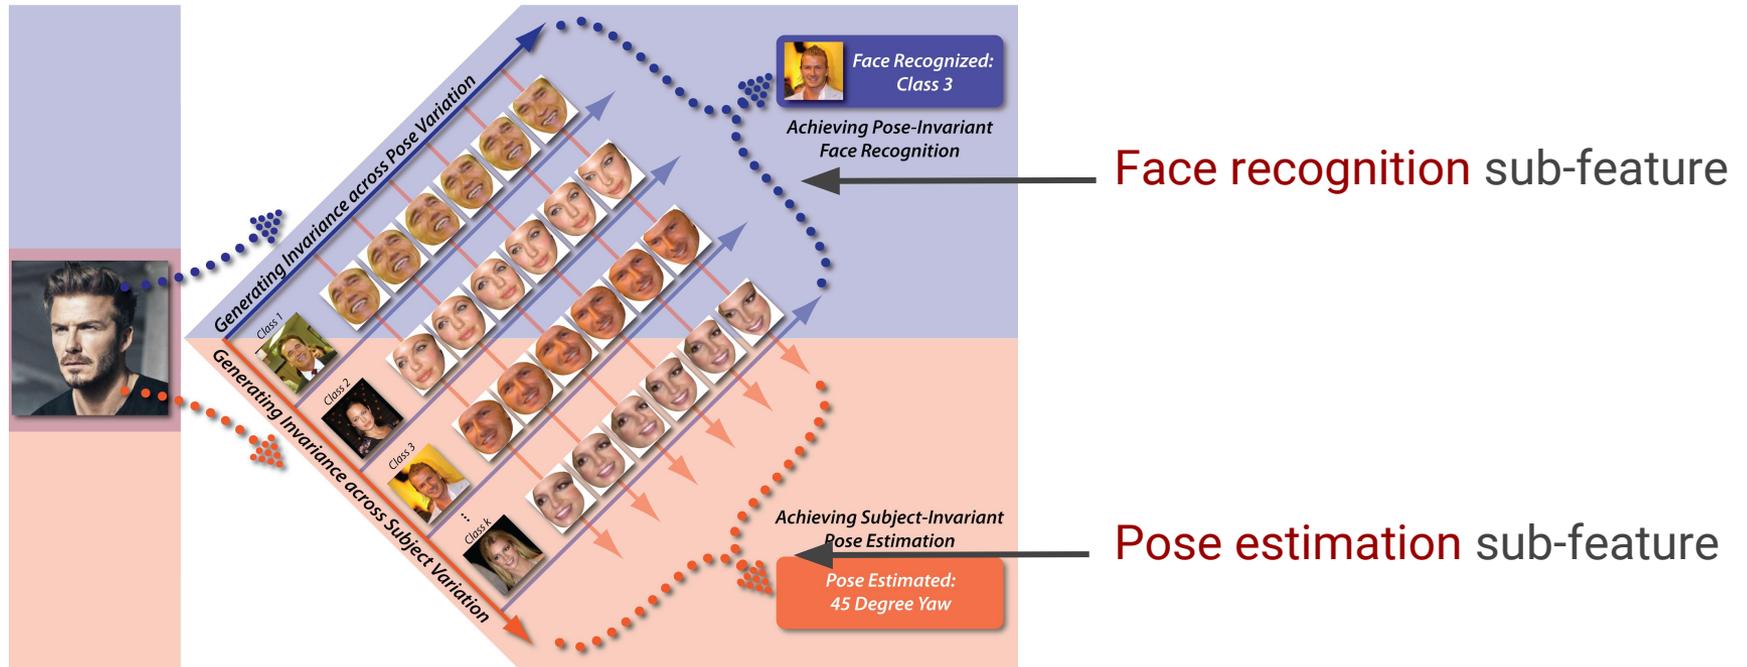
To improve discrimination, we can compute invariant features in the RKHS

We show the **discriminative non-linear templates form a group** in the RKHS, leading to Discriminative Invariant Kernel Features

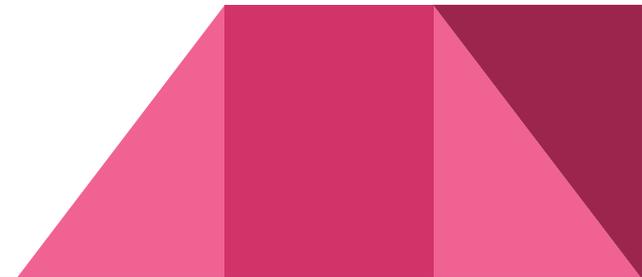


Our System

Final system is versatile



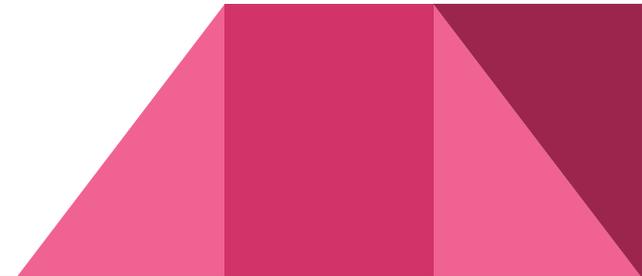
A snapshot of results Semi Synthetic Face Data



A snapshot of results

Semi Synthetic Face Data

- ~153,000 semi-synthetic image dataset of 1000 subjects with 153 poses each.



A snapshot of results

Semi Synthetic Face Data

- ~153,000 semi-synthetic image dataset of 1000 subjects with 153 poses each.
- Images rendered from a 3D model with real texture.



A snapshot of results

Semi Synthetic Face Data

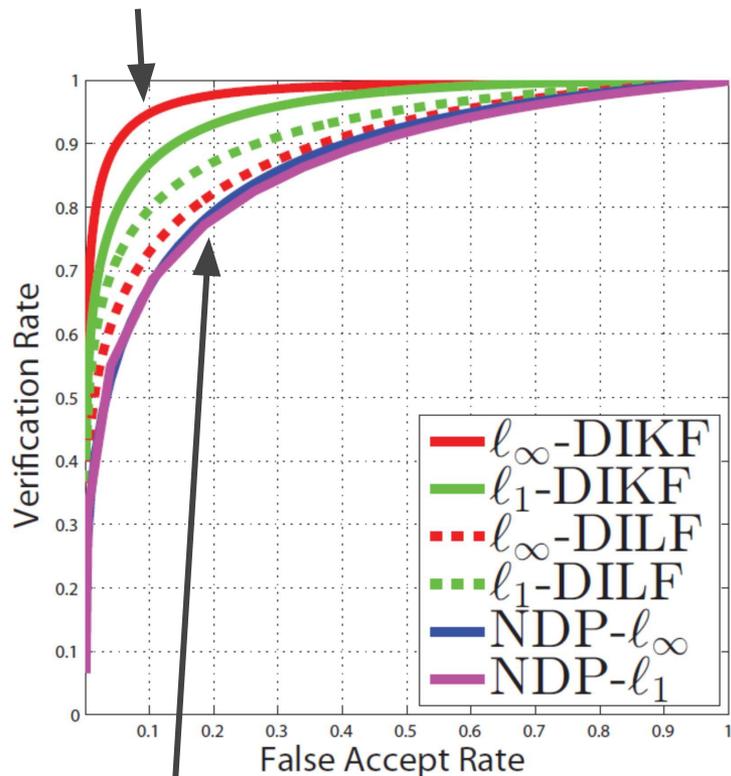
- ~153,000 semi-synthetic image dataset of 1000 subjects with 153 poses each.
- Images rendered from a 3D model with real texture.
- We compare DIKF against sampled templates (NDP) and discriminative linear templates (DILF).



A snapshot of results

Semi Synthetic Face Data

Max Pooled Discriminative Kernel Features



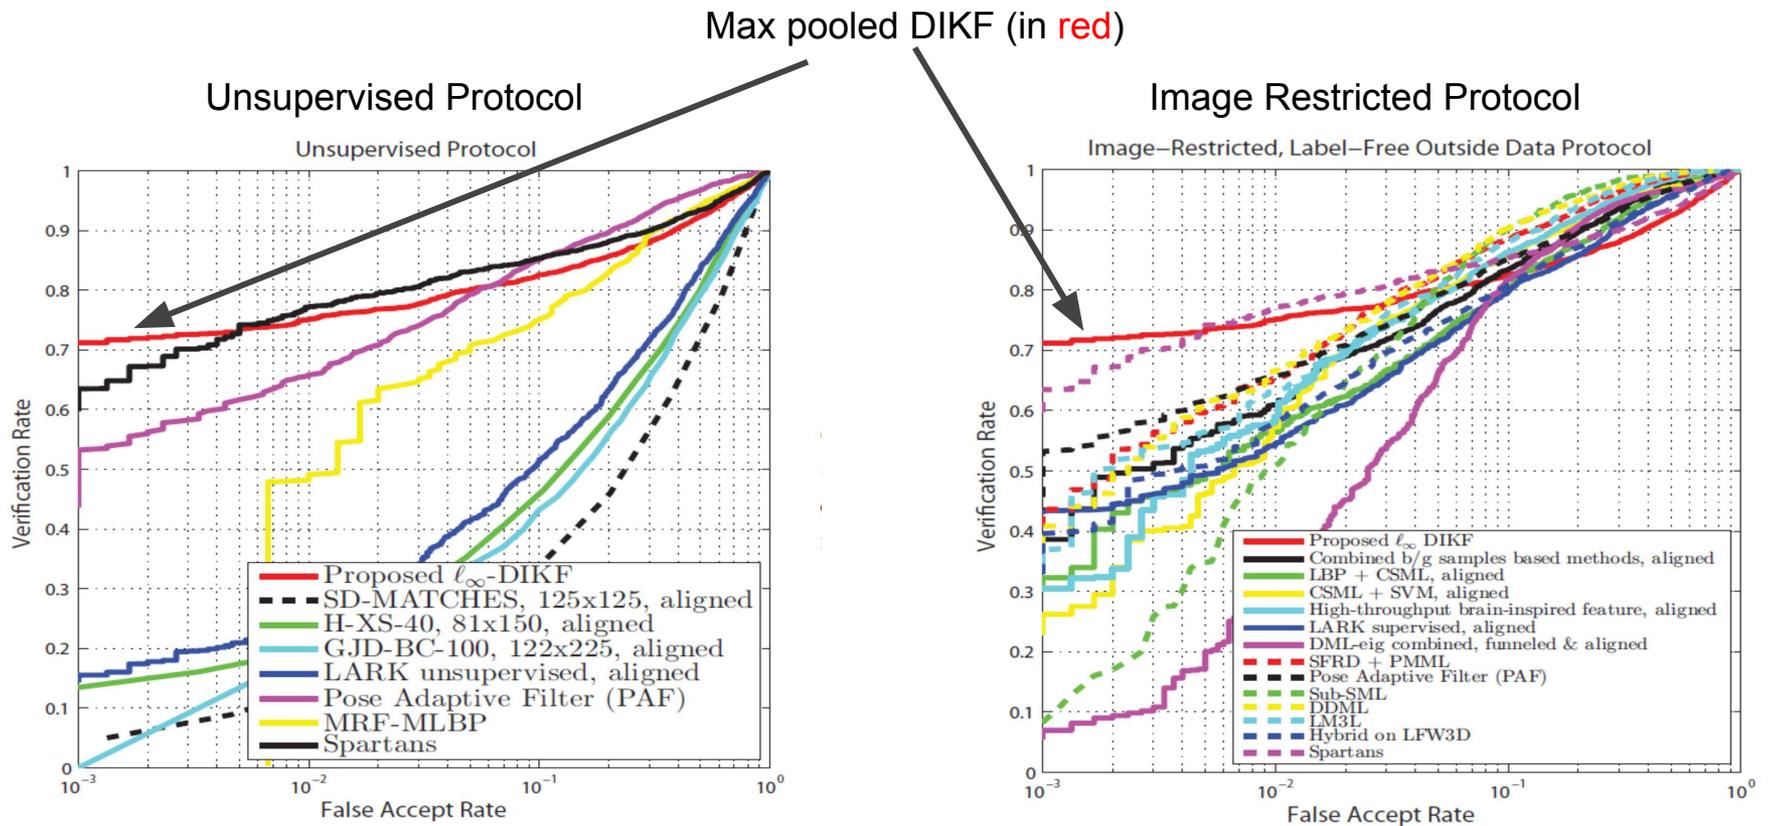
Sampled templates max/mean pooled



A snapshot of results

Labeled Faces in the Wild (LFW)

Max-pooled DIKF **matches state-of-the-art** results on two LFW protocols, despite being **simpler** than competing methods and working on **raw pixels**



Find us at Poster # 32



Dipan K. Pal

ECE PhD, CMU

Visual representation,
invariant features, sparse
signal processing, machine
learning



Felix Juefei Xu

ECE PhD, CMU

Computer vision, machine
learning, unconstrained
face recognition and
biometrics



Marios Savvides

Research Professor, CMU

Director of CyLab
Biometrics Center, CMU