

Poselets


Michael Krainin

CSE 590V

Oct 18, 2011

Person Detection

- Dalal and Triggs '05
 - Learn to classify pedestrians vs. background
 - HOG + linear SVM
 - Doesn't account for variations in body pose and viewpoint



Poselets



**Poselets capture part of the pose from a given
viewpoint**

[Bourdev & Malik, ICCV09]

Poselets



But how are we going to create training examples of poselets?

[Bourdev & Malik, ICCV09]

How do we train a poselet for a given pose configuration?



[Bourdev & Malik, ICCV09]

Finding correspondences at training time

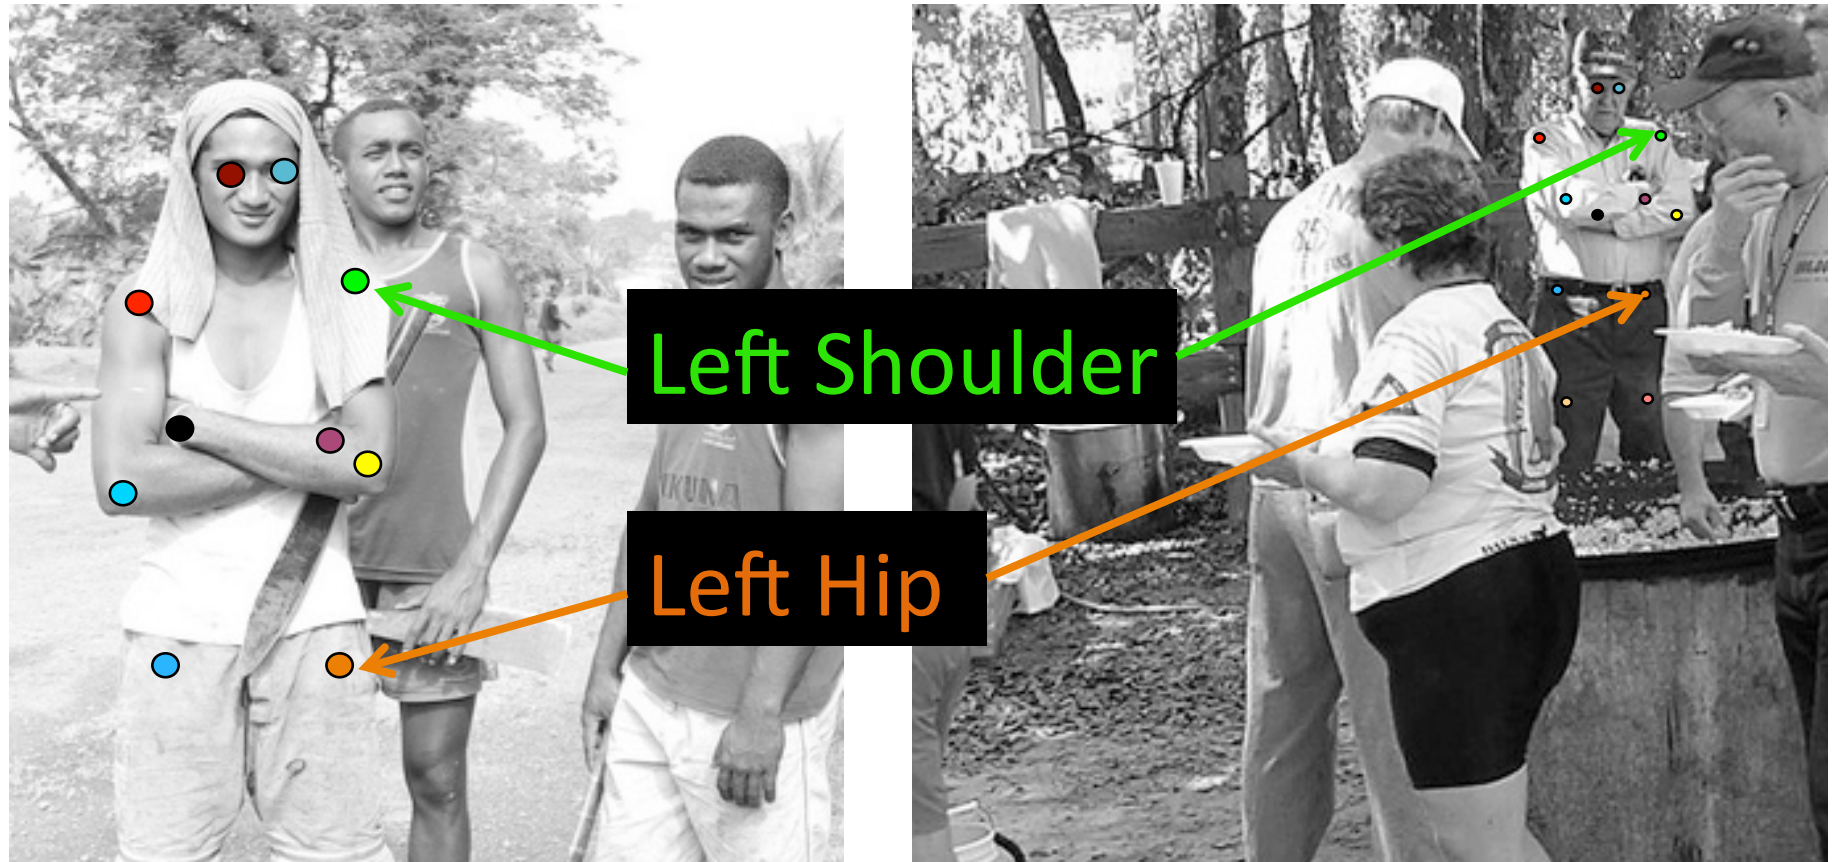


Given part of a human pose



How do we find a similar pose configuration in the training set? [Bourdev & Malik, ICCV09]


Finding correspondences at training time



We use keypoints to annotate the joints, eyes, nose, etc. of people

[Bourdev & Malik, ICCV09]


Finding correspondences at training time



Residual Error



Training poselet classifiers



Residual
Error:

0.15

0.20

0.10

0.85

0.15

0.35

1. Given a seed patch
2. Find the closest patch for every other person
3. Sort them by residual error
4. Threshold them

Training poselet classifiers




1. Given a seed patch
2. Find the closest patch for every other person
3. Sort them by residual error
4. Threshold them
5. Use them as positive training examples to train a linear SVM with HOG features

[Bourdev & Malik, ICCV09]

Which poselets should we train?

- Choose thousands of random windows, generate poselet candidates, train linear SVMs



- Select a small set of poselets that are:
 - Individually effective
 - Complementary


Person Detection Using Poselets

	100 poselets	200 poselets	[12]
VOC 2007	45.6%	46.9%	36.8%
VOC 2008	54.1%	52.6%	43.1%

- [12] Felzenszwalb et al. Object Detection with Discriminatively Trained Part Based Models. PAMI 2010.

Other Uses of Poselets

- Object Segmentation (CVPR 2011)
 - Predict area using “soft masks”
 - Deformation to match image edges
 - Extends poselets to other objects
- Activity Recognition (CVPR 2011)
 - Recognition from a single image
 - Use which poselets fired and to what extent to predict the activity



Describing People: A Poselet-Based Approach to Attribute Classification


L. Bourdev, S. Maji, and J. Malik

ICCV 2011

How poselets help in high-level vision



Gender recognition with poselets



[Bourdev et al., ICCV11]

Attribute Classification Overview



Given a test image



Poselet
Activations



[Bourdev et al., ICCV11]

Features

- Pyramid HOG
- LAB histogram
- Skin features
 - Hands-skin
 - Legs-skin



[Bourdev et al., ICCV11]

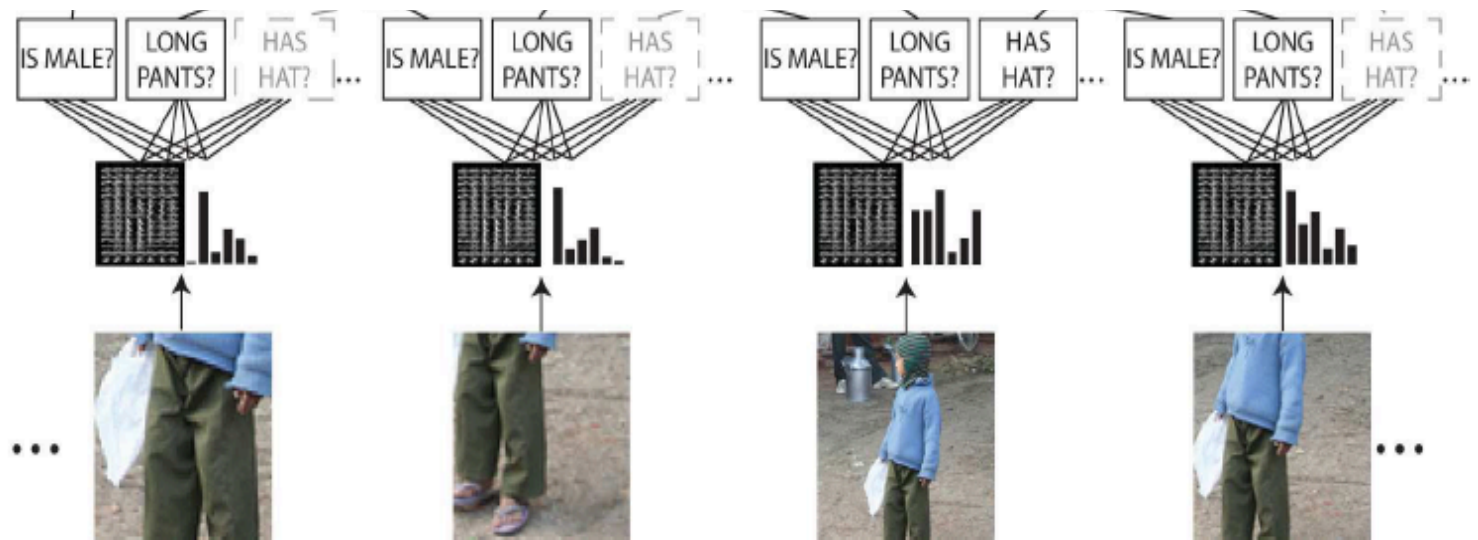
Poselet Level Classification

- Linear SVM + Sigmoid
- Some attributes are associated with a particular body part

Poselet-level
Classifiers

Features

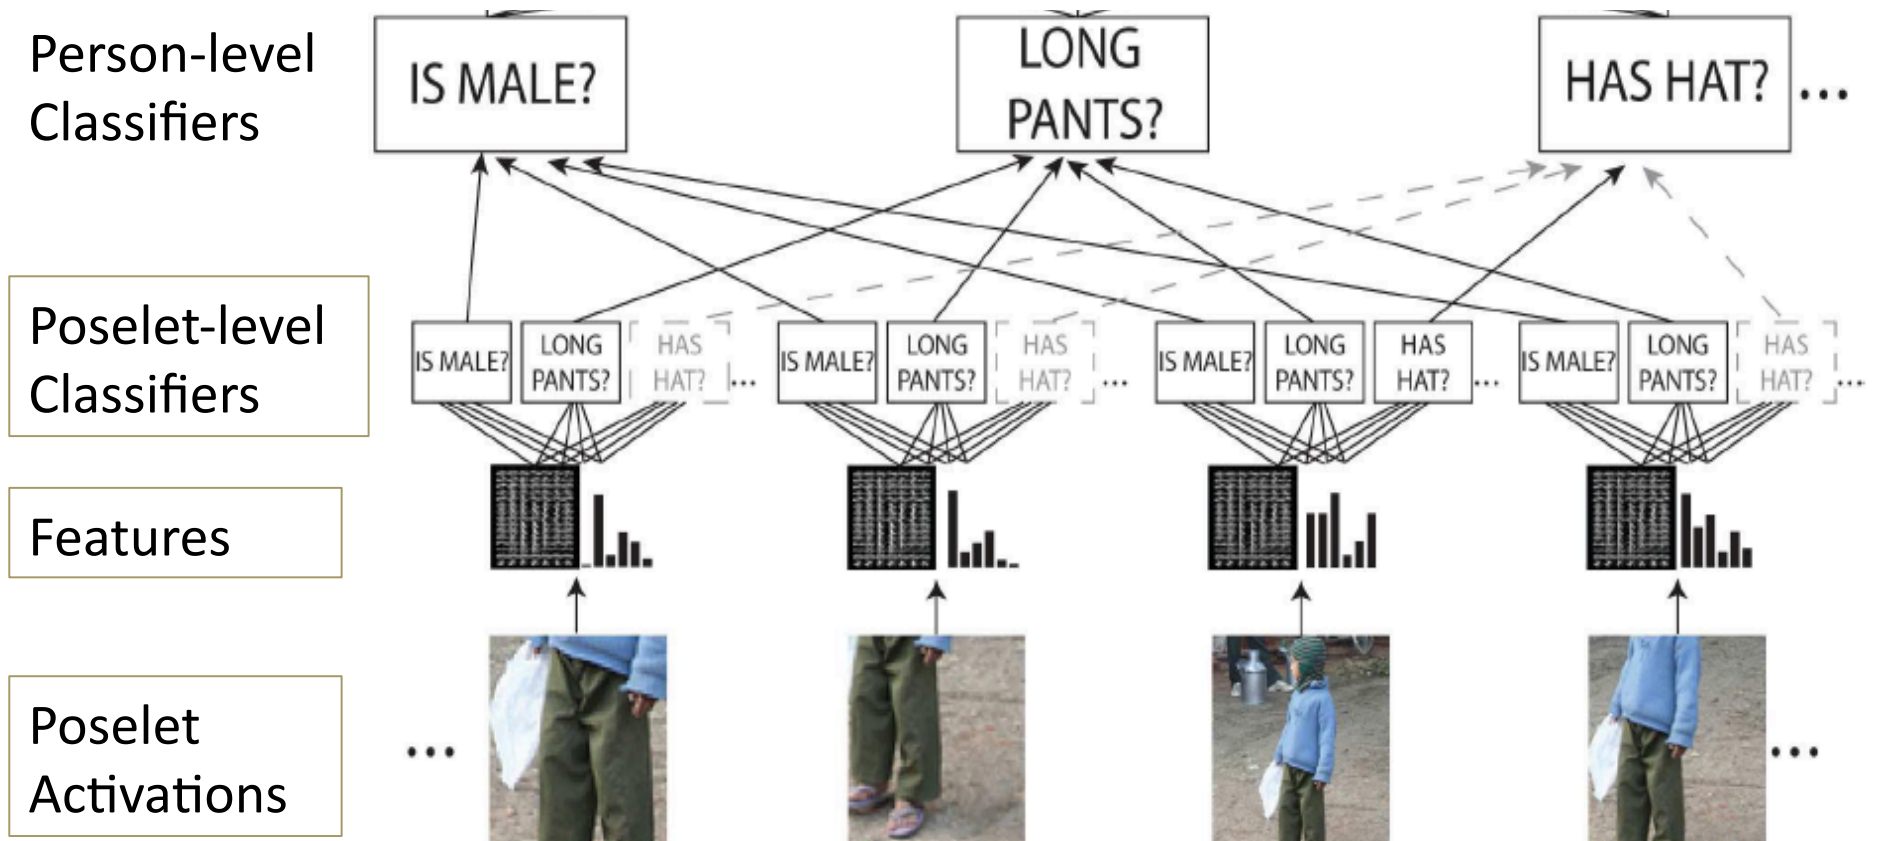
Poselet
Activations



[Bourdev et al., ICCV11]

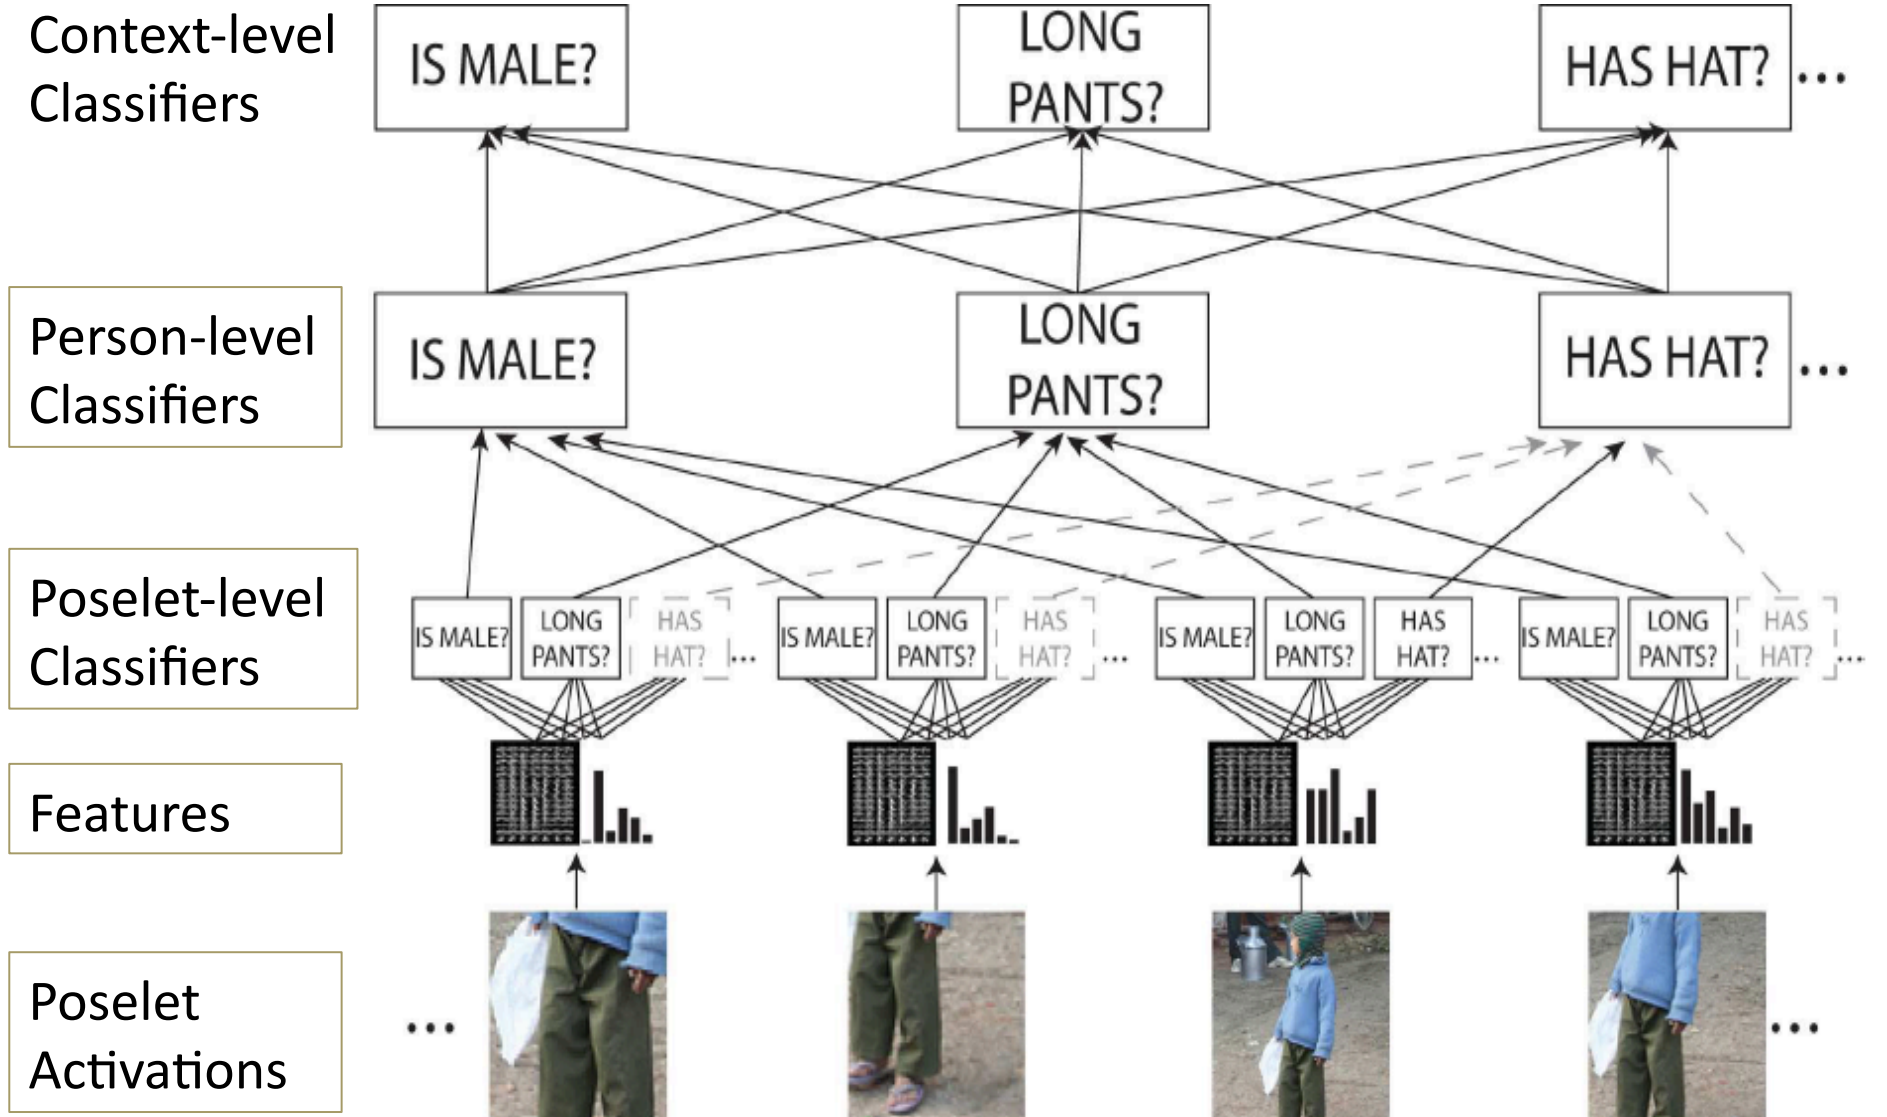
Person-level Classification

- Another Linear SVM + Sigmoid



[Bourdev et al., ICCV11]


Context-Level Classification



[Bourdev et al., ICCV11]

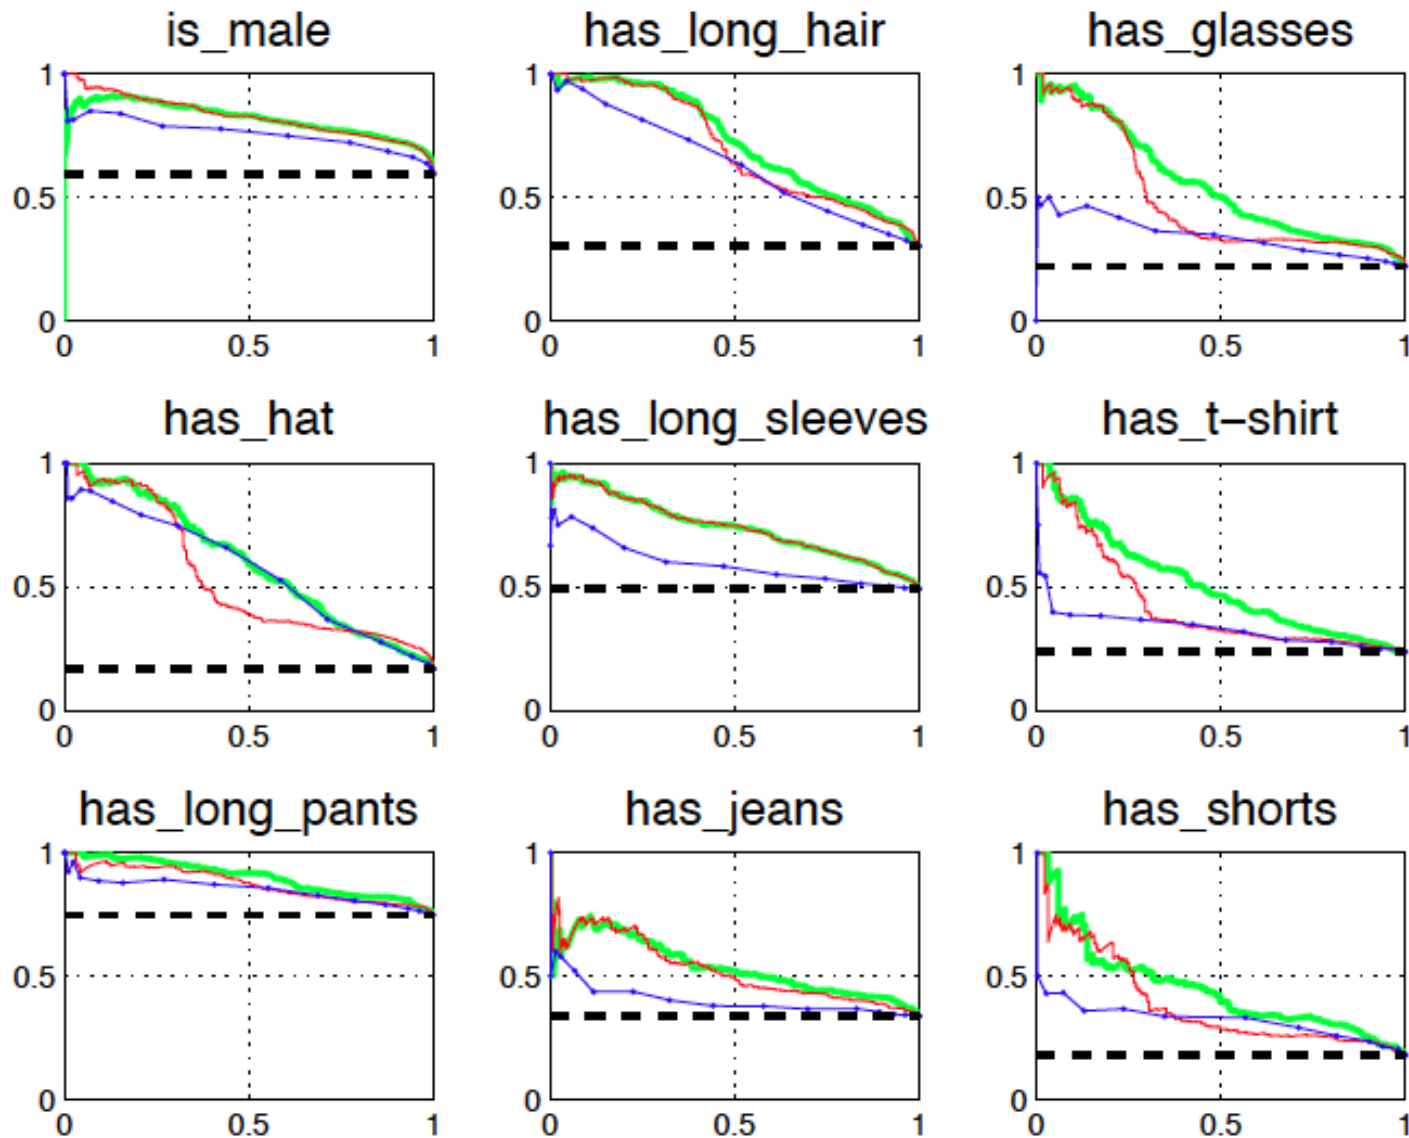
Results

Random 1¼% of the test set



[Bourdev et al., ICCV11]

Precision/recall on our test set



Label Frequency

SPM Poselets-No-context

Poselets-Full

[Bourdev et al., ICCV11]

Gender Recognition vs. Cognitec

- Cognitec is one of the leading face recognition companies (according to latest NIST test)

	Cognitec	Poselets
Scope	Frontal faces only	Any view
Features	Proprietary biometric	Standard HOG
Min resolution	60pixel face width	Any resolution
Performance	75.0% AP	82.4% AP

- Upper bound for any gender recognizer based on frontal faces: Max AP=80.5% vs. 82.4%

Highest-scoring “is male”



Lowest-scoring “is male”

[Bourdev et al., ICCV11]

Highest-scoring “has long hair”



Lowest-scoring “has long hair”

[Bourdev et al., ICCV11]

Highest-scoring “wears a hat”



Lowest-scoring “wears a hat”

[Bourdev et al., ICCV11]

Highest-scoring “wears glasses”



Lowest-scoring “wears glasses”

[Bourdev et al., ICCV11]

Highest-scoring “long sleeves”




Lowest-scoring “long sleeves”

[Bourdev et al., ICCV11]






Little Leaguer



Kids with Santa



The Graduate





Newlyweds

“100 Special Moments” by Jason Salavon

Highly-Weighted Poselets

is-male



has-long-hair



has-glasses




Which poselets are discriminative for gender?


Preferred by human subjects



Preferred by our system



Describing people



Poselets website

<http://eecs.berkeley.edu/~lbourdev/poselets>

- The set of published poselet papers
- H3D data set + Matlab tools
- Java3D annotation tool + video tutorial
- Matlab code to detect people using poselets
- Latest trained poselets