Interclass transfer: why and how

P. Perona - Caltech

(joint work with A. Holub and M. Welling)

NIPS 2005 Workshop on Inter-Class Transfer - Dec 2005

Many categories to be learnt












How many training examples?

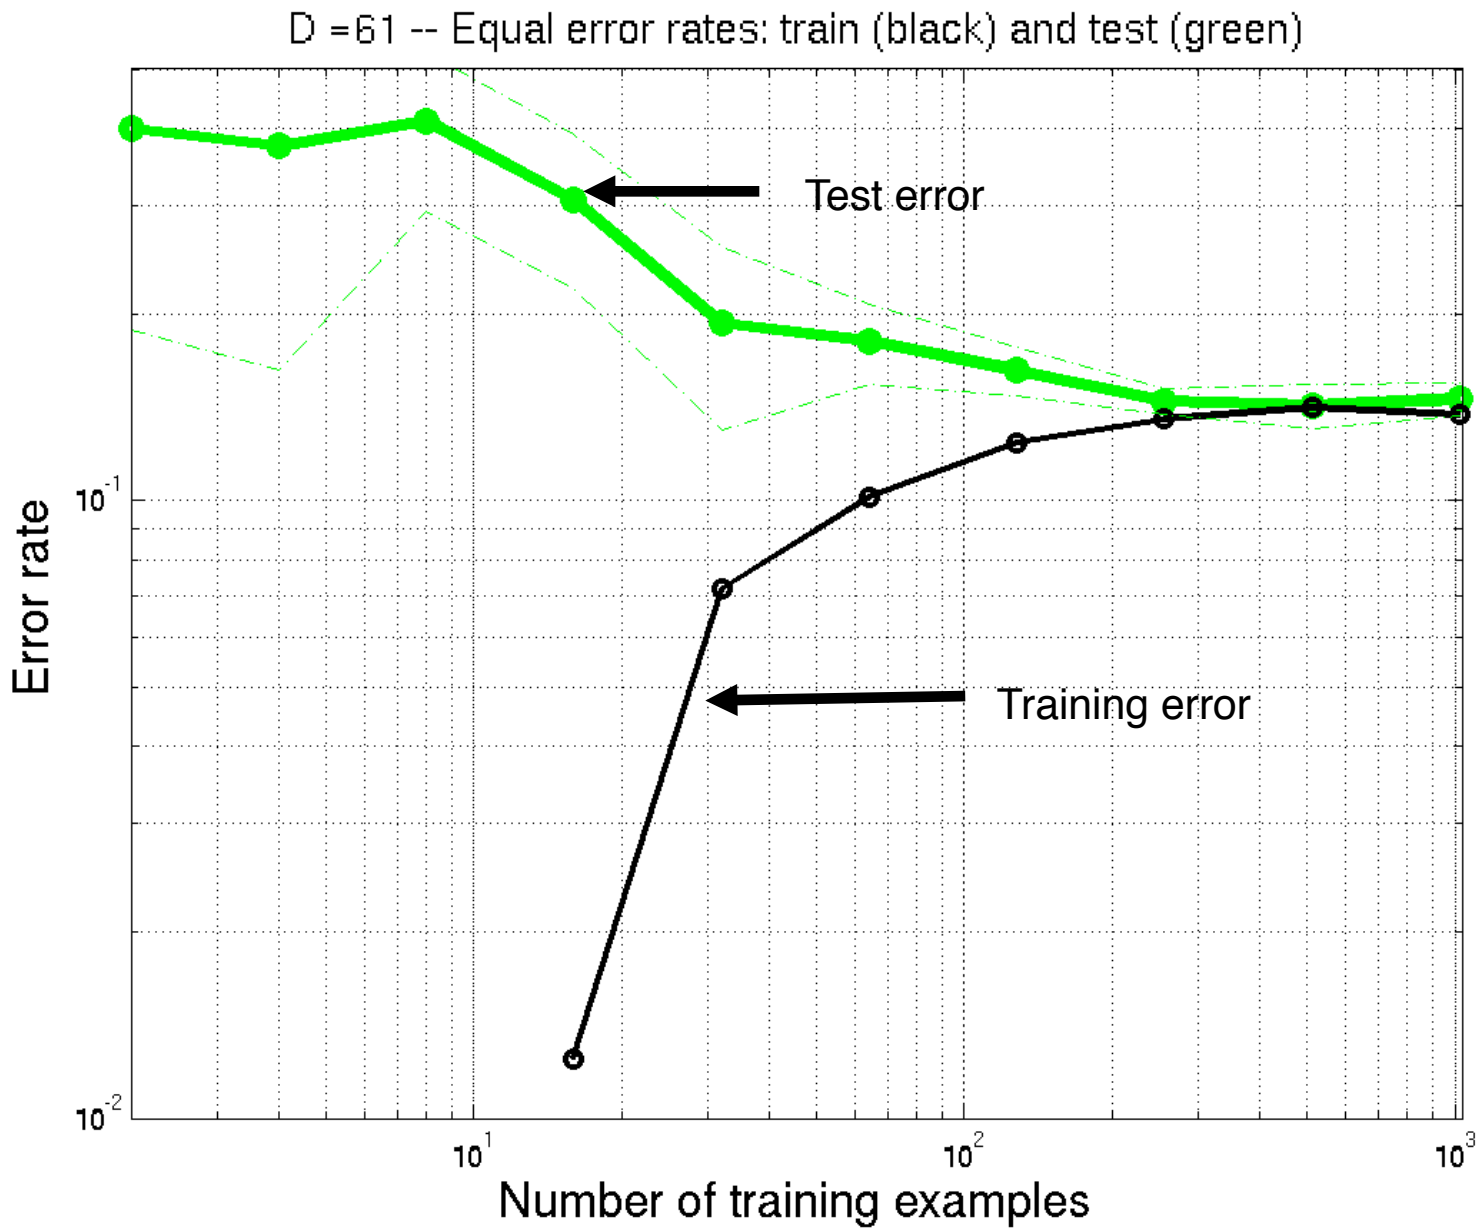

Learn a new category: the pottopod

Find the pottopod

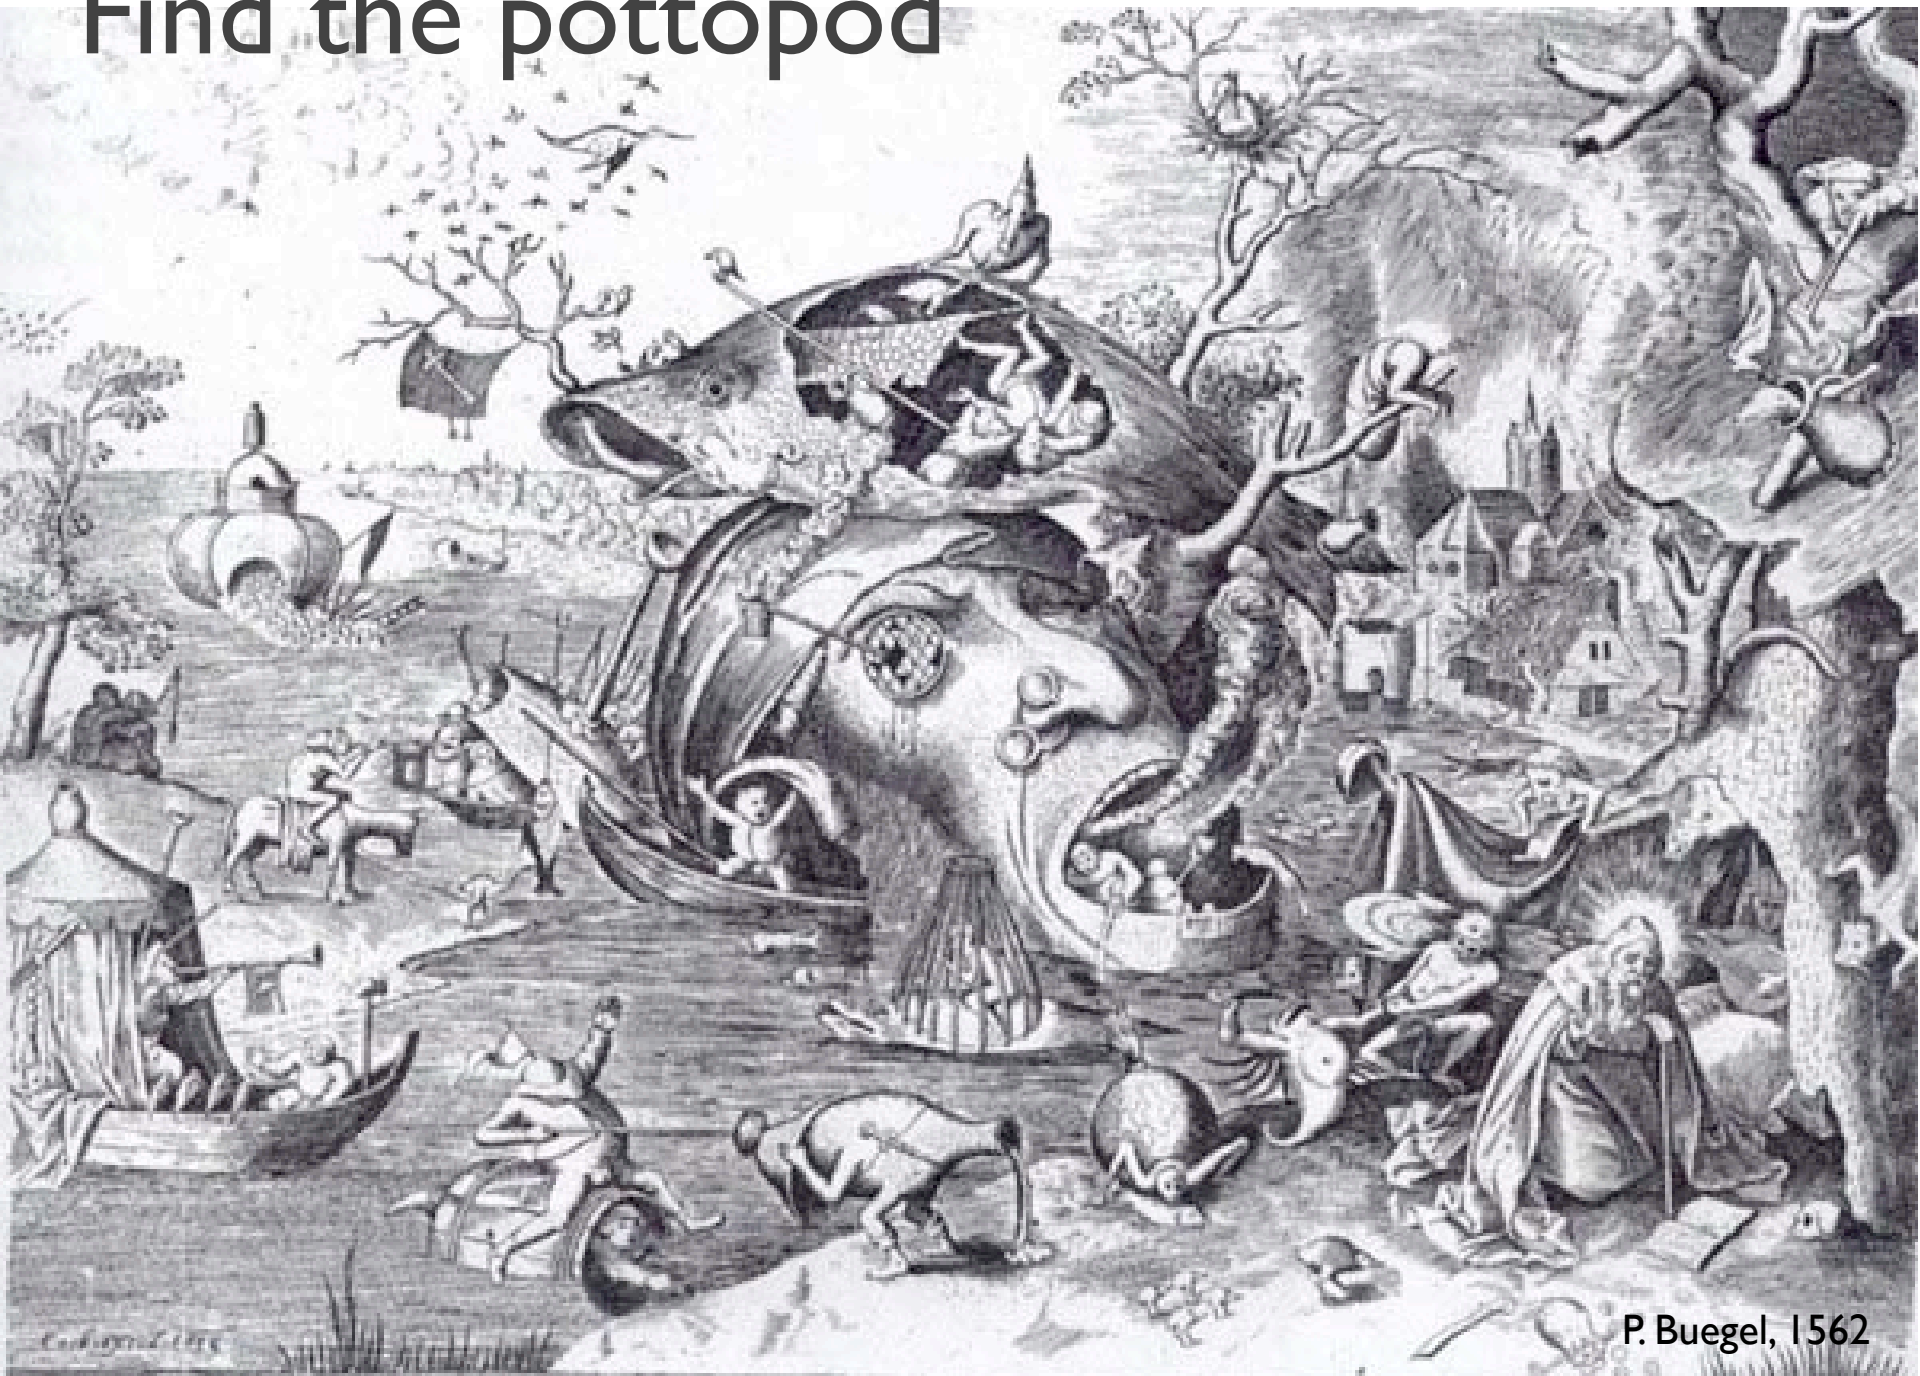

P. Buegel, 1562

Efficient learning

- Can learn a new class with 1 example
- How can this happen?

Training example

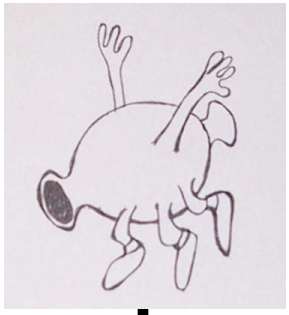

Previously learnt categories

.....

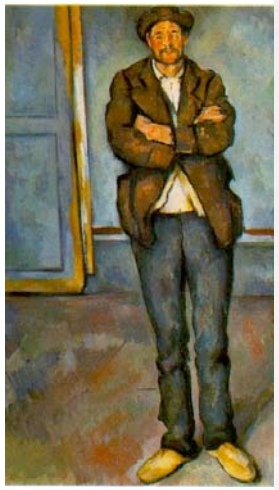

.....

Smart box

New category

Approaches

- **Learn prior from data**
  - Bayesian [Miller et al. '00, Fei Fei et al. '03]
  - Fisher (today's talk) [Holub et al. '05]
- **Hand-crafted prior**
  - Thin-plate + lines [Berg et al. '04,'05]
  - Neocognitron [Poggio et al. '05]

Previously learnt categories

Training example

prior

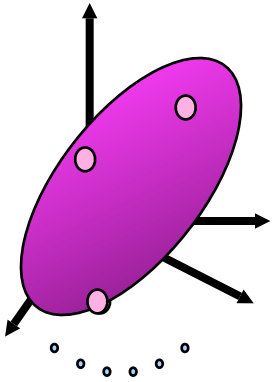

Bayes

New category

[FeiFei et al ‘03]

Priors: geometry, n. parts...

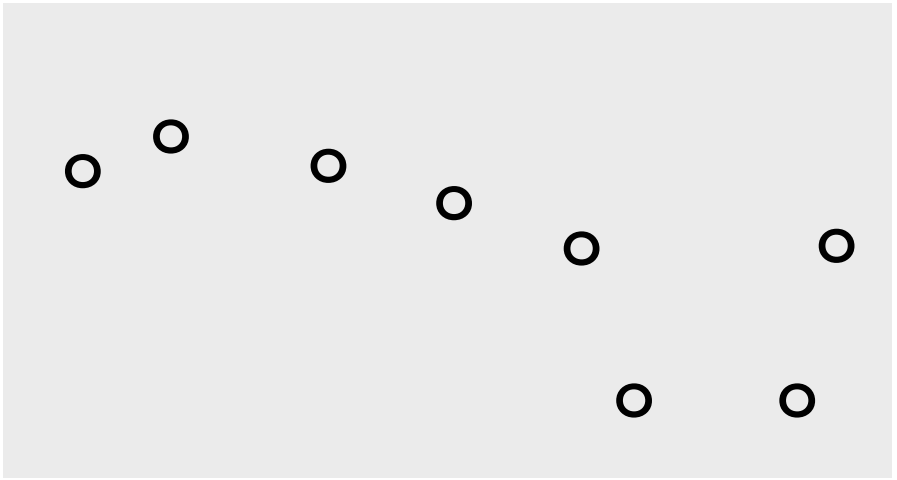

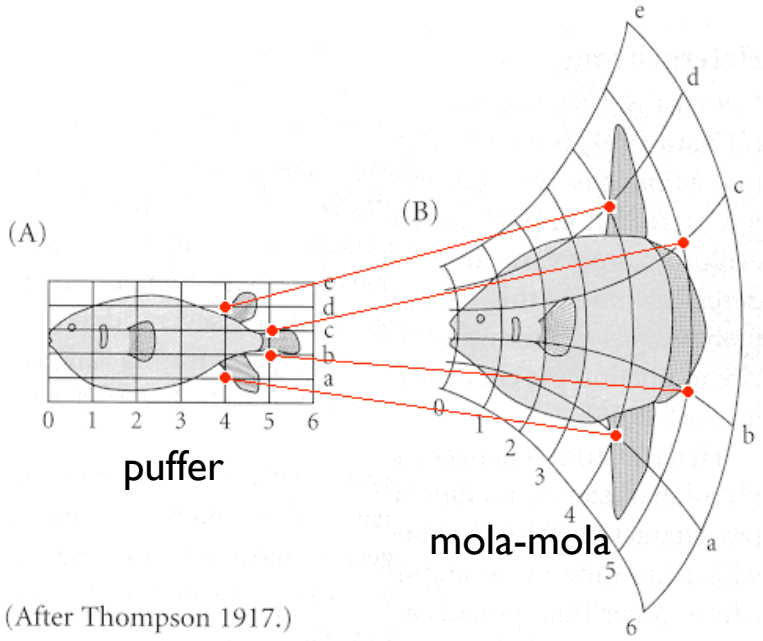

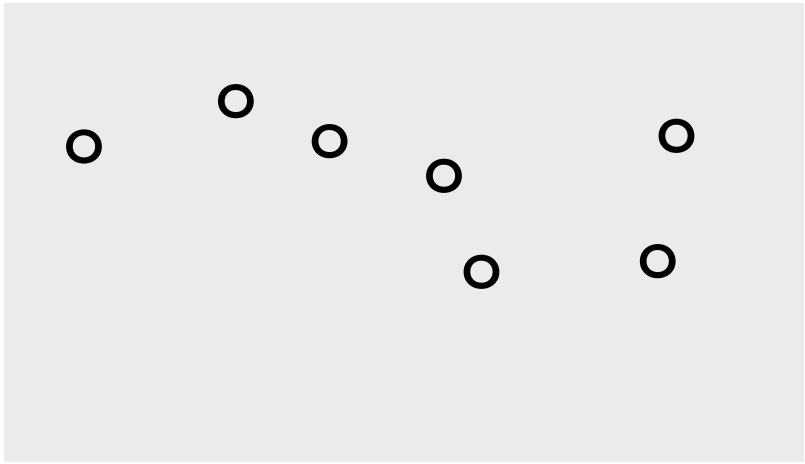

Correspondence Result

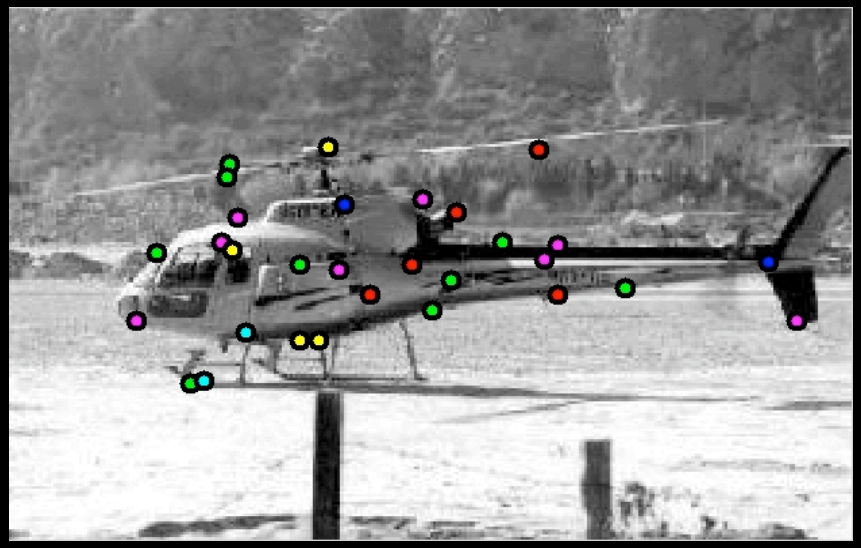

[Berg et al, CVPR05]

Interpolated Correspondence Using Thin Plate Splines

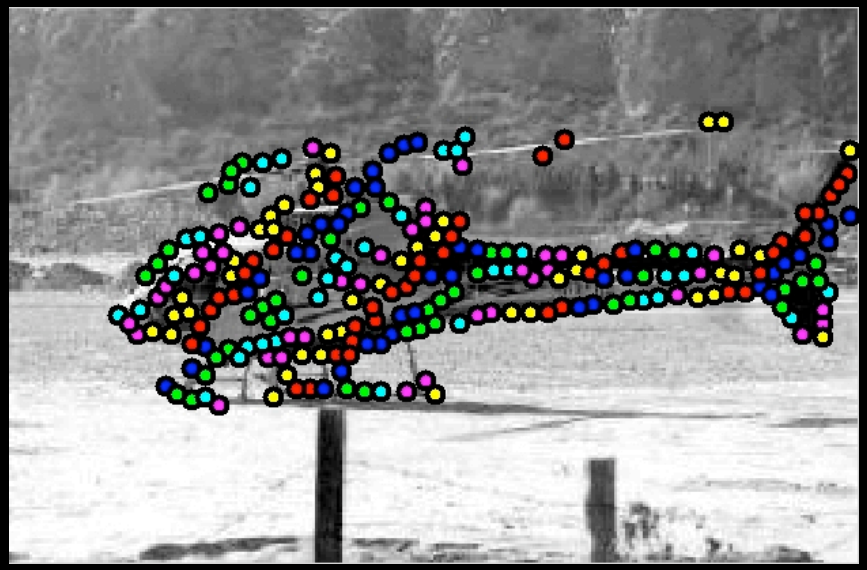

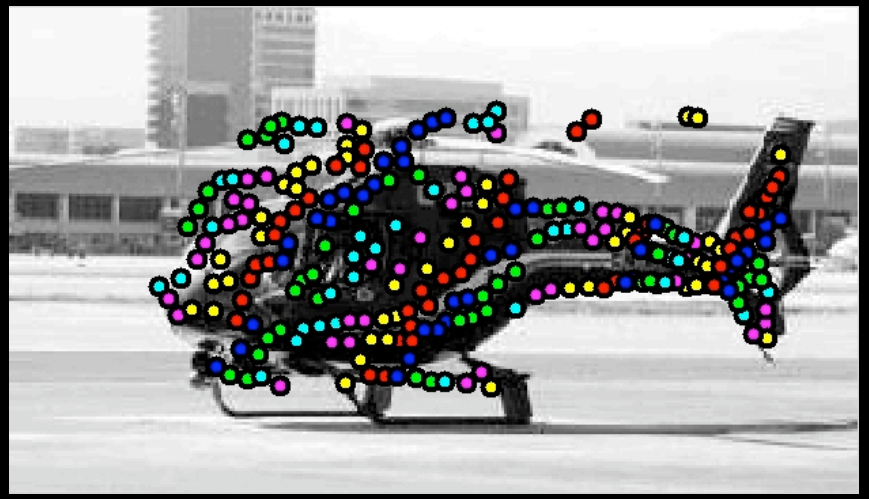

[Berg et al, CVPR05]

Generative vs discriminative

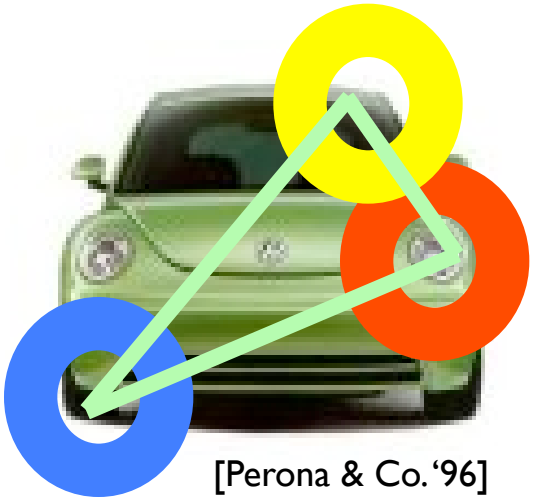

[Perona & Co. '96]

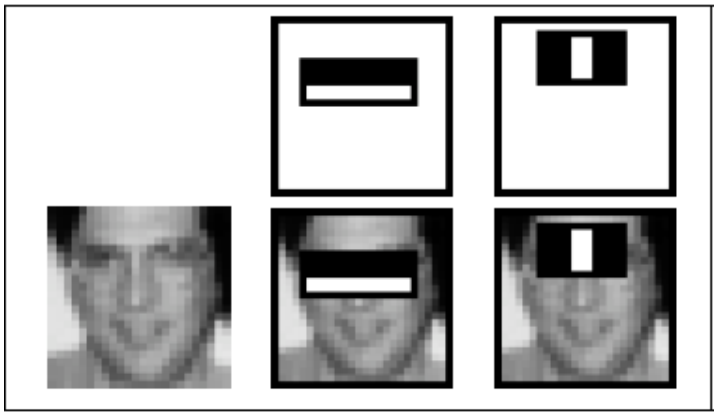

[Viola & Jones, 02]

$f(x|y)$

Correspondence
Occlusion
Incremental learning
Intuitive

$y = f(x)$

Performance
Fewer parameters
Fewer training ex.

Generative-discriminative hybrid

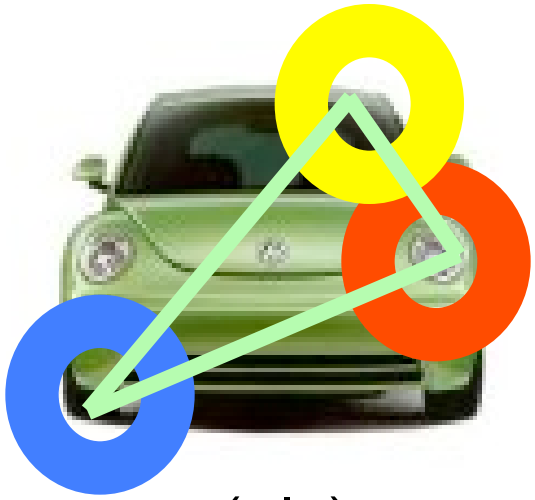

p(x|y)

F1
F2
F3
F4
F5
…

Fisher Scores*

For Image 'x'

Jaakkola and Hausler, 1999

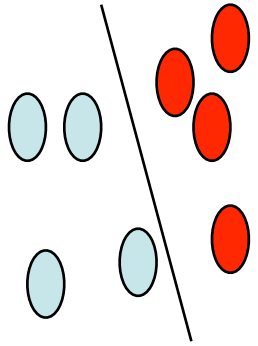

y=f(x)

$$\frac{\partial \log p(x_n|\theta_{\mathrm{ML}})}{\partial \theta}$$

SVM

[Holub, Welling & Perona - ICCV05]

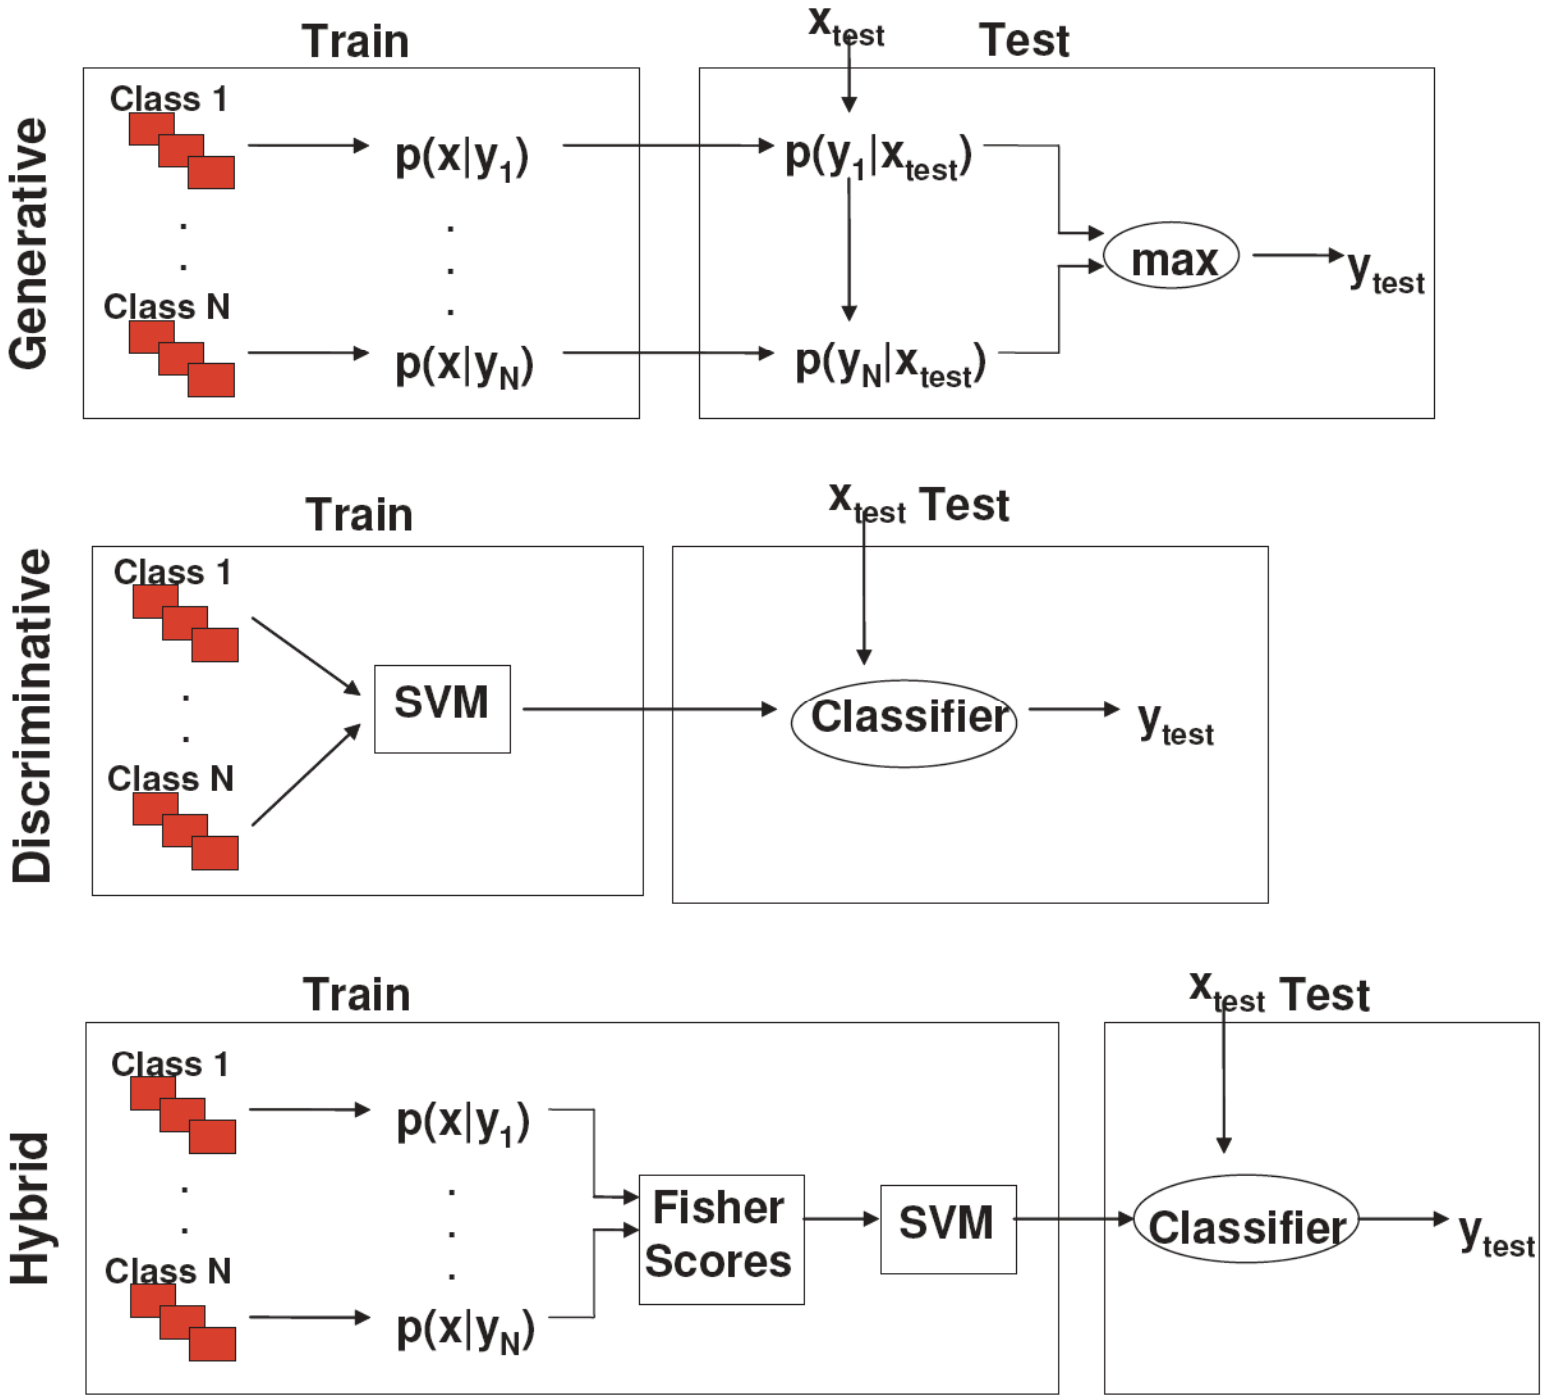
Generative
Train
Class 1
Class N
$p(x|y_1)$
$p(x|y_N)$
x_{test}
Test
$p(y_1|x_{test})$
$p(y_N|x_{test})$
max
y_{test}
Discriminative
Train
Class 1
Class N
SVM
x_{test} Test
Classifier
y_{test}
Hybrid
Train
Class 1
Class N
$p(x|y_1)$
$p(x|y_N)$
Fisher Scores
SVM
x_{test} Test
Classifier
y_{test}

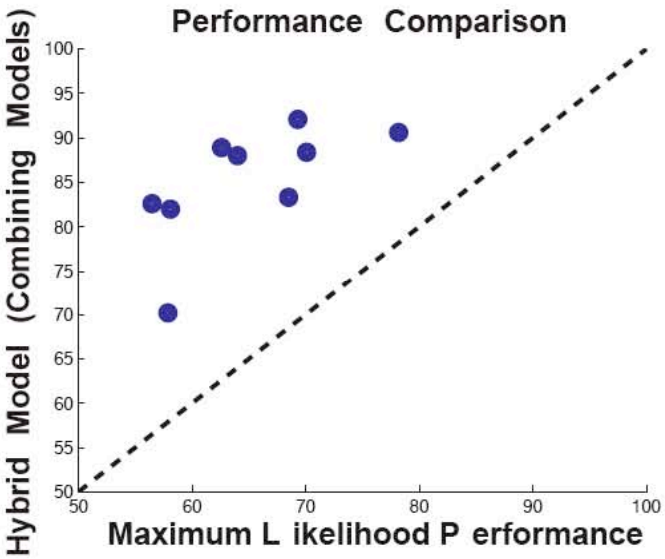

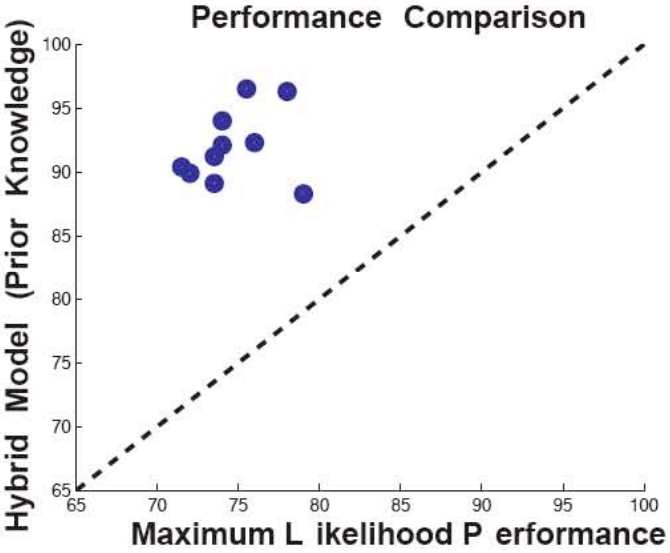

<5 Training Examples

Making use of unlabeled data

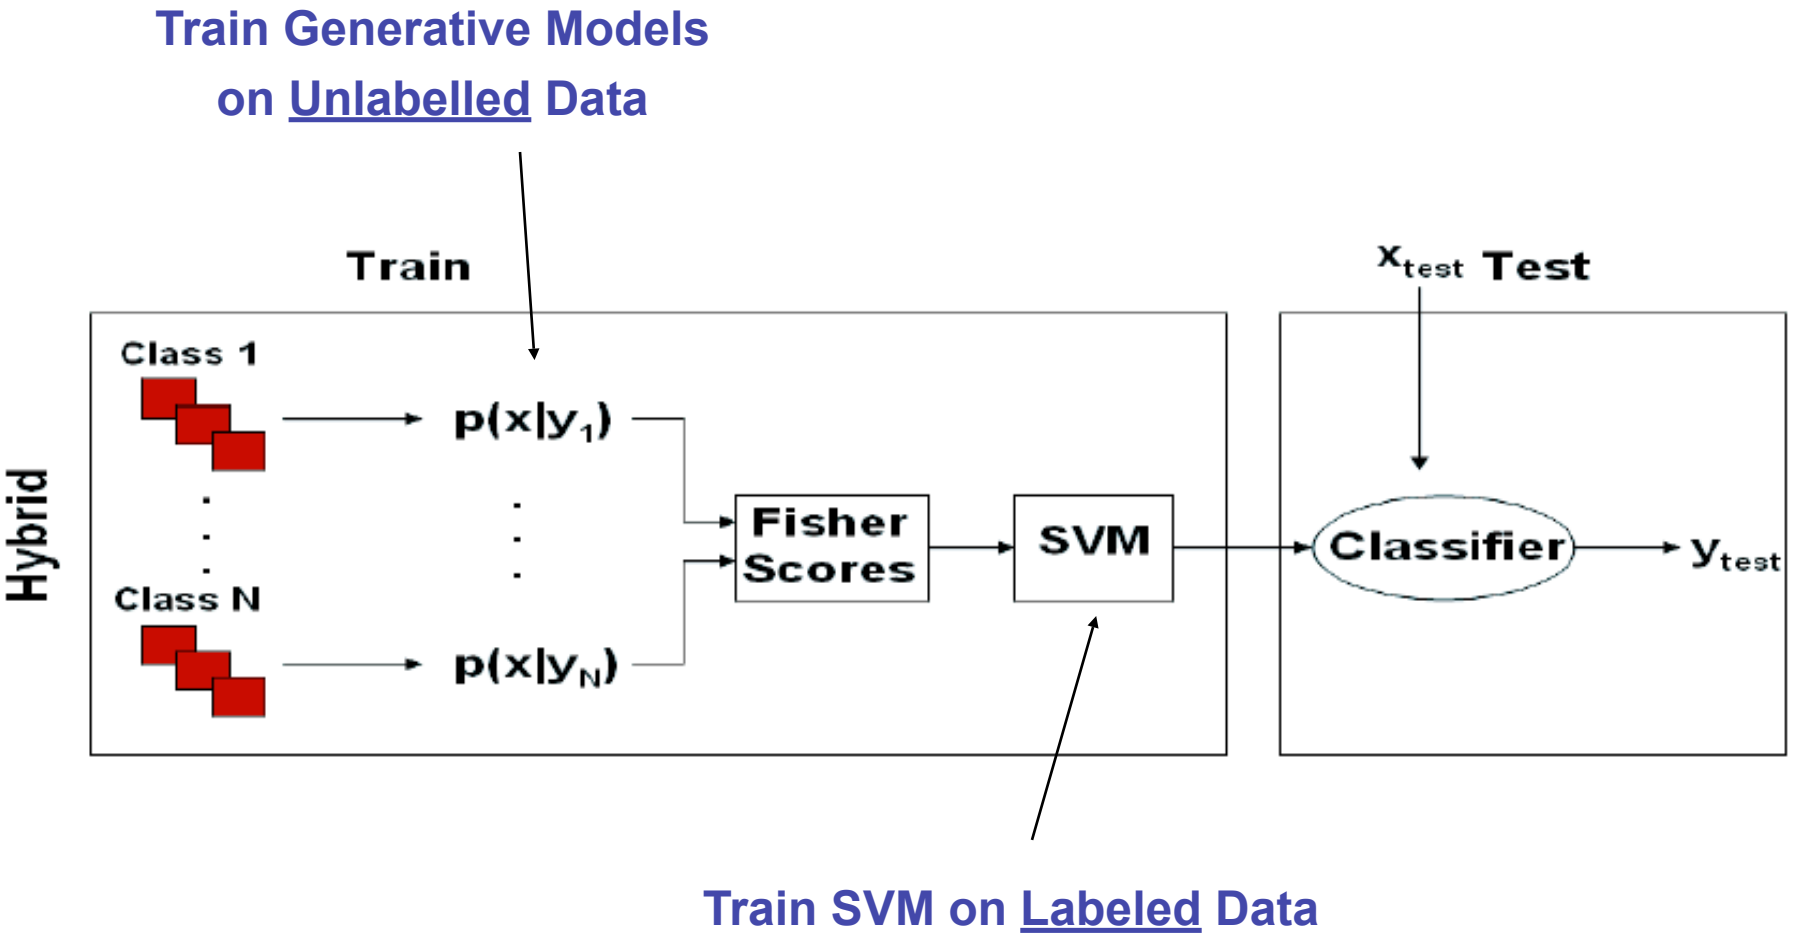

Where is the prior information?

Discrimination Task:

Vs

Possible Unlabelled Data-Sets for Prior:

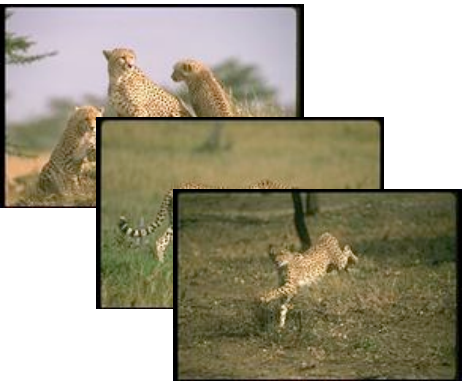

Leopards

Unlabeled faces

Background

The following instructions show how to make an exact copy of a CD or DVD. You can also use Sonic RecordNow for other purposes, including creating CDs from audio files on your computer and creating MP3 CDs. For instructions, see the Sonic RecordNow documentation that came with your computer. Open Sonic RecordNow, click the question mark icon in the upper-right corner of the window, and then click RecordNow Help or RecordNow Tutorial.

Text

Detected Features

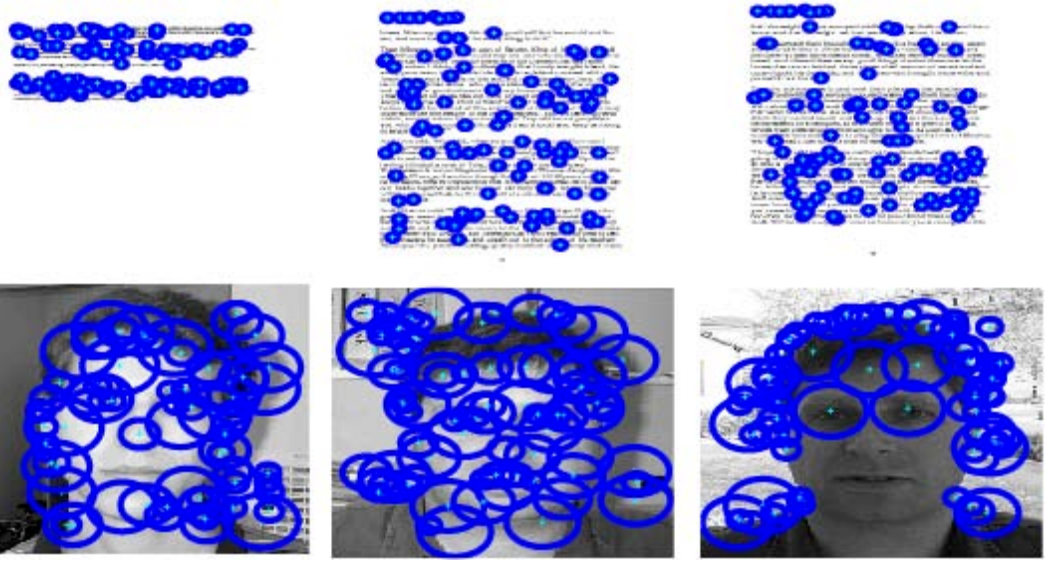

3 Training Examples

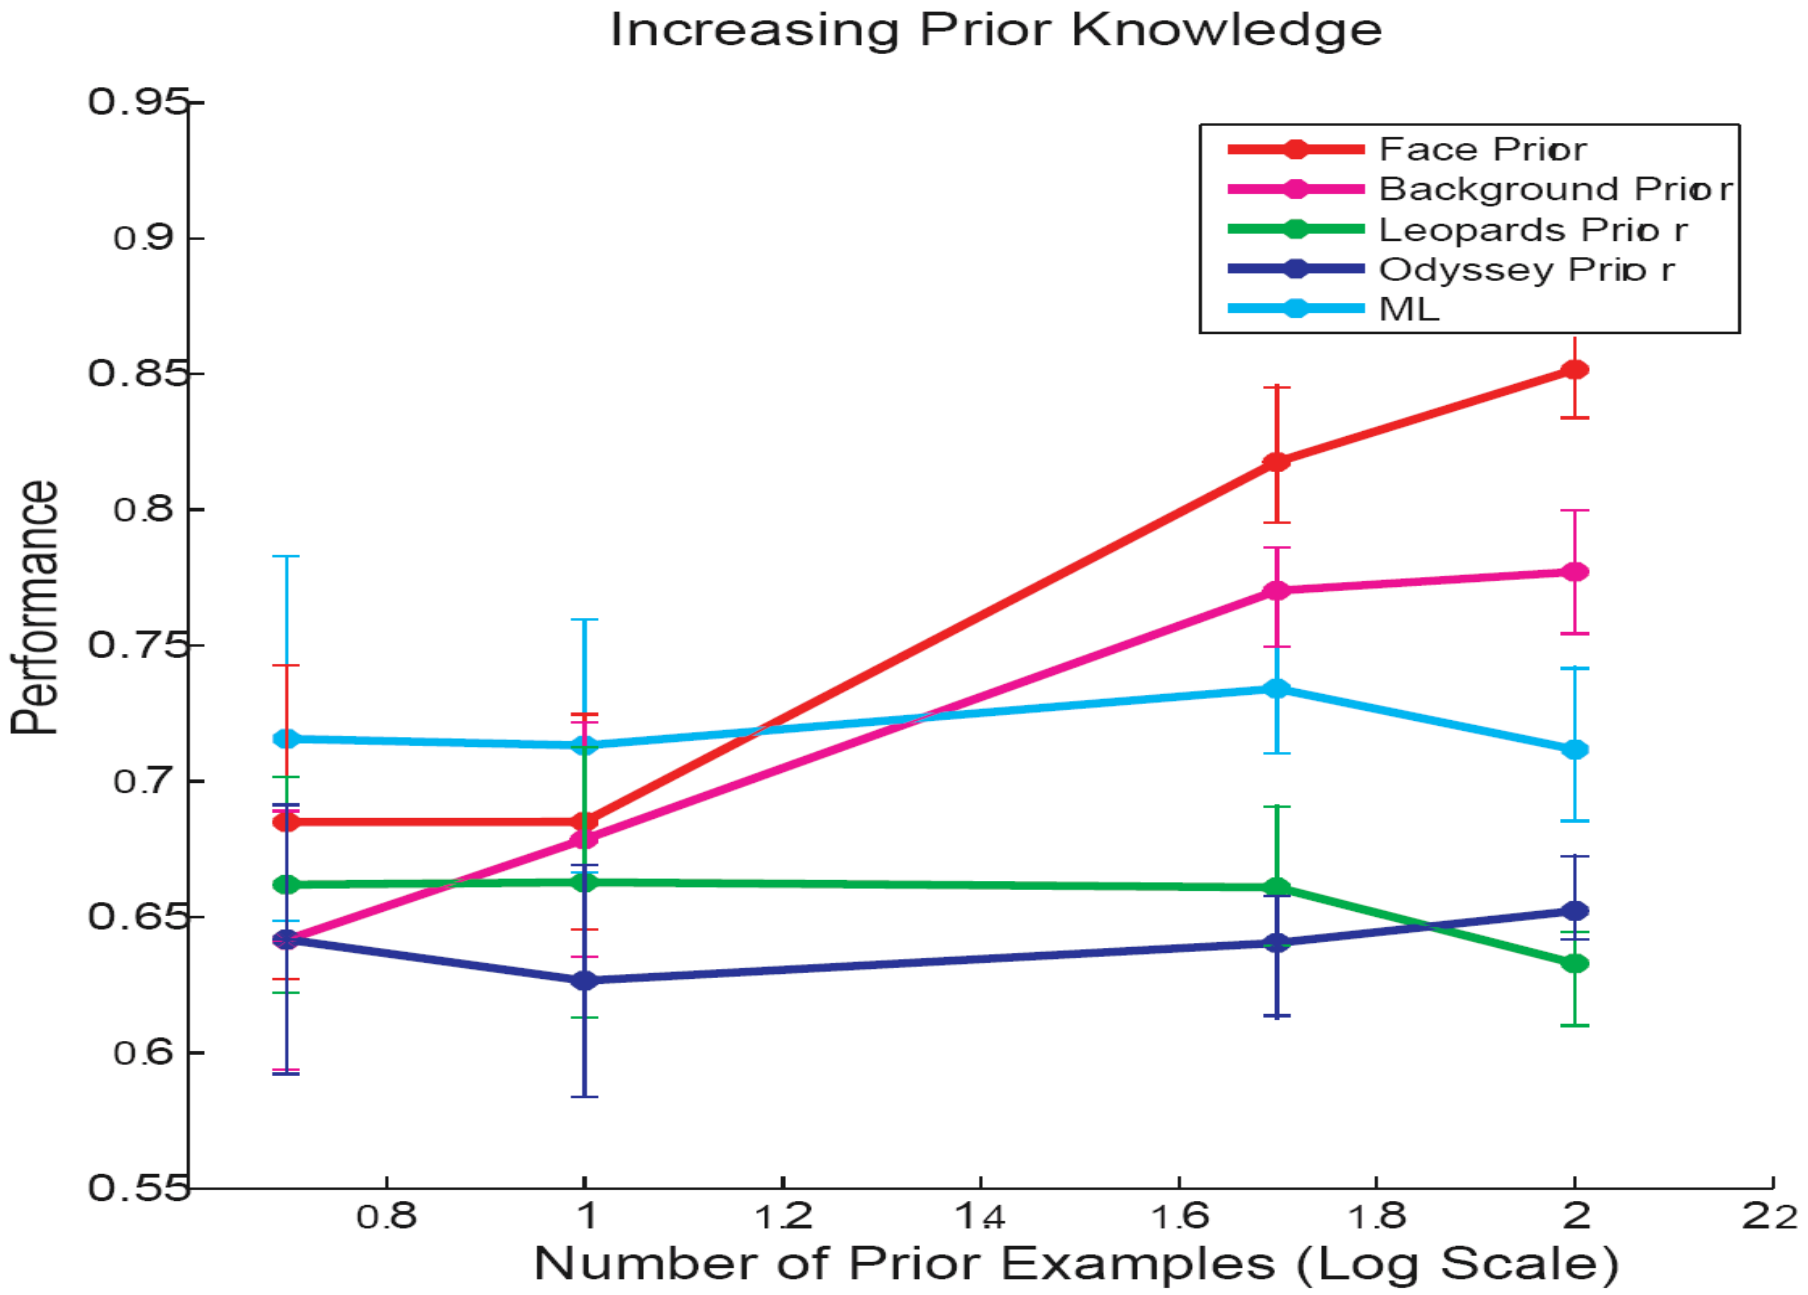

2-Class Discrimination Task

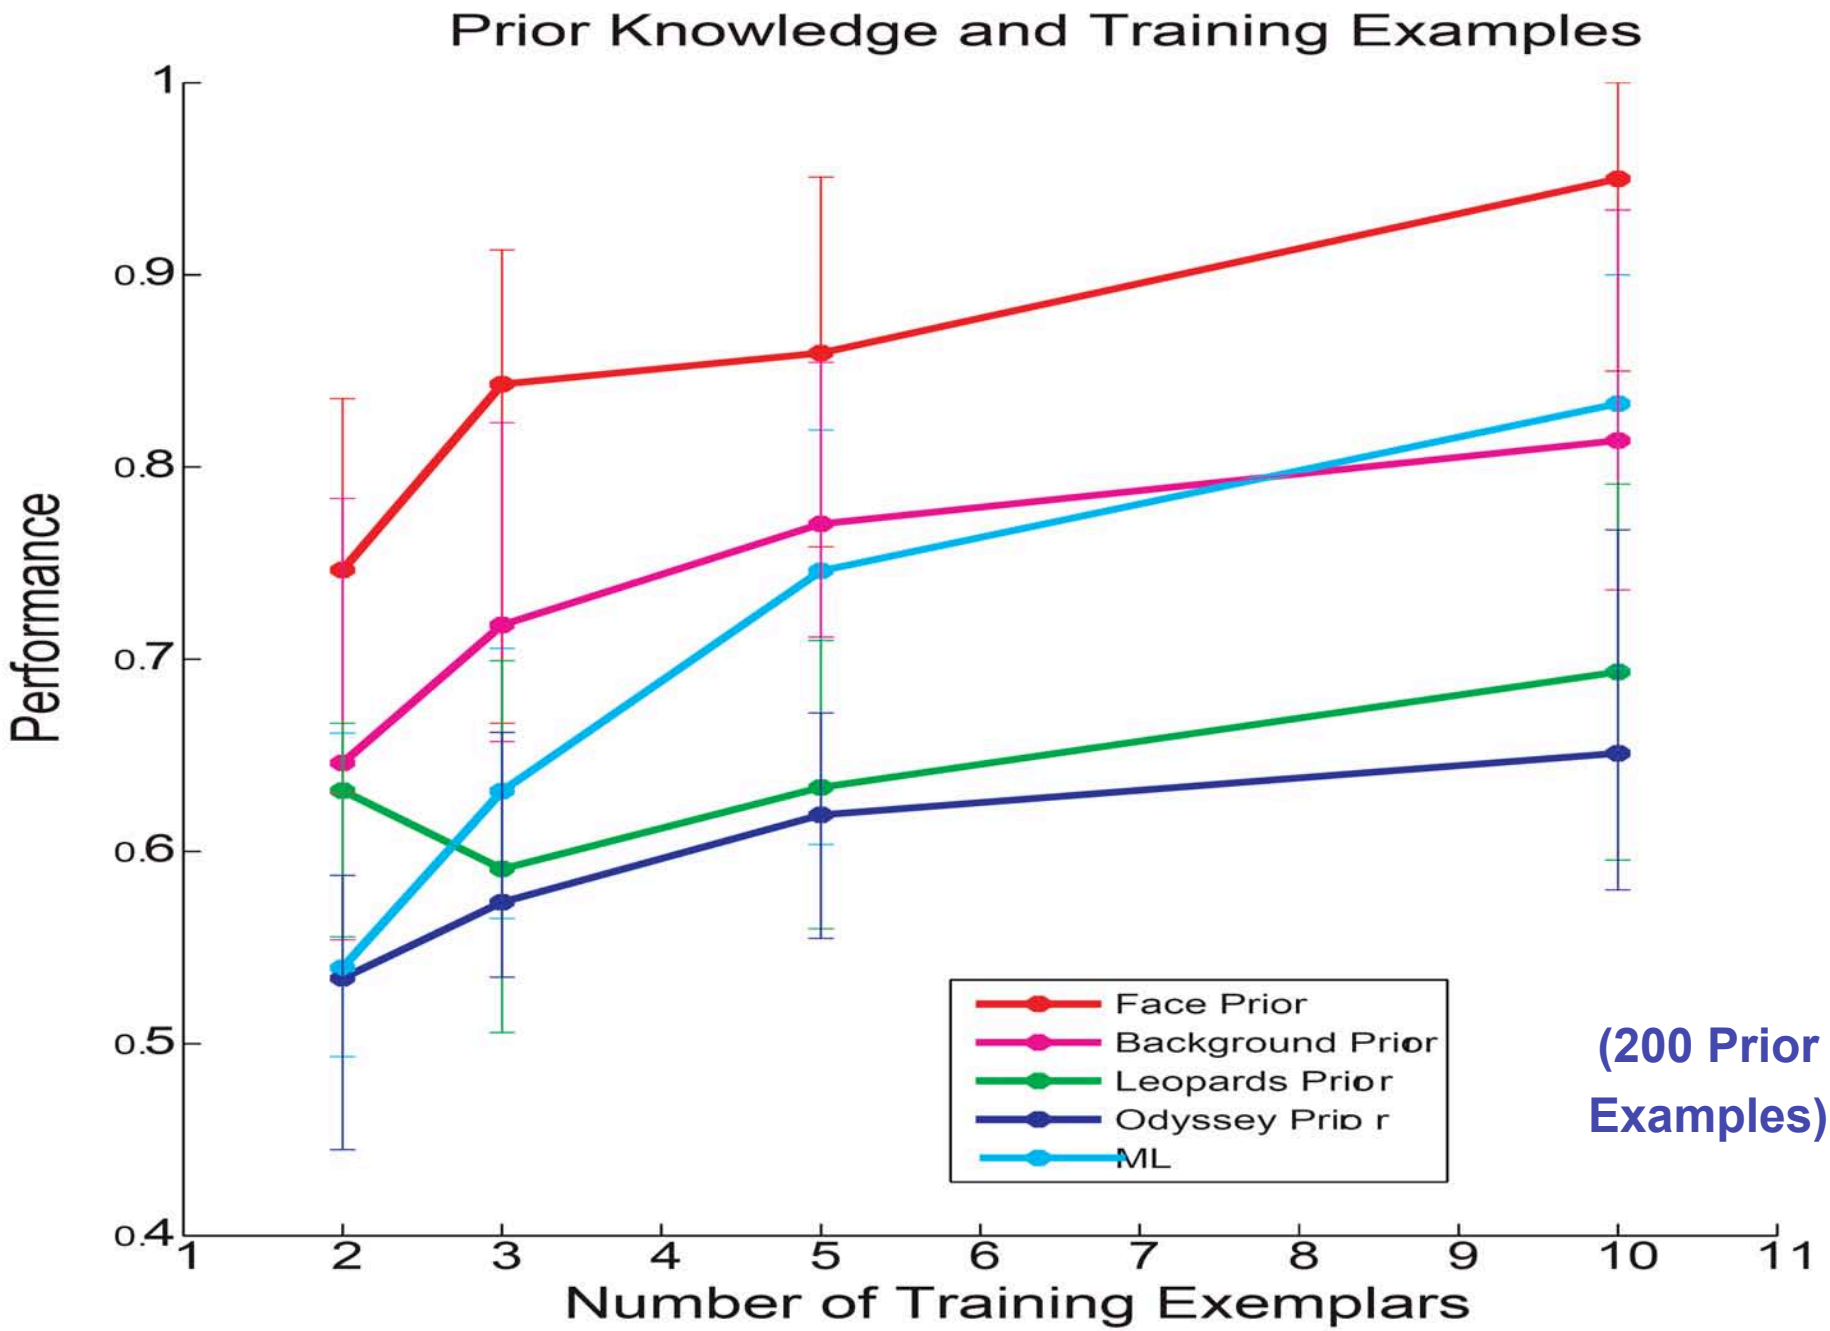

(200 Prior Examples)

Caltech101

- Train a general kernel on all training images.
- Extract fisher scores for each class from this general kernel.
- One vs. One classification on caltech101.

Caltech 101 Categories






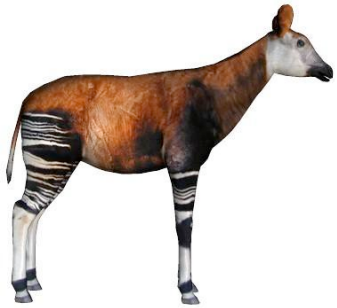

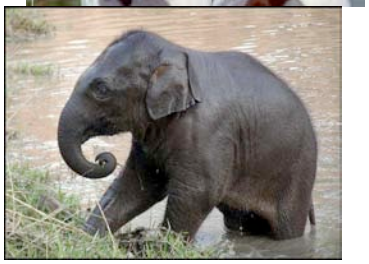






















Average Picture of Caltech 101

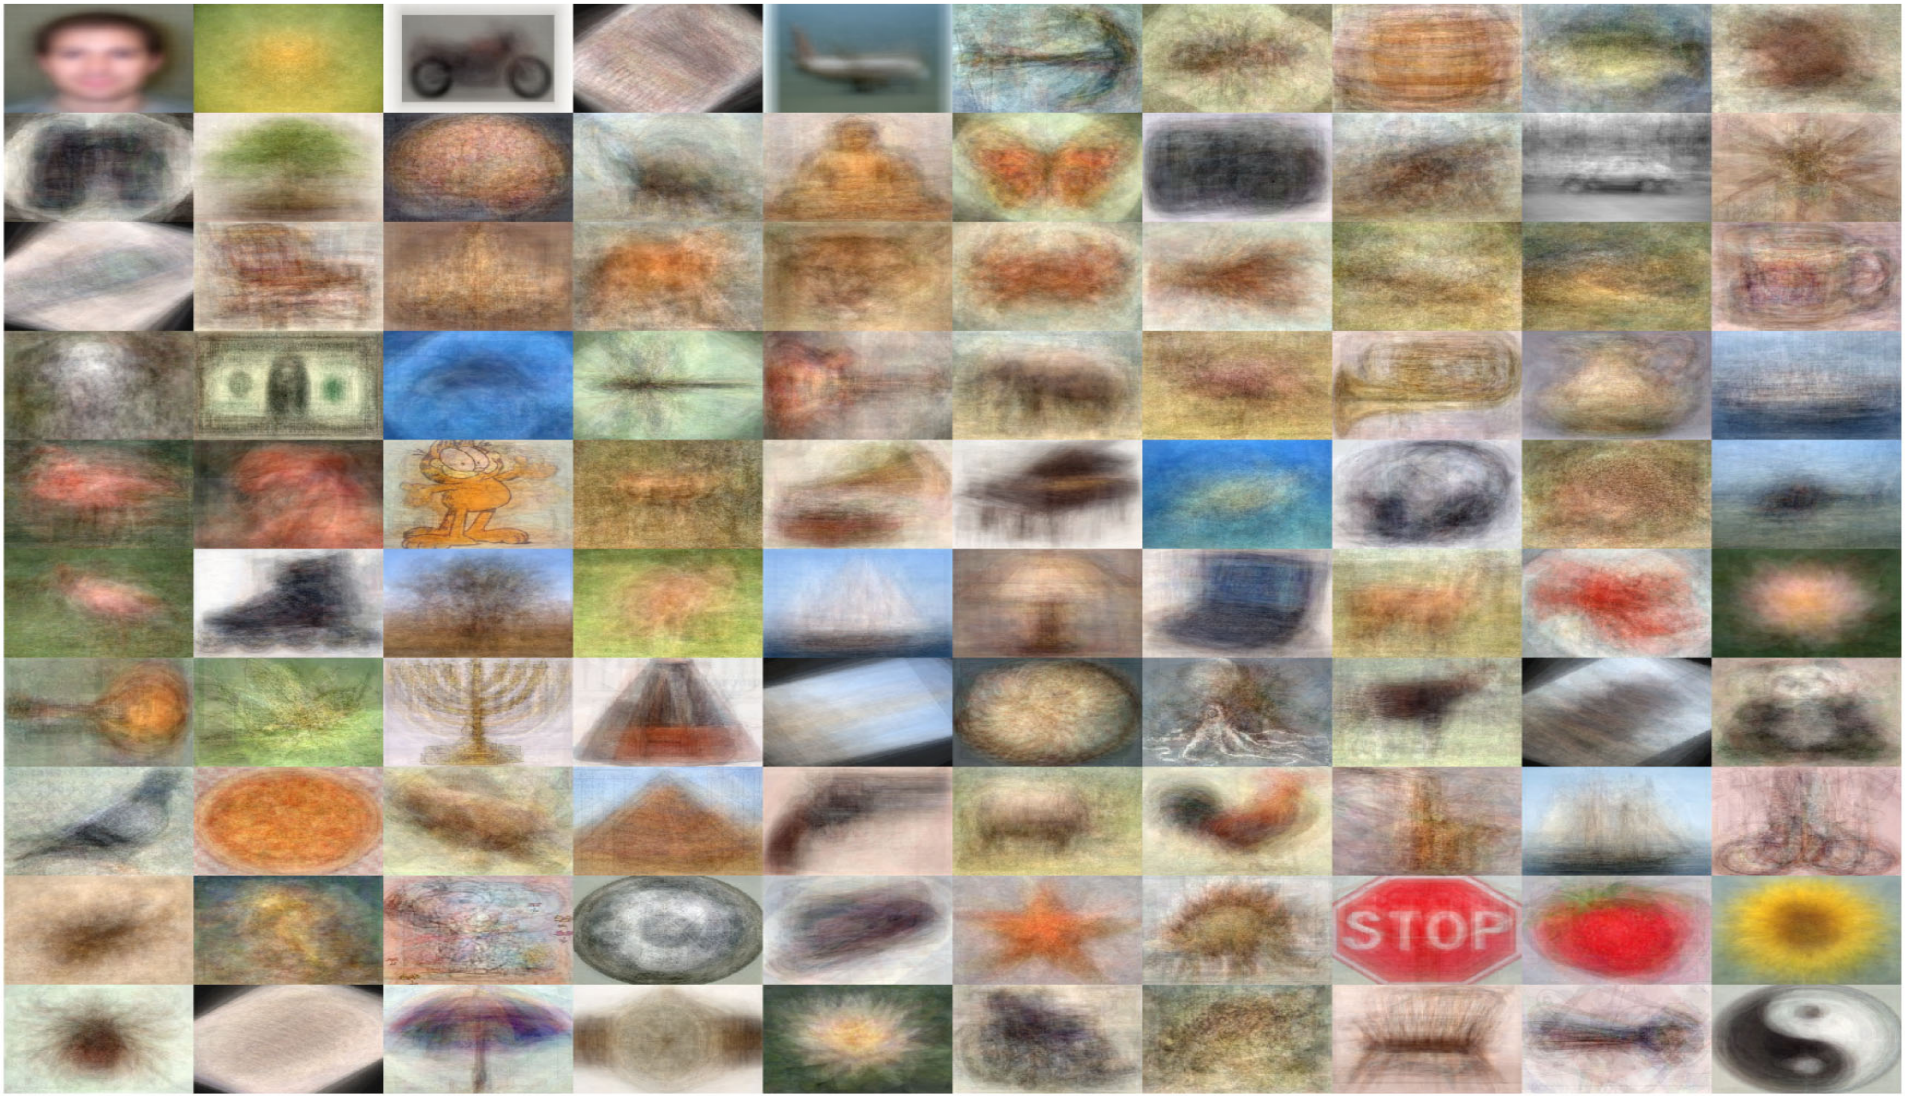

Antonio Torralba's Image of Average Caltech 101

Classification Results – Caltech 101

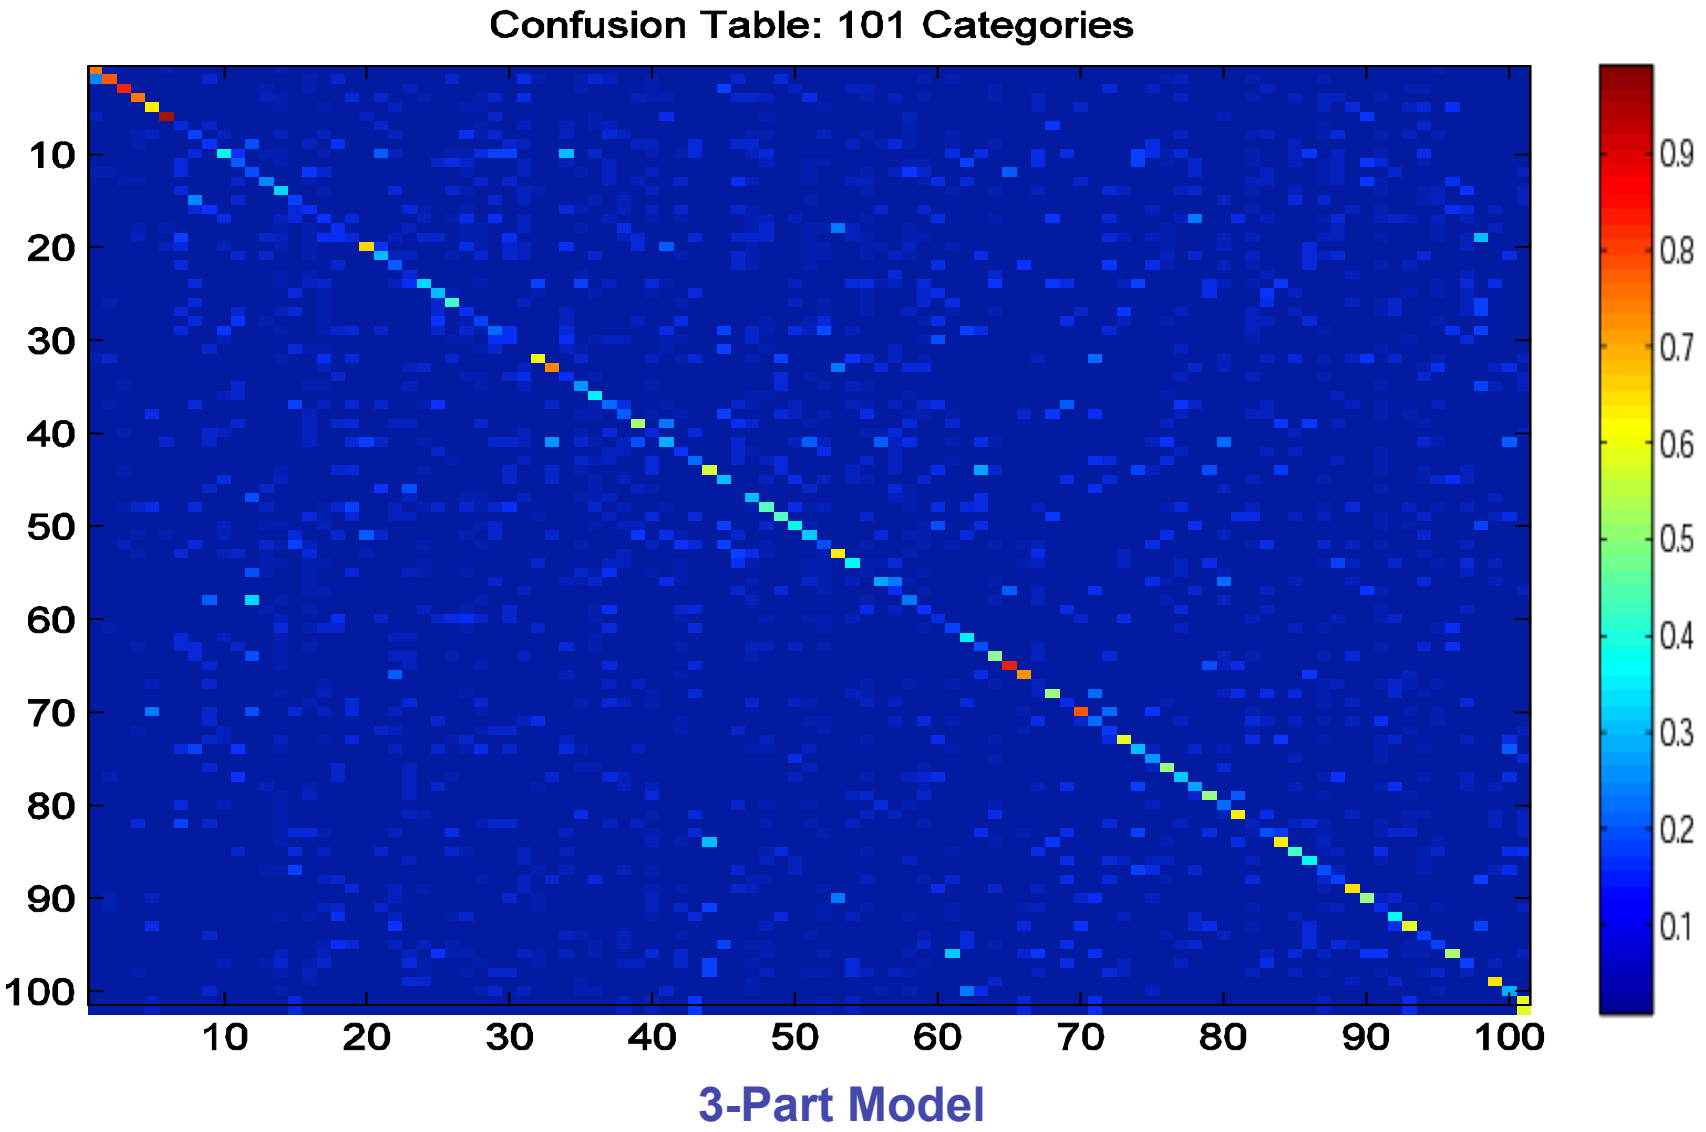

3-Part Model

(3 Detectors Used), avg. performance: 35%.

Appearance vs. Shape

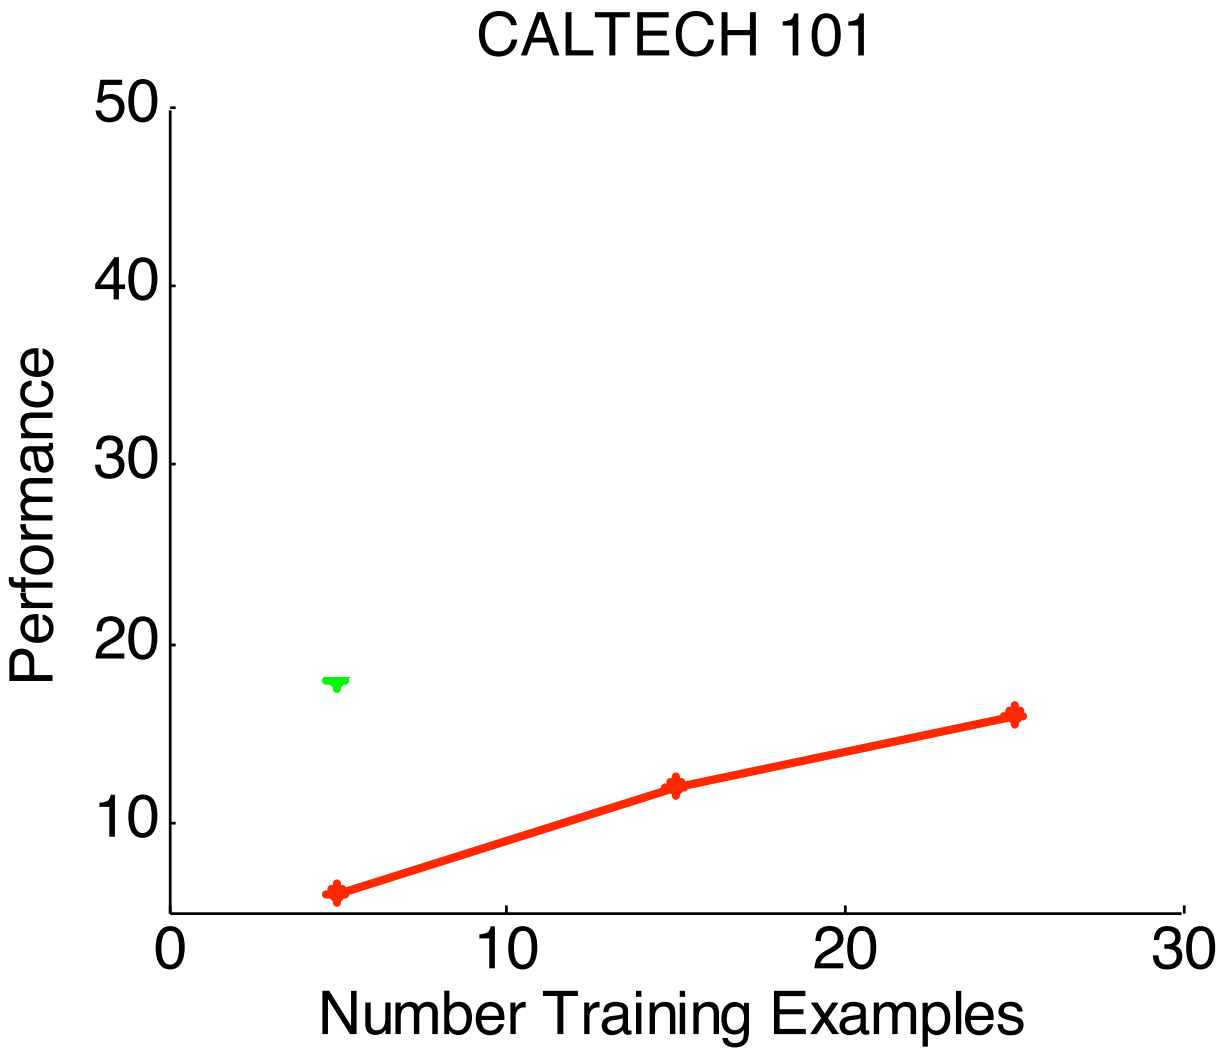

Caltech101 Performance

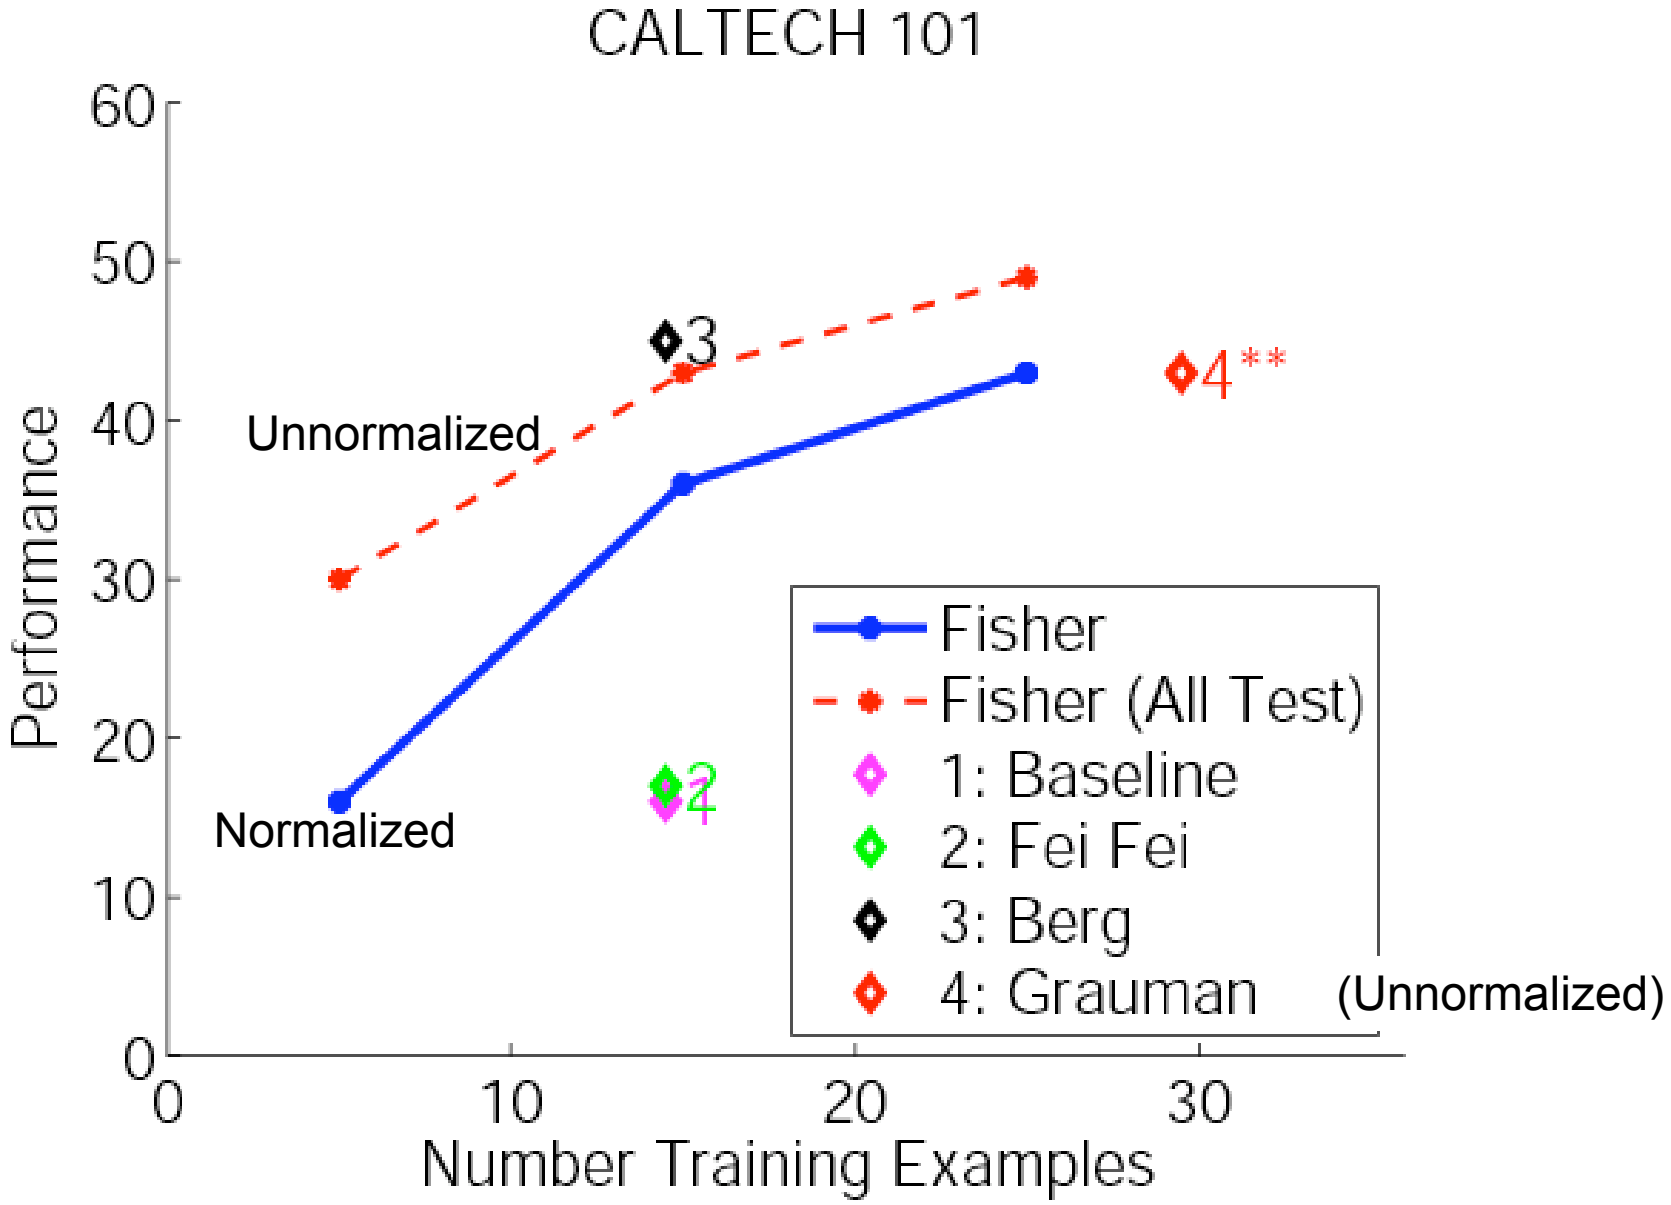

Conclusions

- Inter-class transfer is very useful
- Many mechanisms possible
- Generative-discriminative hybrid allows use of unlabeled data