The Open Set Recognition Problem

Walter J. Scheirer
School of Engineering and Applied Sciences,
Department of Cellular and Molecular Biology, and
Center for Brain Science, Harvard University

Benchmarks in computer vision

Assume we have examples from all classes:

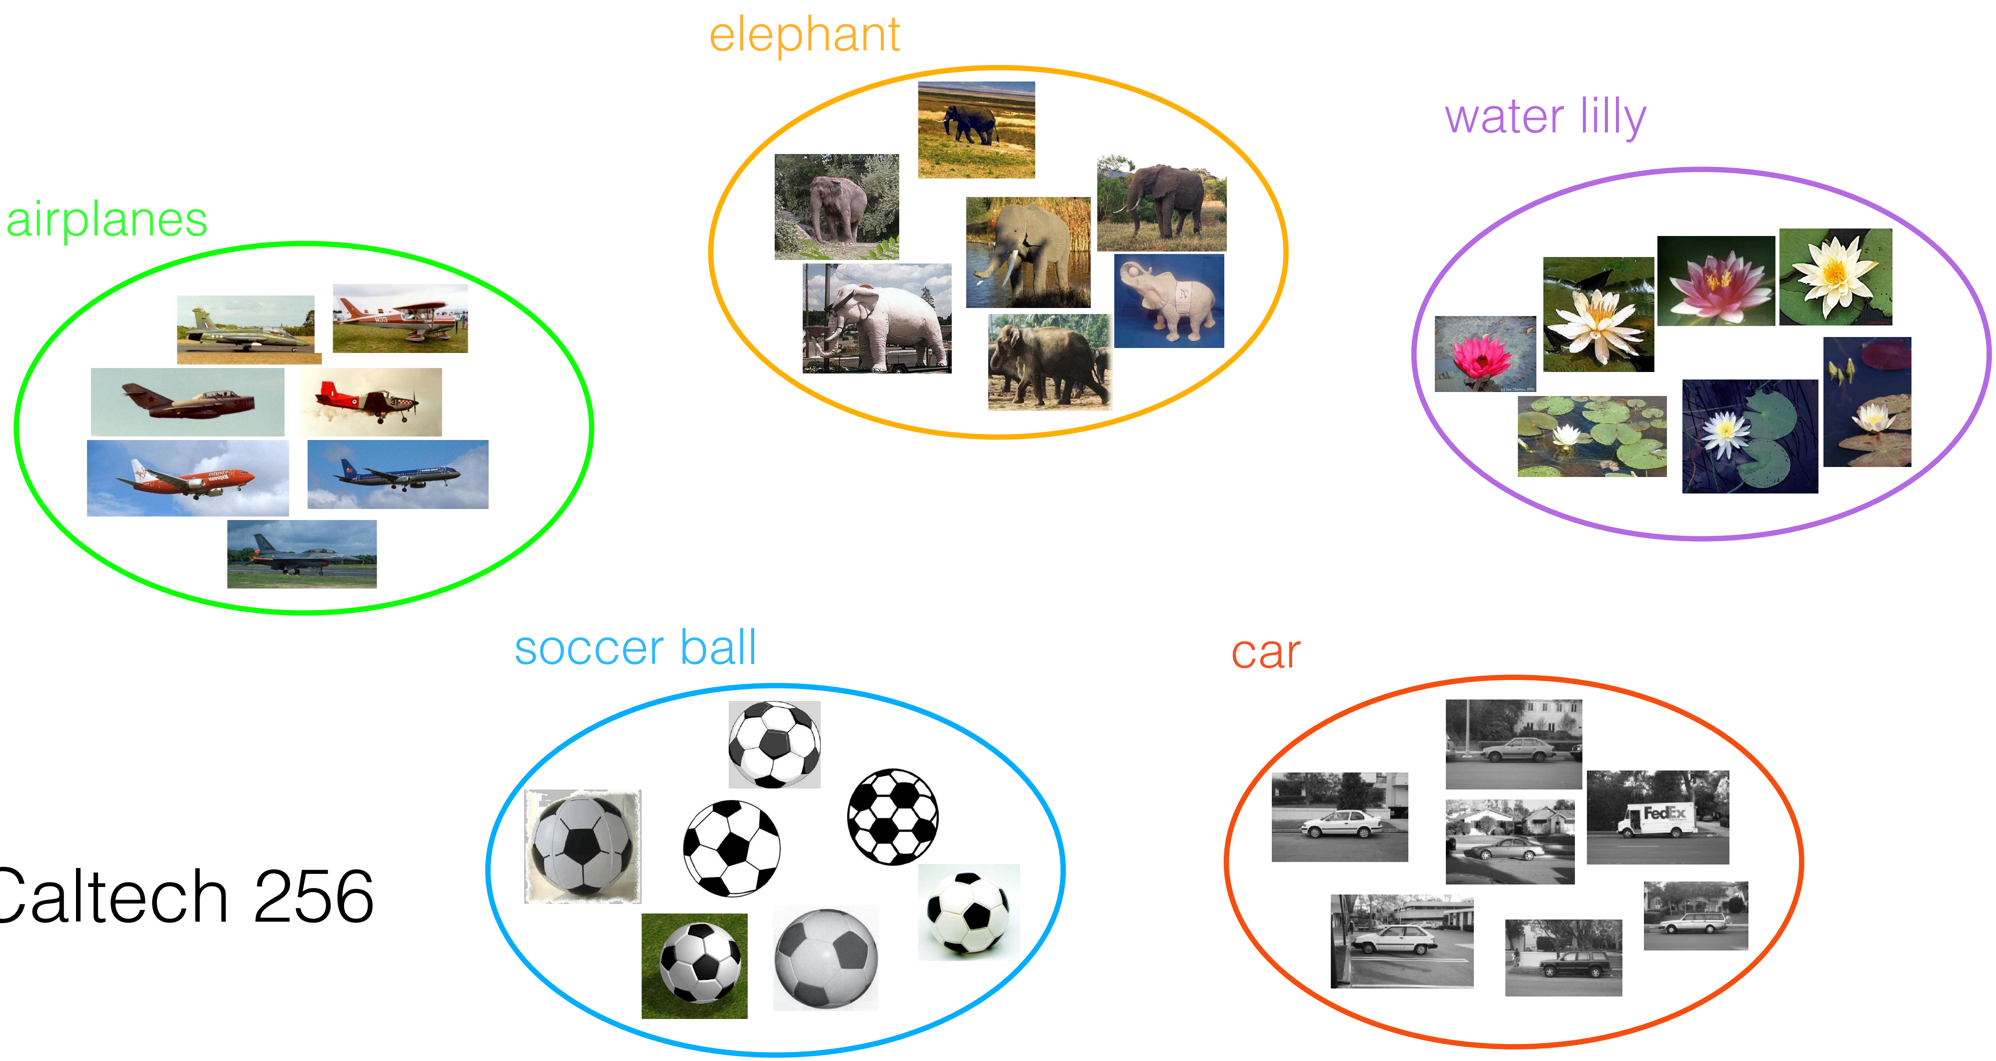

Caltech 256

Out in the real world…

Detect the cars in this image

while rejecting the trees, signs, telephone poles…

M. Milford, E. Vig, W.J. Scheirer, D.D. Cox, "Condition Invariant Top-Down Visual Place Recognition," ICRA Submission 2013.

"All positive examples are alike; each negative example is negative in its own way"

Zhao and Huang (with some help from Tolstoy)
CVPR 2001

X. Zhou and T. Huang, "Small Sample Learning during Multimedia Retrieval using BiasMap," in IEEE CVPR, 2001.

What is the general object recognition problem?

- Duin and Pekalska*: how one should approach multi-class recognition is still an open issue
  - Is it a series of binary classifications?
  - Is it a search performed for each possible class?
  - What happens when some classes are ill-sampled, not sampled at all or undefined?

R. P. Duin and E. Pekalska, "Open Issues in Pattern Recognition," in *Computer Recognition Systems*, M. Kurzynski, E. Puchala, M. Wozniak, and A. Zolnierek, Eds. Springer, 2005, pp. 27–42.

Vision problems in order of "openness"

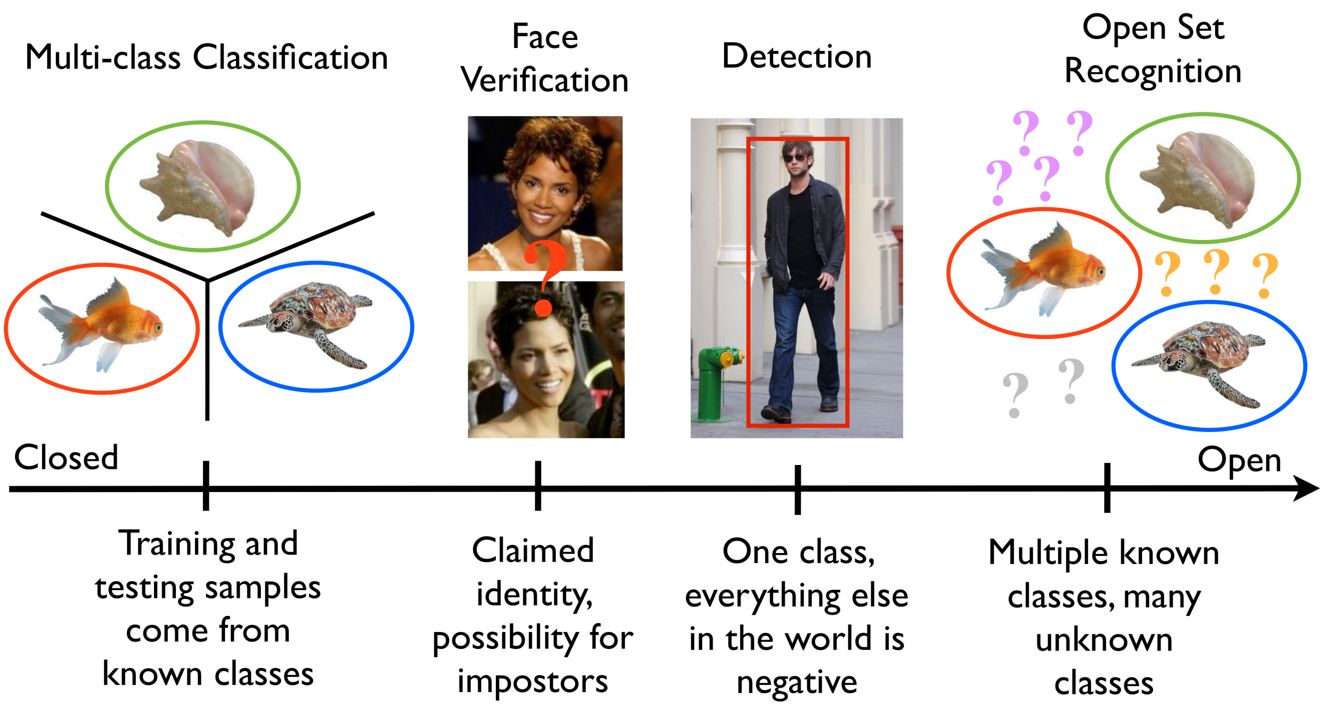

W.J. Scheirer, A. Rocha, A. Sapkota, and T. Boult, "Towards Open Set Recognition," IEEE T-PAMI, 35(7) July 2013.

Let's formalize openness

$$\text{openness} = 1 - \sqrt{\frac{2 \times |\text{training classes}|}{|\text{testing classes}| + |\text{target classes}|}}$$

Examples of openness values

	Targets	Training	Testing	Openness
Typical Multi-class	x	x	x	0%
Face Verification	12	12	50	38%
Typical Detection	1	100,000	1,000,000	55%
Object Recognition	88	12	88	63%
Object Recognition	88	6	88	74%
Object Recognition	212	6	212	83%

Fundamental multi-class recognition problem

$$\operatorname*{argmin}_{f} \left\{ R_{\mathcal{I}}(f) := \int_{\mathbb{R}^d \times \mathbb{N}} L(x, y, f(x)) P(x, y) \right\}$$

Ideal Risk   Loss Function   Joint Distribution

Undefined for
open set recognition!

A. Smola, “Learning with Kernels,” Ph.D. dissertation, Technische Universität Berlin, Berlin, Germany, November 1998.

Open Space

Open Space

- Open space is the space far from known data
- We need to address the infinite half-space problem of linear classifiers
- Principle of Indifference*
  - If there is no known reason to assign probability, alternatives should be given equal probability
  - One problem: we need the distribution to integrate to 1!

J.M. Keynes, A Treatise on Probability. Macmillan & Company, Limited, 1921.

Open Space Risk

Open Space Risk: the relative measure of open space to the full space

open space

$$R_{\mathcal{O}}(f) = \frac{\int_{\mathcal{O}} f(x)dx}{\int_{S_o} f(x)dx}$$

Open space + positive training examples

The open set recognition problem

Preliminaries

Space of positive class data: $\mathcal{P}$

Space of other known class data: $\mathcal{K}$

Positive training data: $\hat{V} = \{v_1, \ldots, v_m\}$ from $\mathcal{P}$

Negative training data: $\hat{K} = \{k_1, \ldots, k_n\}$ from $\mathcal{K}$

Unknown negatives appearing in testing: $\mathcal{U}$

Testing data: $\mathcal{T} = \{t_1, \ldots, t_z\}$, $t_i \in \mathcal{P} \cup \mathcal{K} \cup \mathcal{U}$

Assume the problem openness is > 0

The open set recognition problem

Minimize open set risk:

$$\underset{f \in \mathcal{H}}{\operatorname{argmin}} \left\{ R_{\mathcal{O}}(f) + \lambda_r R_{\mathcal{E}}(f(\hat{V} \cup \hat{K})) \right\}$$

Open Space Risk Associated with $\mathcal{U}$

Regularization Constant

Empirical Risk Function

What options do we have to solve this problem?

Binary Classification

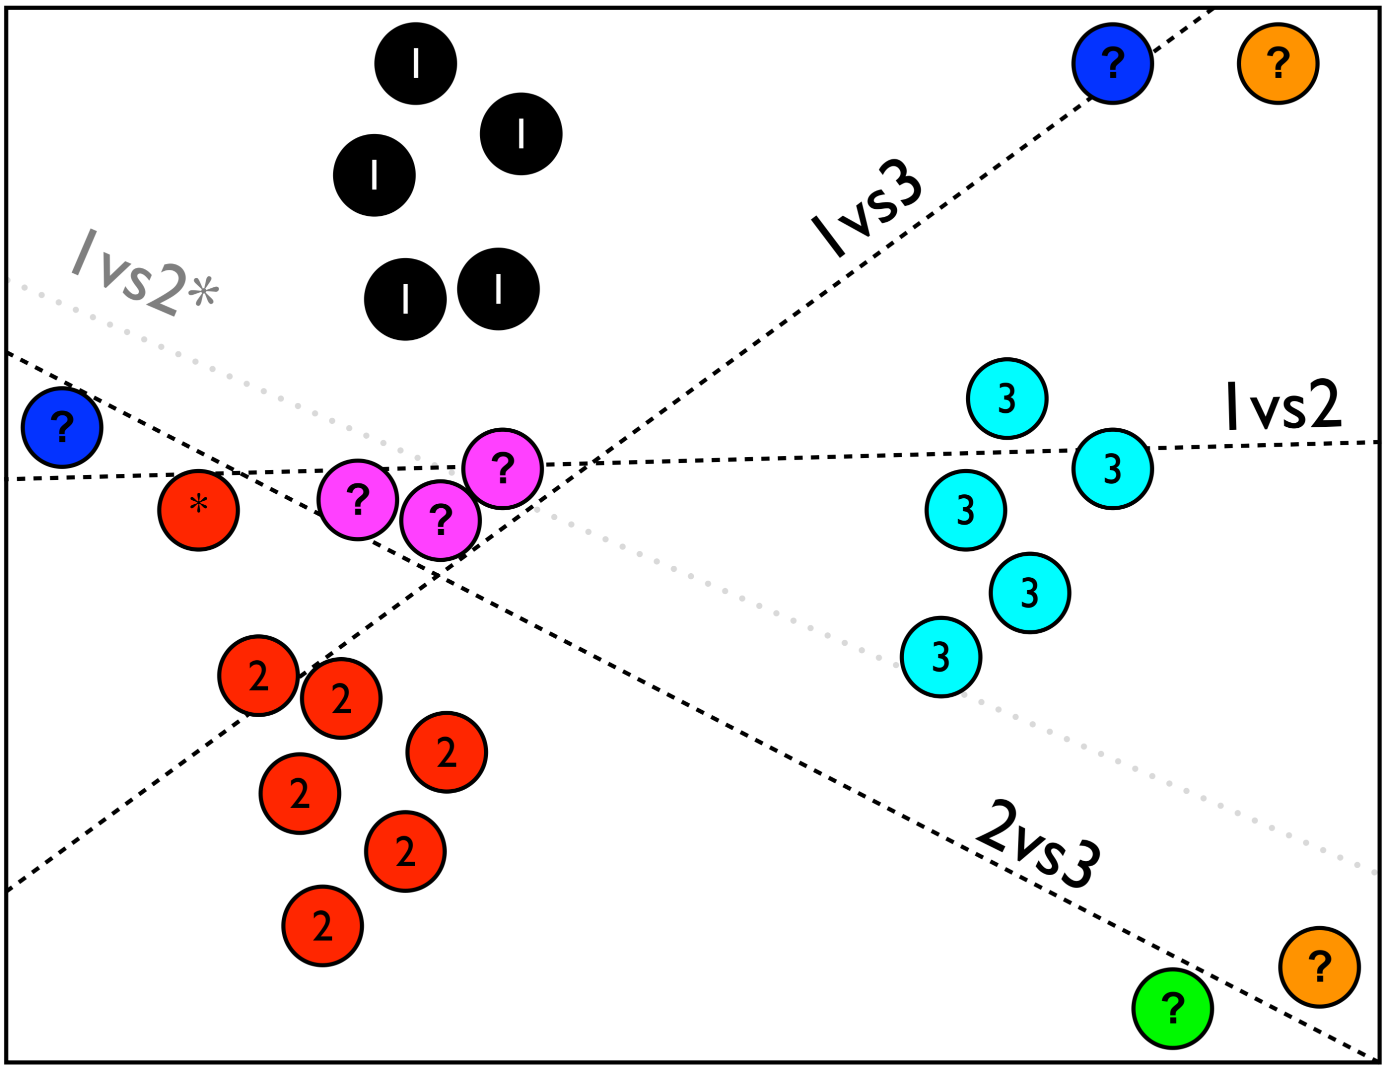

Multi-class 1-vs-All Classification

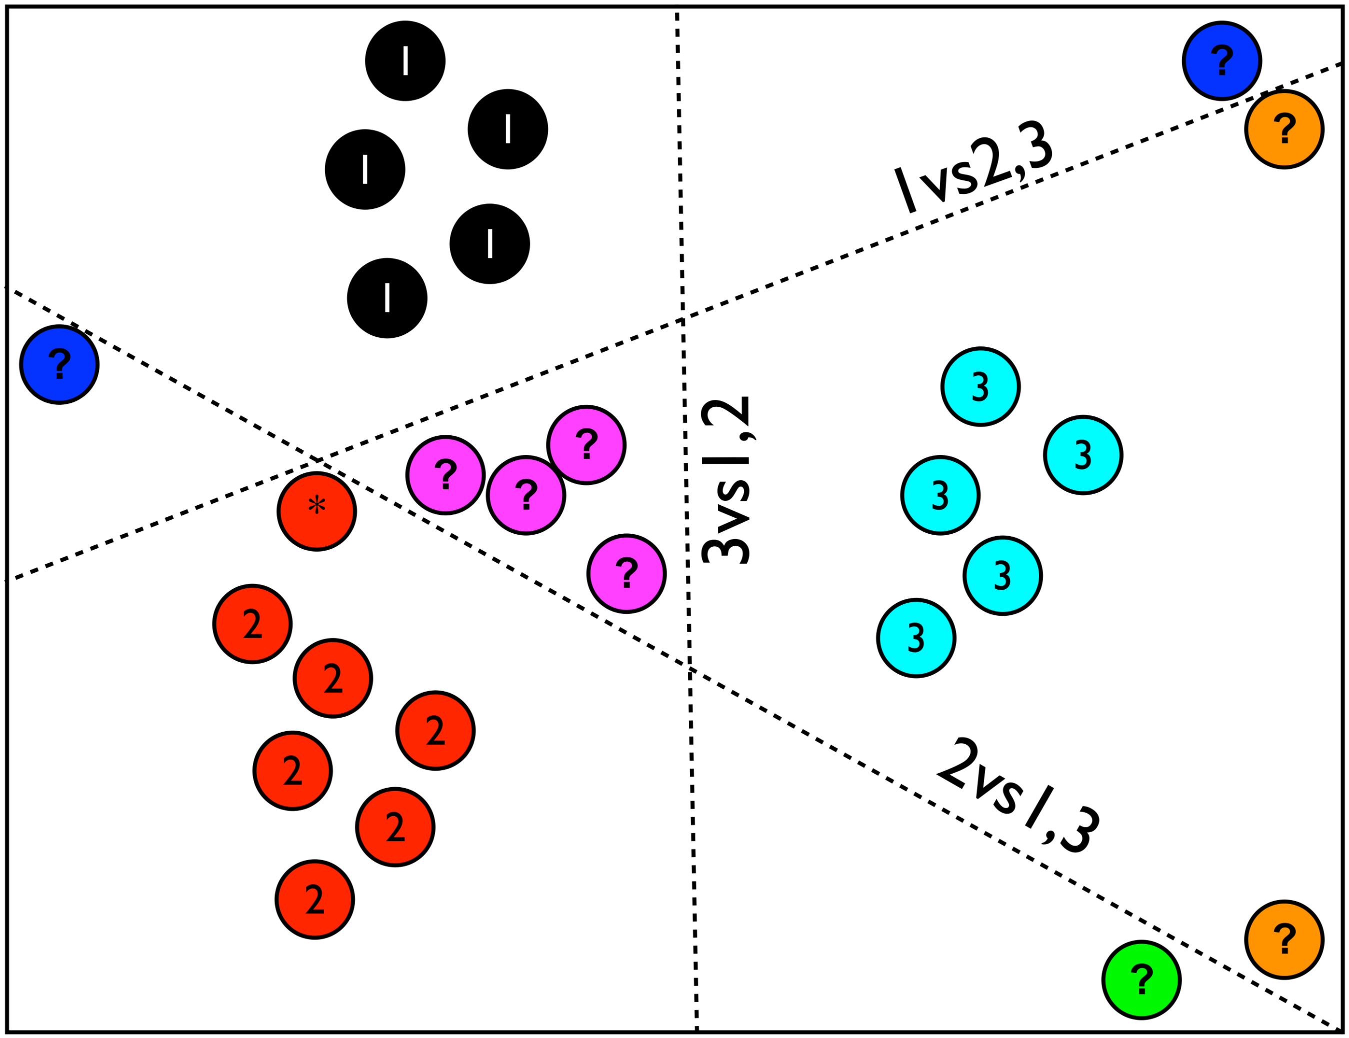

1-class Classification

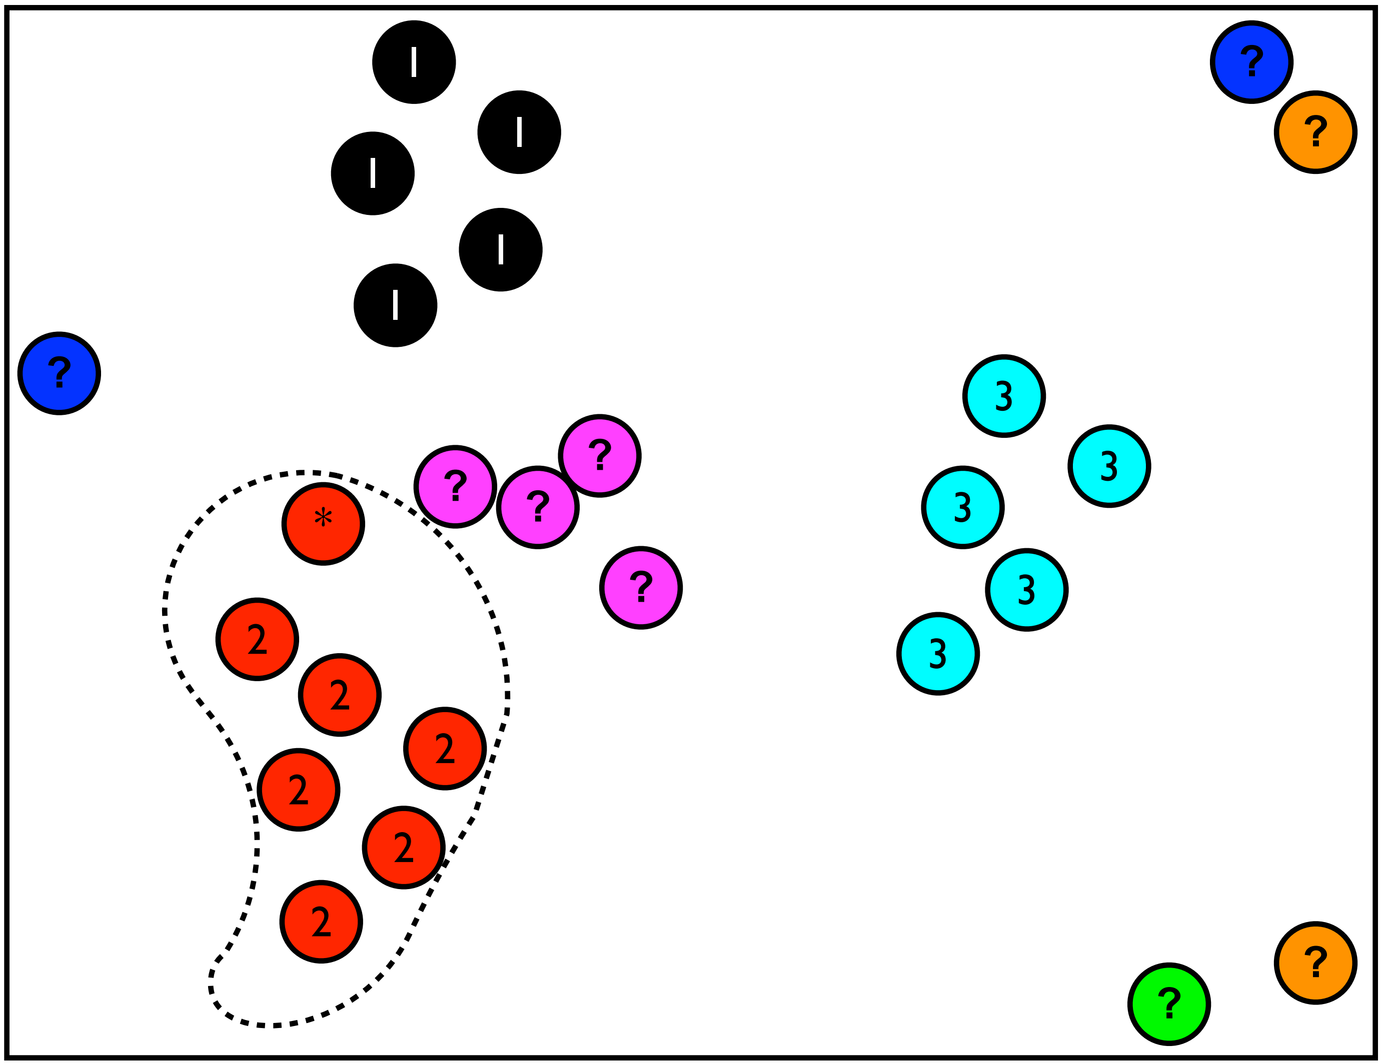

B. Schölkopf, J. Platt, J. Shawe-Taylor, A. Smola, and "R. Williamson. Estimating the Support of a High-dimensional Distribution," Technical report, Microsoft Research, 1999.

Why didn't the 1-class SVM catch on?

- Zhou and Huang *Multimedia Systems* 2003
  - Kernel and parameter selection
    - Gaussian kernels lead to over-fitting
    - Parameters chosen in *ad hoc* fashion
    - An issue in other domains too!

X. Zhou and T. Huang, "Relevance Feedback in Image Retrieval: A Comprehensive Review," *Multimedia Systems*, vol. 8, no. 6, pp. 536–544, 2003.

Other approaches

- M. Rohrbach, M. Stark, and B. Schiele, “Evaluating Knowledge Transfer and Zero-Shot Learning in a Large-Scale Setting,” in IEEE CVPR, 2011.
- C. H. Lampert, H. Nickisch, and S. Harmeling, “Learning To Detect Unseen Object Classes by Between-Class Attribute Transfer,” in IEEE CVPR, 2009.
- E. Bart and S. Ullman, “Single-example Learning of Novel Classes Using Representation by Similarity,” BMVC, 2005.
- M. Palatucci, D. Pomerleau, G. Hinton, and T.M. Mitchell, “Zero-shot Learning with Semantic Output Codes,” NIPS, 2009.
- L. Wolf, T. Hassner, and Y. Taigman, “The One-shot Similarity Kernel,” ICCV 2009.
- G. Heidemann, “Unsupervised Image Categorization,” Image and Vision Computing, vol. 23, no. 10, pp. 861–876, October 2004.

Let's include open space risk in our optimization problem

Slab Model

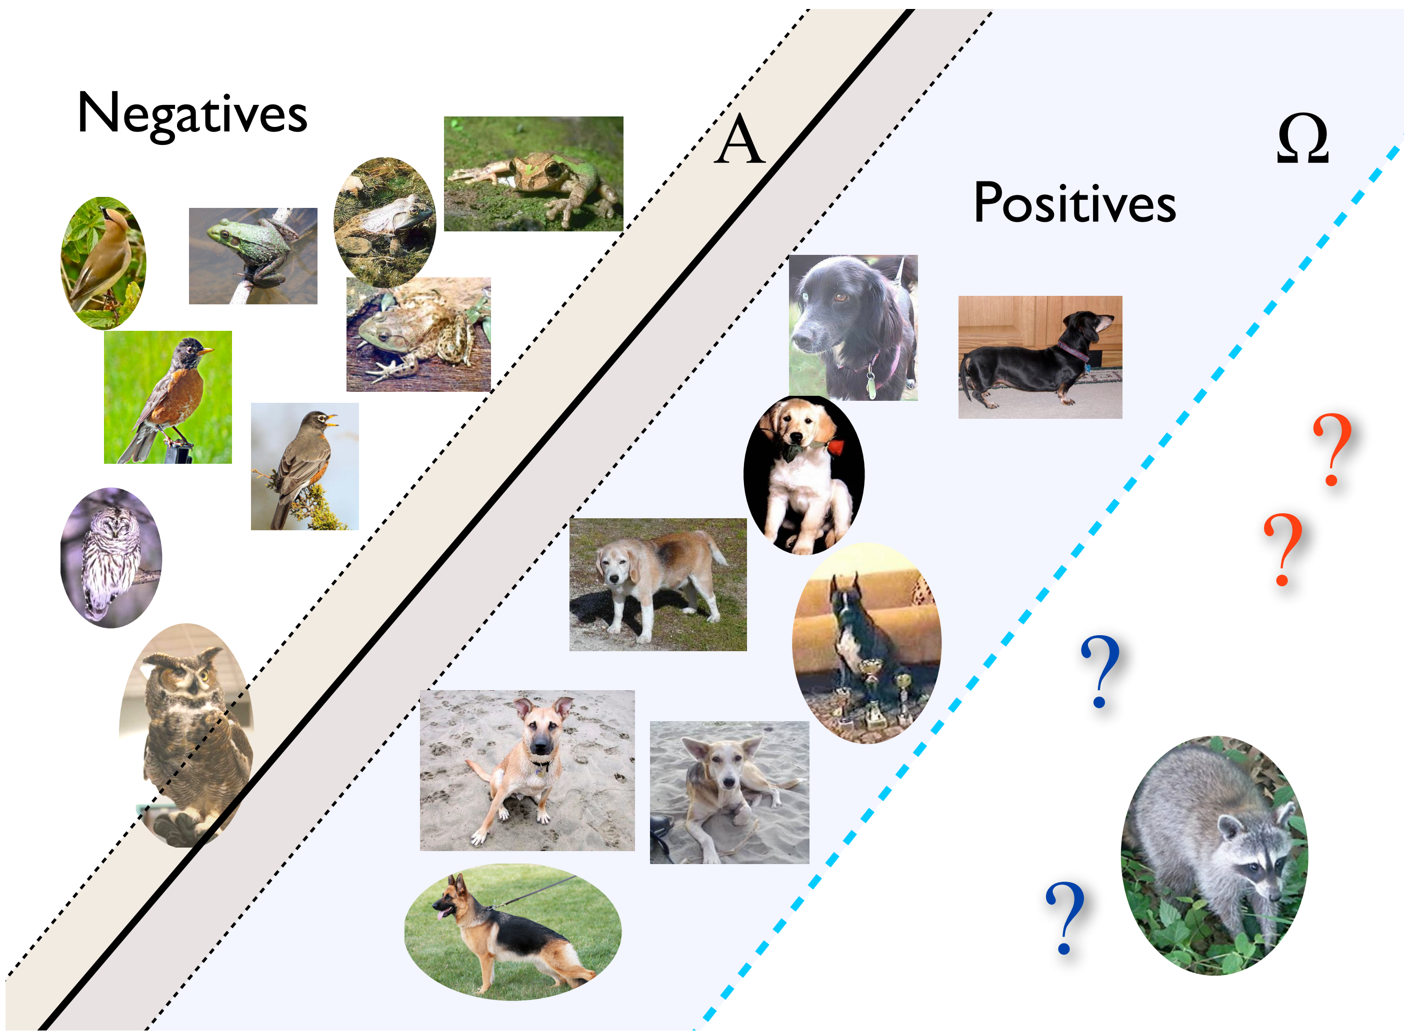

Base Linear 1-vs-Set Machine

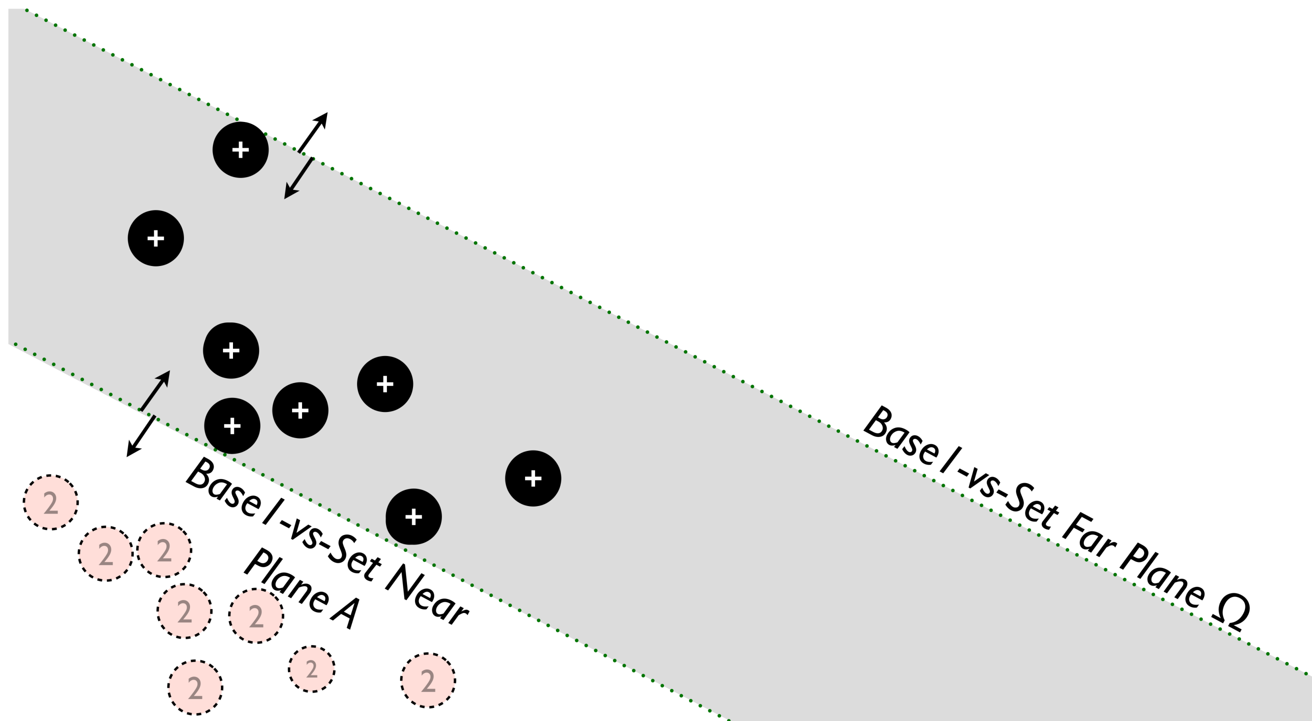

Generalization

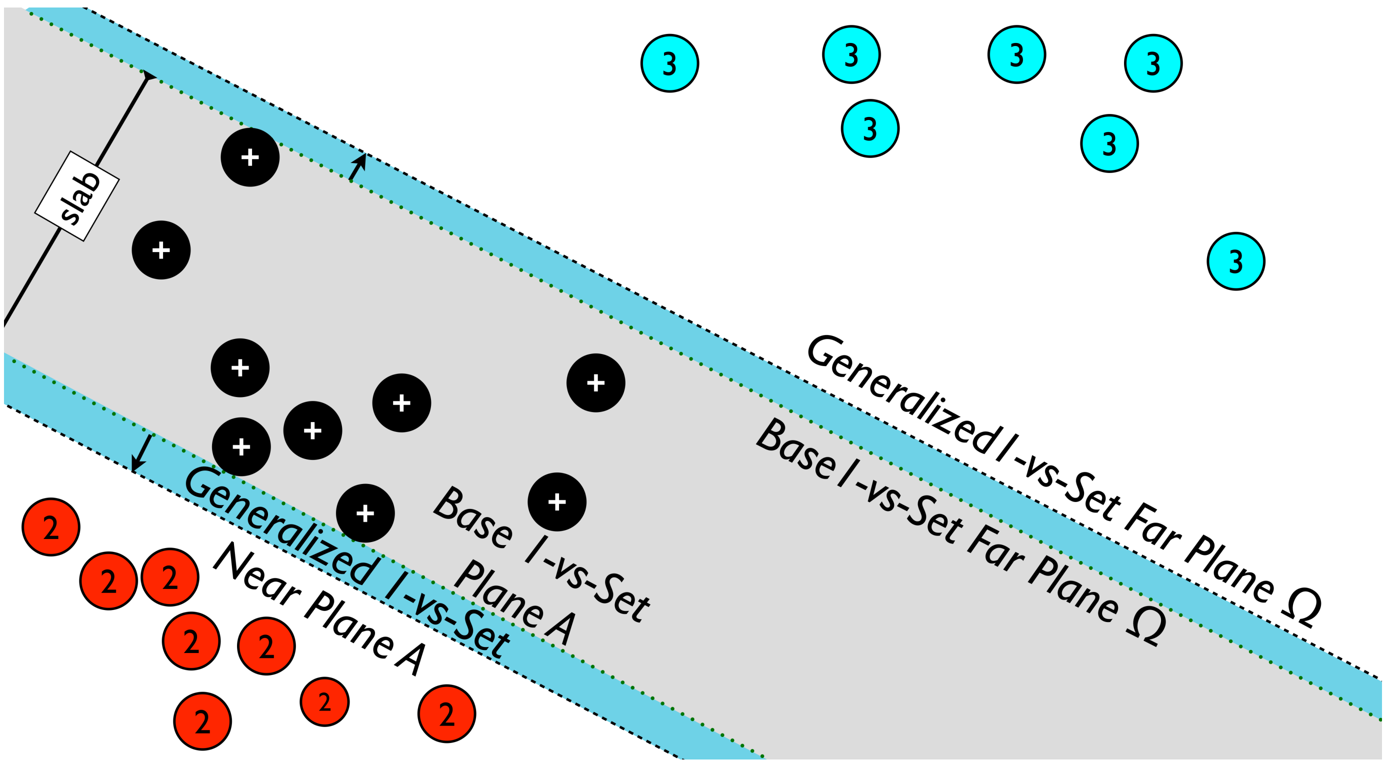

Specialization

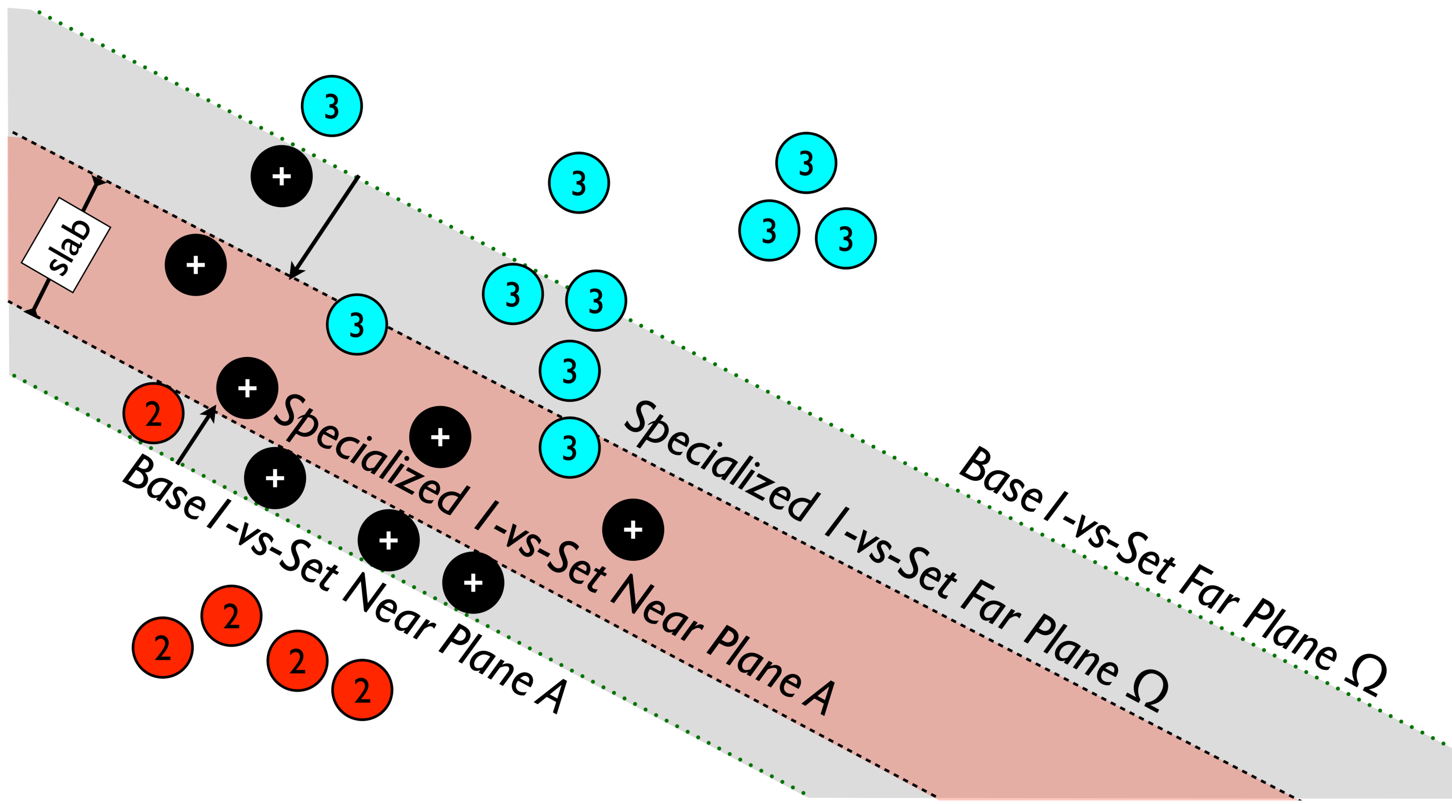

Open space risk for linear slab model

δ_A Marginal distance of near plane

δ_Ω Marginal distance of far plane

δ^+ Separation needed to account for all positive data

$$\frac{\delta_\Omega - \delta_A}{\delta^+}$$

Overgeneralization risk

$$\frac{\delta^+}{\delta_\Omega - \delta_A}$$

Overspecialization risk

Open space risk for linear slab model

Two additional terms

Importance of open space around A

Importance of open space around Ω

$$R_\varsigma = \frac{\delta_\Omega - \delta_A}{\delta^+} + \frac{\delta^+}{\delta_\Omega - \delta_A} + p_A \omega_A + p_\Omega \omega_\Omega$$

Margin around A

Margin around Ω

Sketch of the 1-vs-Set Machine training algorithm

1. Train a linear SVM f using $\hat{V}$ and $\hat{K}$

2. Generate decision scores for each training point in $\hat{V}$ and $\hat{K}$

3. Sort decision scores, where s_k is the minimum and s_j is the maximum

4. Initialize A to margin plane of f, and Ω to s_j

5. Iteratively move A to s_{k+1} or s_{k-1}, Ω to s_{j-1} or s_{j+1} to minimize $R_\varsigma(f) + \lambda_r R_\mathcal{E}$

1-vs-Set Machine Plane Refinement

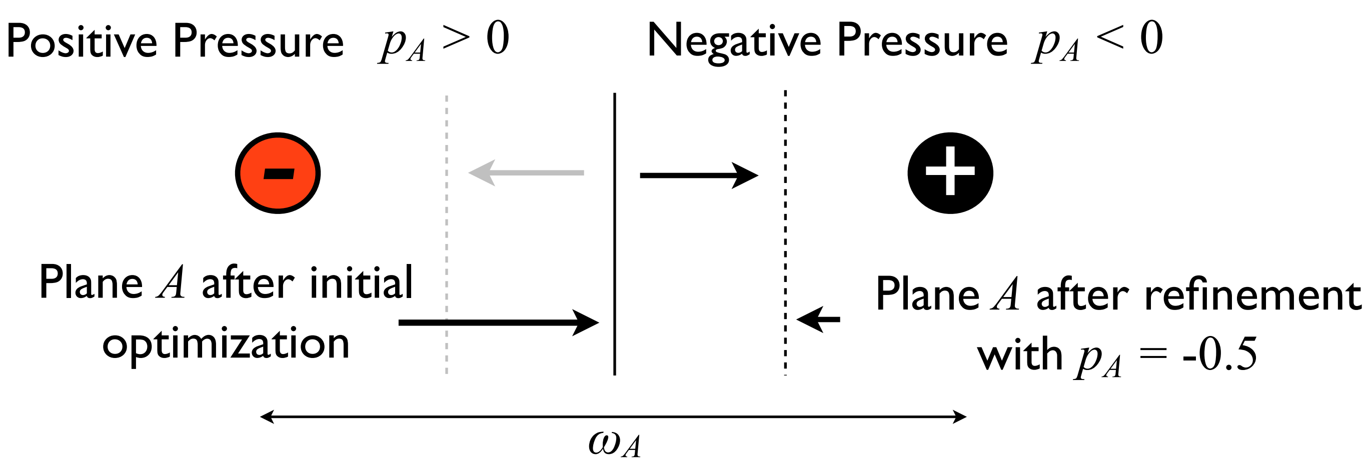

1-vs-Set Machine Prediction

function PREDICT(t_x, f, A, Ω)
    if $(A \leq f(t_x)$ **and** $f(t_x) \leq \Omega)$ **then Return** $+1$
    else Return -1
    end if
end function

How can we evaluate open set recognition in a controlled manner?

Accuracy as a statistic for open set problems

$$\text{Accuracy} = \frac{TP + TN}{TP + TN + FP + FN}$$

Imagine the following case:

1/100 TP correct
100,000/100,000 TN correct
99.9% accuracy!

F-measure as a statistic for open set problems

Consistent point of comparison across inconsistent precision and recall numbers:

$$\text{F-measure} = 2 \times \frac{\text{Precision} \times \text{Recall}}{\text{Precision} + \text{Recall}}$$

Open Set Object Recognition

Cross-data set methodology*

Training: Caltech 256



known classes

Testing: Caltech 256 + ImageNet



known classes +



unknown classes

Open Universe of 88 classes: 1 positive class, n training classes, 87 negative testing classes (532,400 images)

Open Universe of 212 classes: 1 positive class, n training classes, 211 negative testing classes (13,610,400 images)

A. Torralba and A. A. Efros, "Unbiased Look at Dataset Bias," in IEEE CVPR 2011.

Features

Histogram of Oriented Gradients

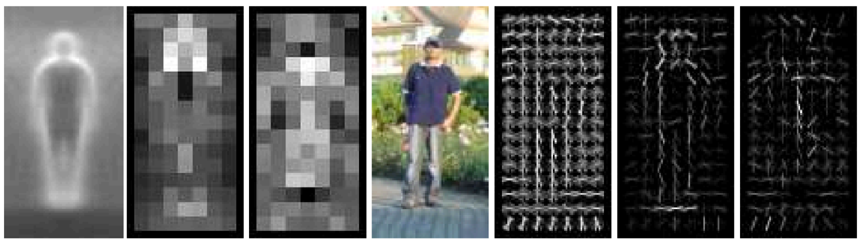

(Dalal and Triggs 2005) © 2005 IEEE

N. Dalal and B. Triggs, “Histograms of Oriented Gradients for Human Detection,” in IEEE CVPR, 2005

LBP-like descriptor

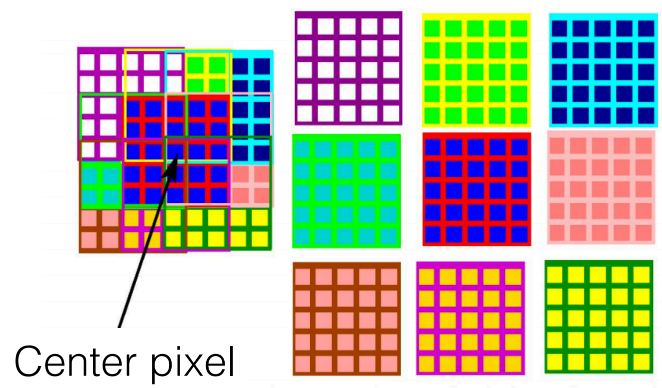

A. Sapkota, B. Parks, W.J. Scheirer, and T. Boult, “FACE-GRAB: Face Recognition with General Region Assigned to Binary Operator

1-vs-Set Machine vs. Typical SVMs

2-tailed paired t-test	binary 1-vs-Set	binary linear	binary RBF	1-class 1-vs-Set	1-class linear	1-class RBF
binary 1-vs-Set (HOG 88)		**	**	**	**	**
binary linear (HOG 88)	—		—	++	++	++
binary RBF (HOG 88)	—	++		++	++	++
1-class 1-vs-Set (HOG 88)	—	—	—		**	—
1-class linear (HOG 88)	—	—	—	—		—
1-class RBF (HOG 88)	—	—	—	—	++	
binary 1-vs-Set (HOG 212)		**	*	**	**	**
1-class 1-vs-Set (HOG 212)	—	—	—		—	*
binary 1-vs-Set (LBP-like 88)		**	**	**	**	**
1-class 1-vs-Set (LBP-like 88)	—	—	—		**	—
binary 1-vs-Set (LBP-like 212)		*	—	**	**	**
1-class 1-vs-Set (LBP-like 212)	—	—	—		**	—

** 1-vs-Set Machine is statistically significant at $p < 0.01$

* 1-vs-Set Machine is statistically significant at $p < 0.05$

++ Baseline Machine is statistically significant at $p < 0.01$

— No statistical significance

Top 25 classes for the open universe of 88 classes

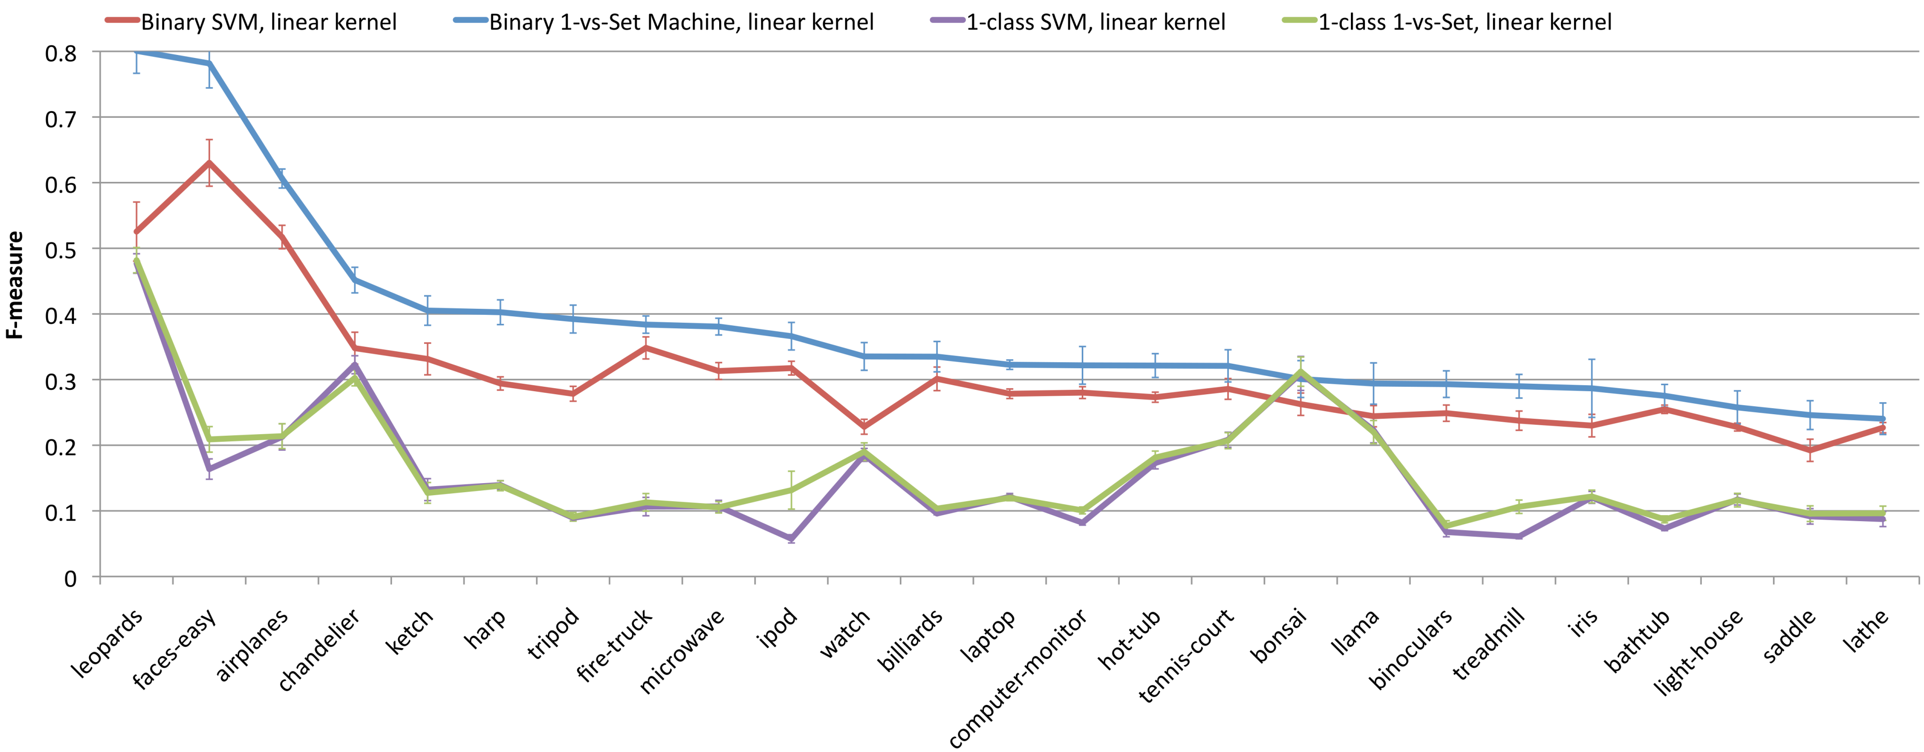

Top 25 classes for the open universe of 88 classes

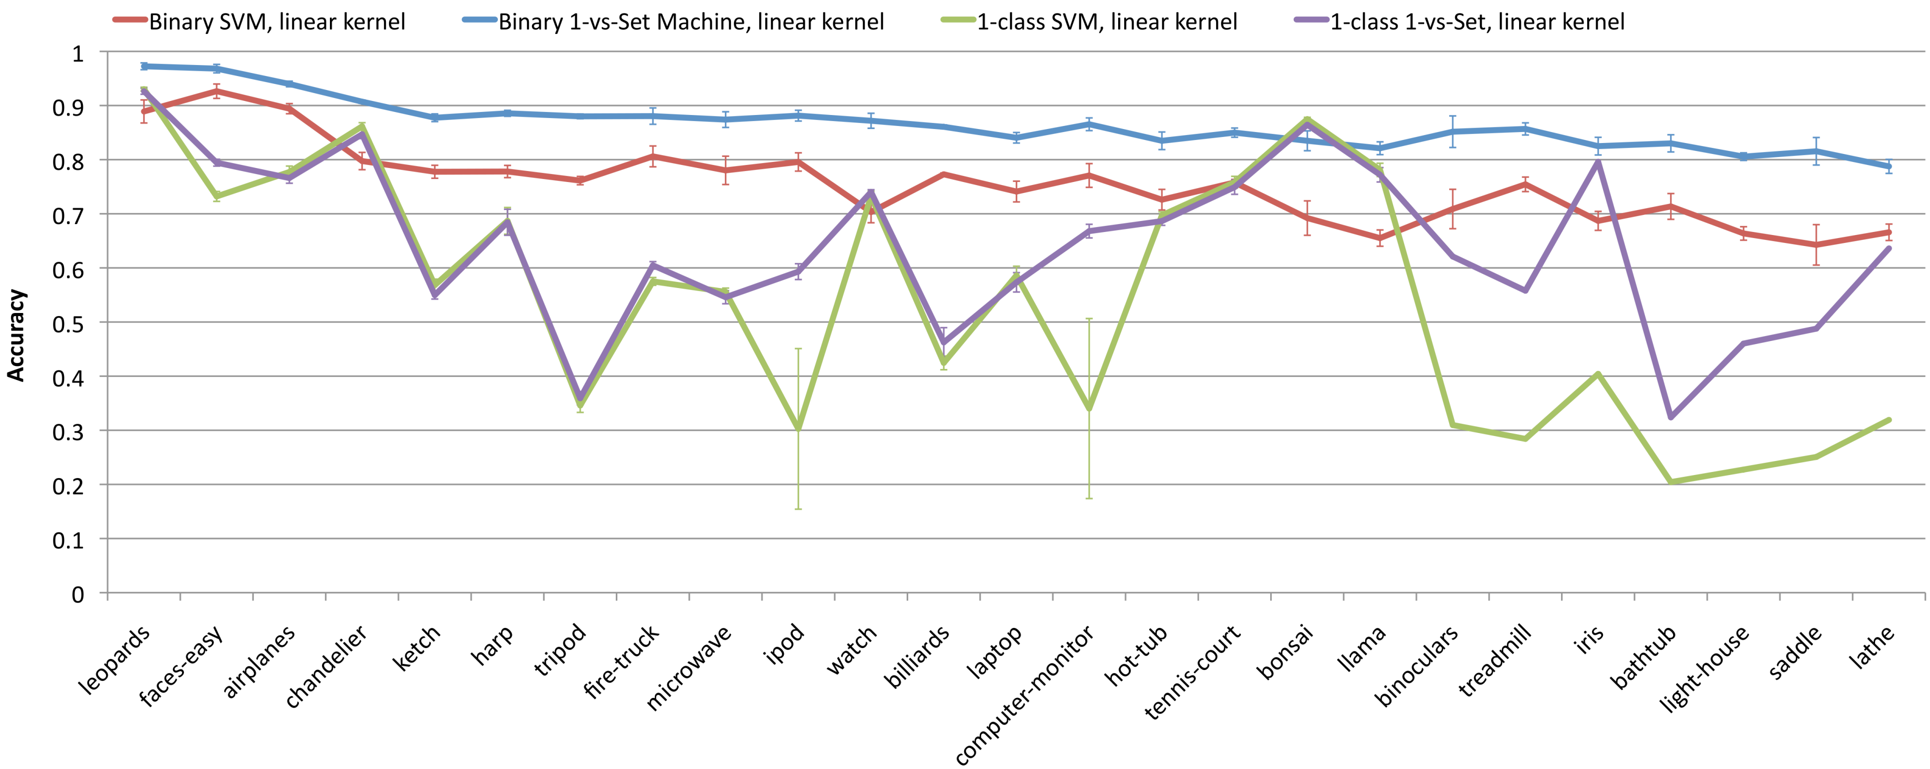

F-measure as a function of openness

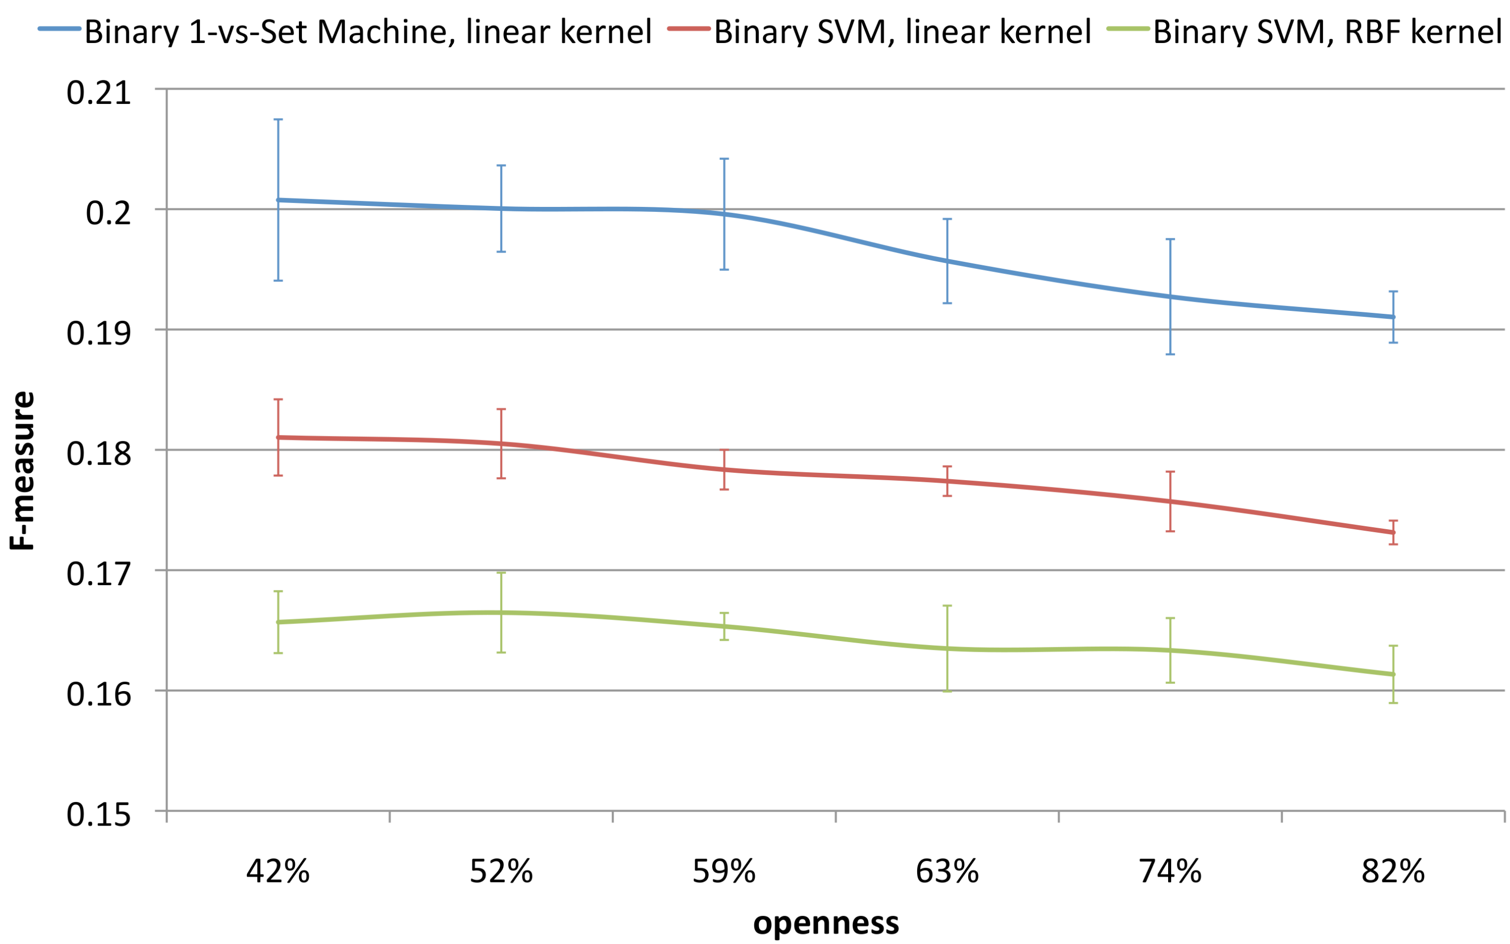

Near and far plane pressures for open universe of 88 classes

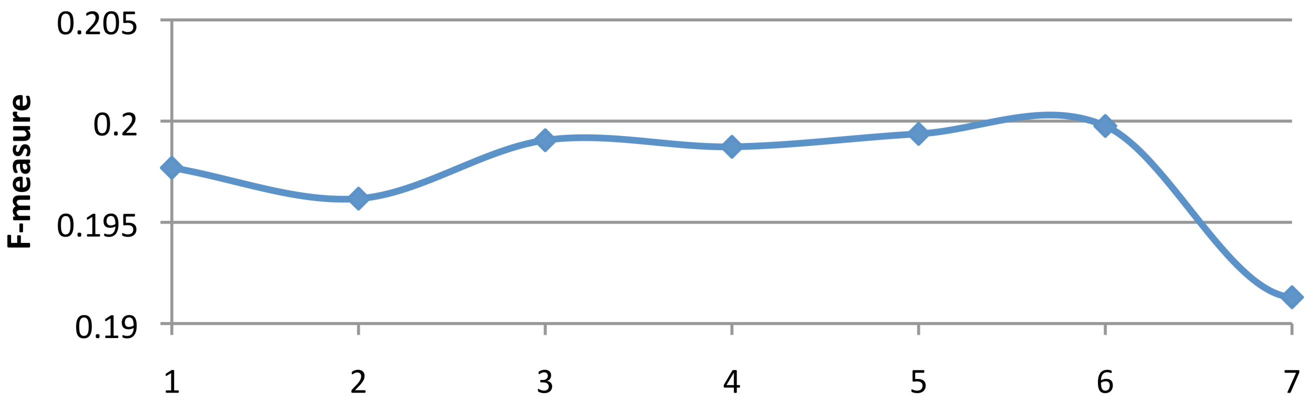

$p_\Omega \in \{1, \ldots, 7\}$ **at** $p_A = 1.6$

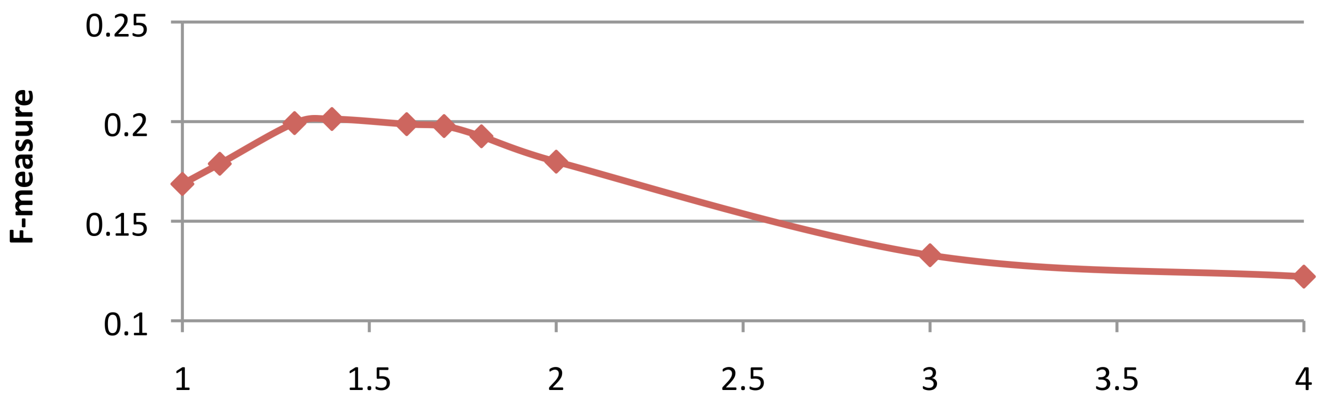

$p_A \in \{1, \ldots, 4\}$ **at** $p_\Omega = 4$

The second plane has an impact on recognition performance

Open Set Face Verification

Labeled Faces in the Wild

Genuine Pair

Impostor Pair

Impostor Pair

Impostor Pair

Gallery classes: 12 people with at least 50 images

Impostor classes: 82 other people in LFW

1,316 test images across all classes

Features: LBP-like and Gabor*

N. Pinto, J. J. DiCarlo, and D. D. Cox, “How Far Can You Get with a Modern Face Recognition Test Set Using Only Simple Features?” in IEEE CVPR, 2009.

Open set face verification

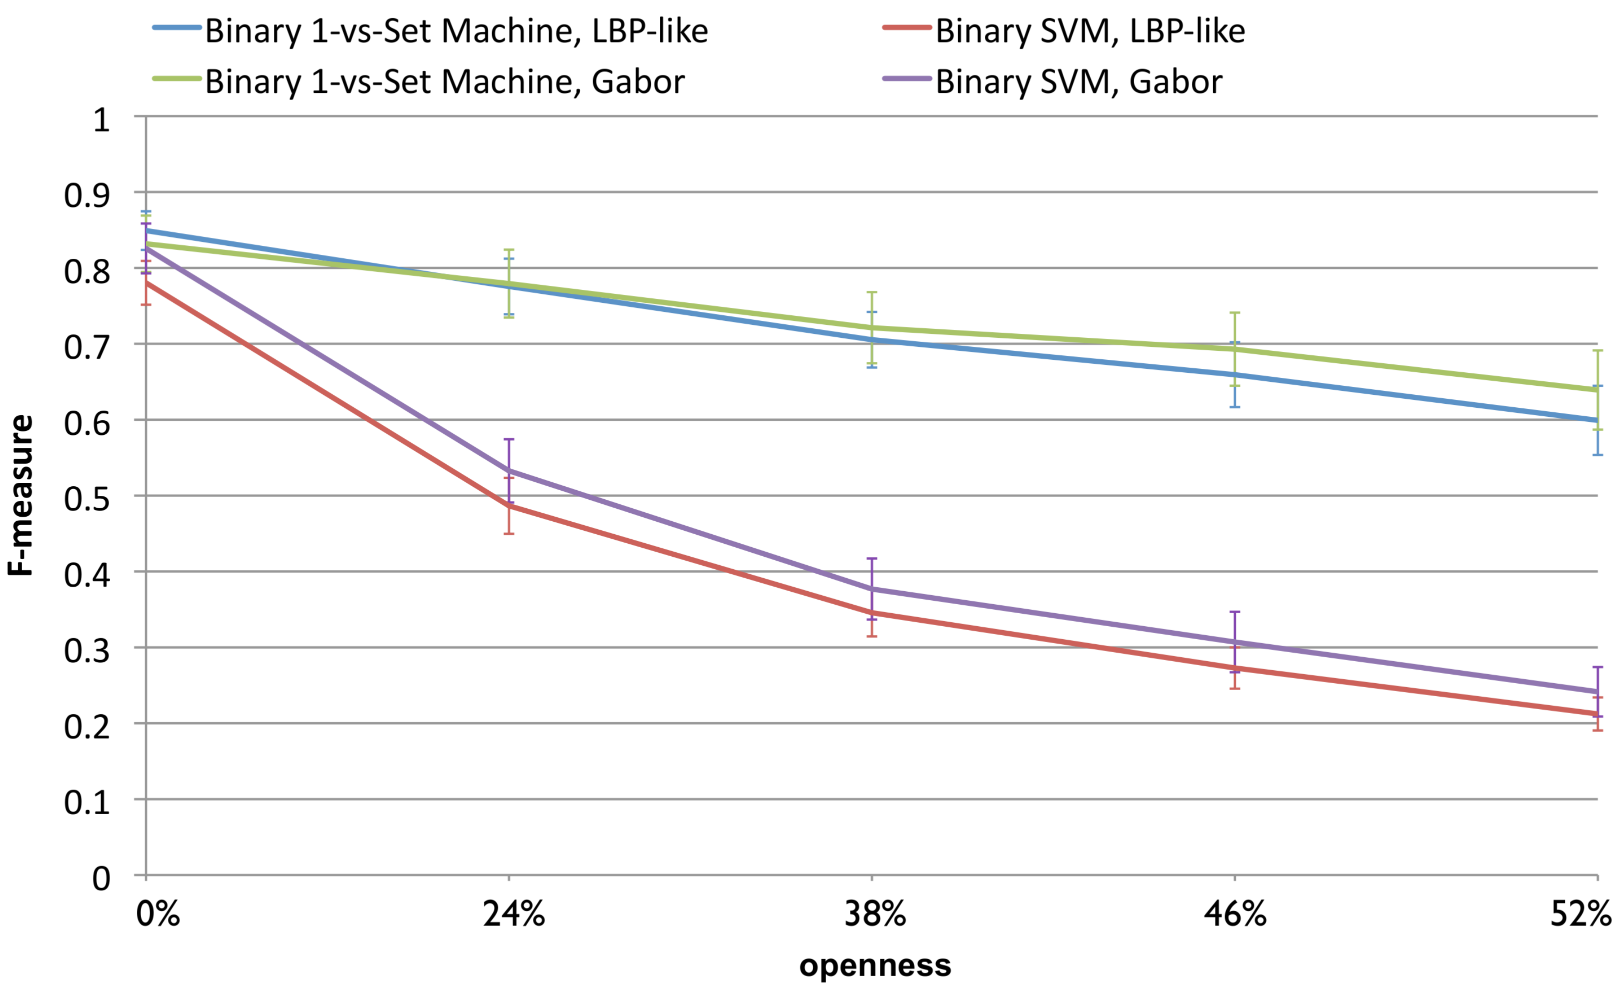

Further Reading

- W.J. Scheirer, A. Rocha, A. Sapkota, and T. Boult, “Towards Open Set Recognition,” IEEE T-PAMI, 35(7) July 2013.
- F. Costa, E. Silva, M. Eckmann, W.J. Scheirer, and A. Rocha, “Open Set Source Camera Attribution and Device Linking,” Pattern Recognition Letters, Accepted 2013.
- M.J. Wilber, W.J. Scheirer, P. Leitner, B. Heflin, J. Zott, D. Reinke, D. Delaney, T.E. Boult, “Animal Recognition in the Mojave Desert: Vision Tools for Field Biologists,” IEEE WACV, 2013.
- B. Heflin, W.J. Scheirer, and T.E. Boult, “Detecting and Classifying Scars, Marks, and Tattoos Found in the Wild,” IEEE BTAS, 2012.

Code

1-vs-Set Machine on GitHub:
https://github.com/tboult/libSVM-onevset

Data sets:
http://www.metarecognition.com/openset/