Multi-Class Open Set Recognition Using Probability of Inclusion

Lalit P. Jain[1], Walter J. Scheirer[2], and Terrance E. Boult[1,3]

[1] University of Colorado at Colorado Springs | [2]Harvard University | [3]Securics, Inc.

Overview

The perceived success of recent visual recognition approaches has largely been derived from their performance on classification tasks, where all possible classes are known at training time.

ImageNet Classification

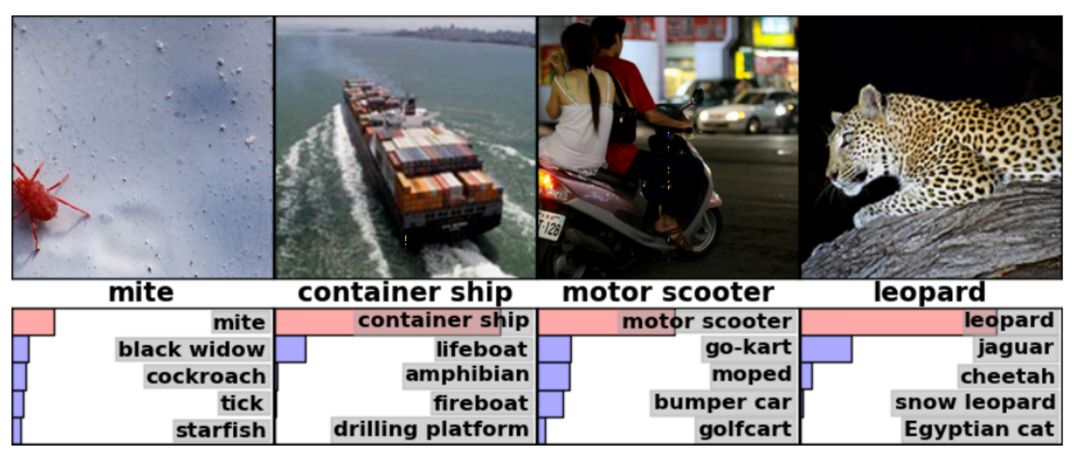

Krizhevsky et al. 2012

ImageNet Detection

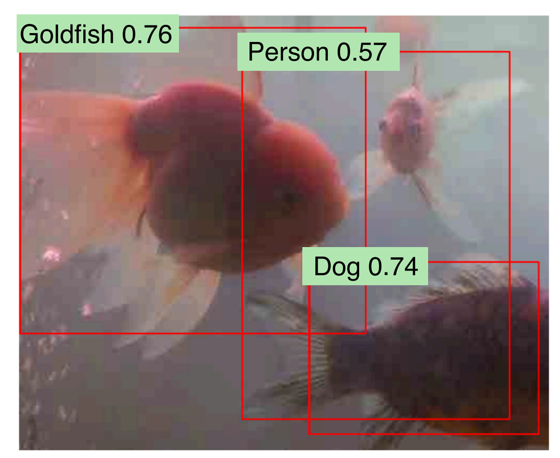

Girshick et al. 2014

Even with state-of-the-art representation learning, results are disappointing for **open set** recognition, where the possibility of encountering classes not present during training exists in testing:

- ILSVRC2014 Best Classification Result: 6.66% Error Rate
- ILSVRC2014 Best Detection Result: **43.9 mean AP**

How can supervised learning directly address open set recognition?

Previous Approaches	Drawbacks
Novel Class Detection	No multi-class recognition
Outlier Rejection	Works with same training/testing backgrounds
Sigmoid Calibration	Sigmoid function often the wrong model

Proposed Approach: fit a robust **single-class probability model** over the positive class scores from a discriminative binary classifier

- 1-Class RBF SVM
- Binary RBF SVM
- Binary + P_I-SVM
- All three models
- Binary SVM SVs
- 1-Class SVM SVs

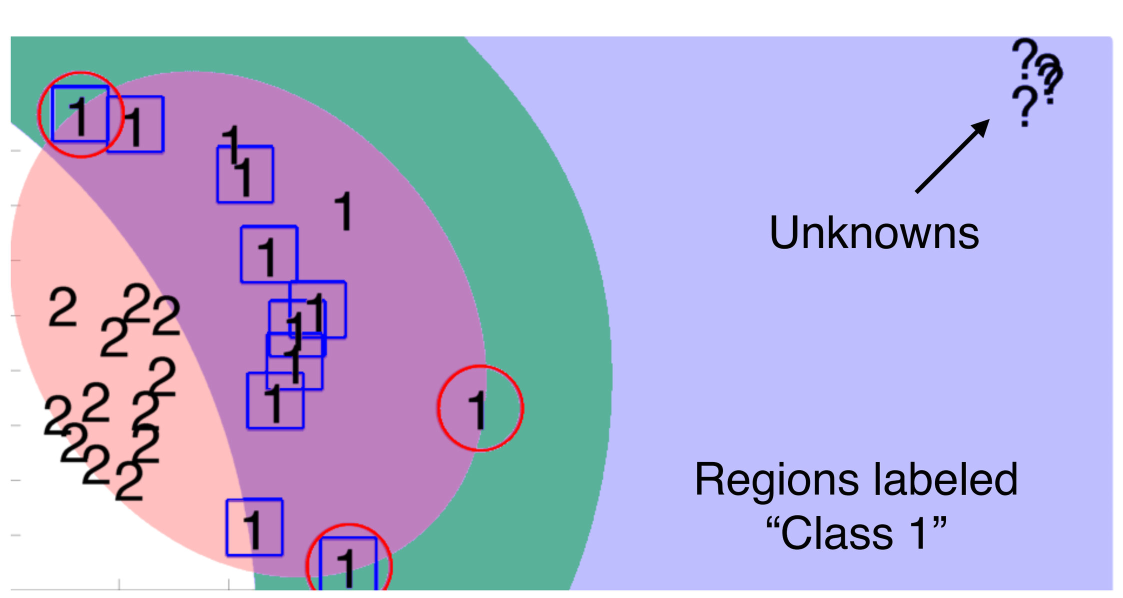

Benefits

- Use of binary model helps discriminate positive class from known negatives
- Use of single-class probability model prevents misclassification of unknowns
- Does not directly use negative data for probability estimation
- Does not assume any particular distribution over all scores
- Fast and easy to compute

The P_I-SVM Algorithm

Objective: compute a per class unnormalized posterior probability estimate for any input sample.

Observation: probability calibration is well-defined near a decision boundary; **extreme value theory** can be invoked to model the probability of inclusion P_I

$$P_I(y|x,\theta_y) = \xi\rho(y)P_I(x|y,\theta_y) = \xi\rho(y)(1-e^{-(\frac{x-\tau_y}{\lambda_y})^{\kappa_y}})$$

EVT params. constant (1) prior Weibull CDF

$$y^* = \underset{y\in\mathcal{C}}{\operatorname{argmax}}\, P_I(y|x,\theta_y) \quad \text{subject to} \quad P_I(y^*|x,\theta_{y^*}) \geq \delta.$$

P_I-SVM Training

1. Train a 1-vs-Rest RBF SVM $h(x)$ for a set of known classes
2. For each class:
   1. Generate scores for available positive training data for class using $h(x)$
   2. Estimate size of tail for sorted scores near the decision boundary
   3. Fit a Weibull distribution to the tail

P_I-SVM Probability Estimation

1. For each input sample:
   1. Generate a score for each known class using $h(x)$
   2. Apply class-specific Weibull CDFs to each score and multiply by $\rho(y)$

How can we estimate the tail size for Weibull fitting?

Consider points within distance ϵ of the SVM decision boundary as extrema

$$B^+(x;\epsilon) = \begin{cases} 1 & \text{if } h(x) \leq \epsilon \\ 0 & \text{otherwise} \end{cases} \quad \text{and} \quad \hat{T}^+_\epsilon = \max(3, \psi \times |\alpha^+|)$$

\# Positive SVs

Cross-Class-Validation for Open Set Evaluation

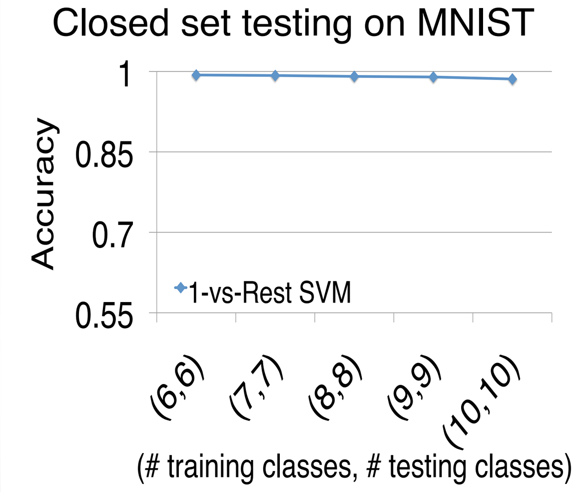

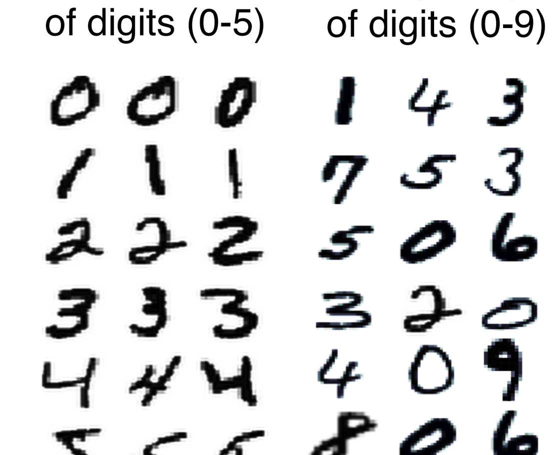

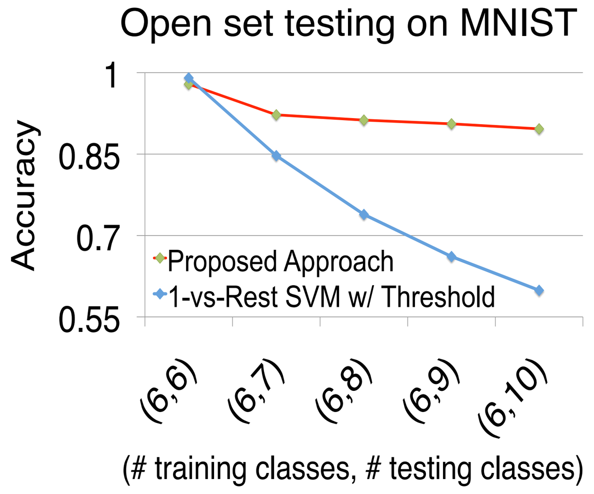

Simulate unknown classes of open set recognition by defining number of: Training classes (t), target known classes (η), classifiers (m), validation classes (e)

$$\text{openness} = 1 - \sqrt{(2\times t)/(m\times\eta+e)}$$

Experiments

Open Set Object Detection for Caltech-256 / ImageNet

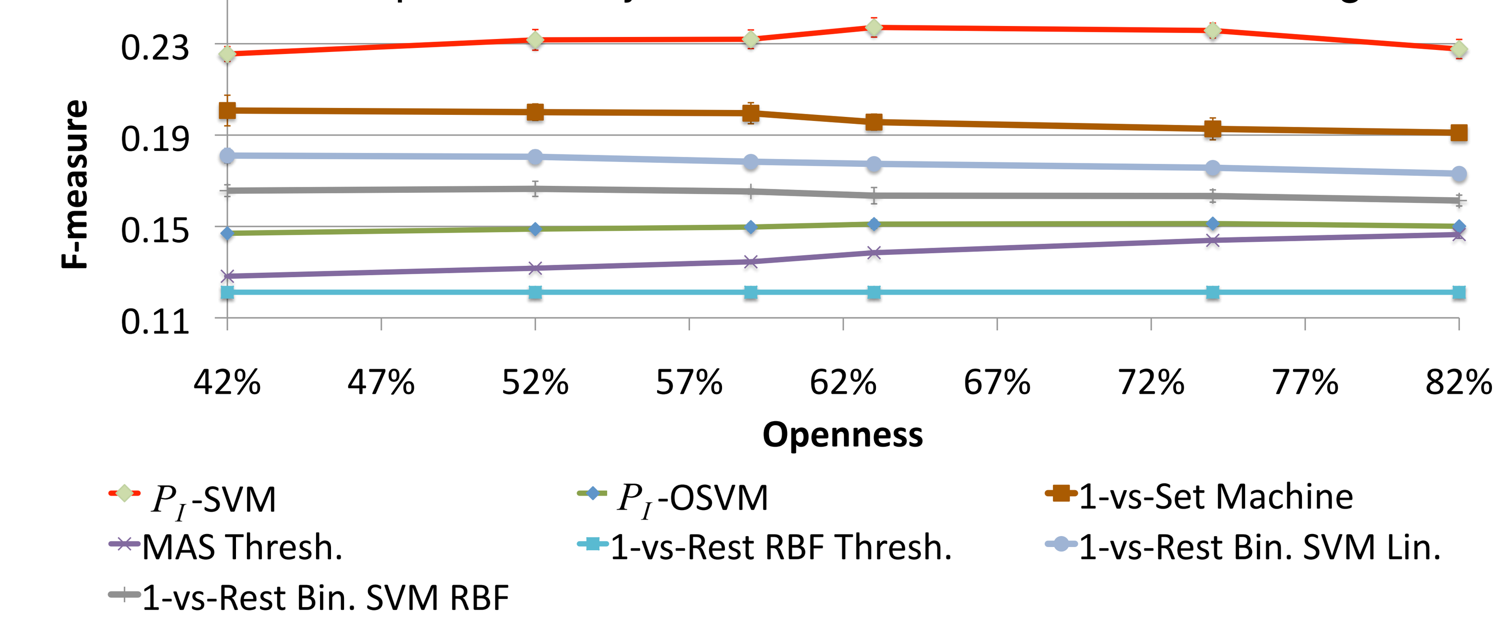

Multi-Class Open Set Recognition for LETTER

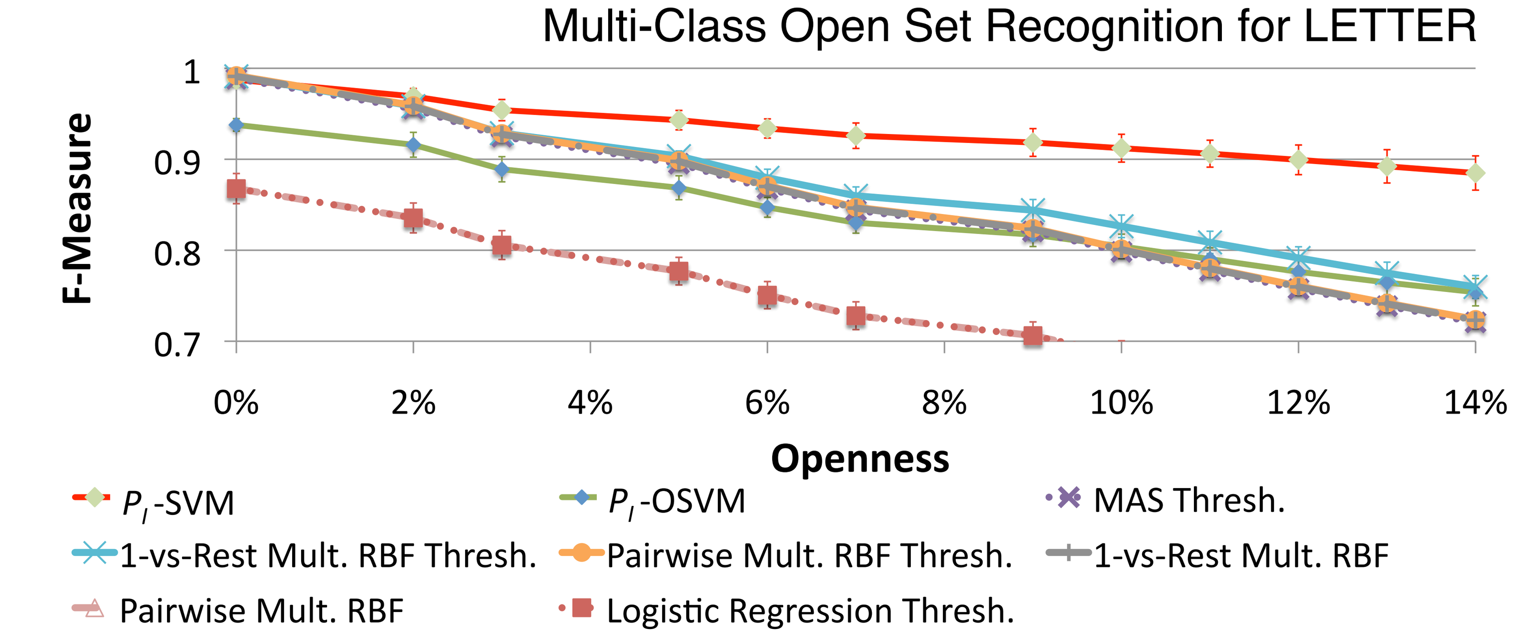

Multi-Class Open Set Recognition for MNIST

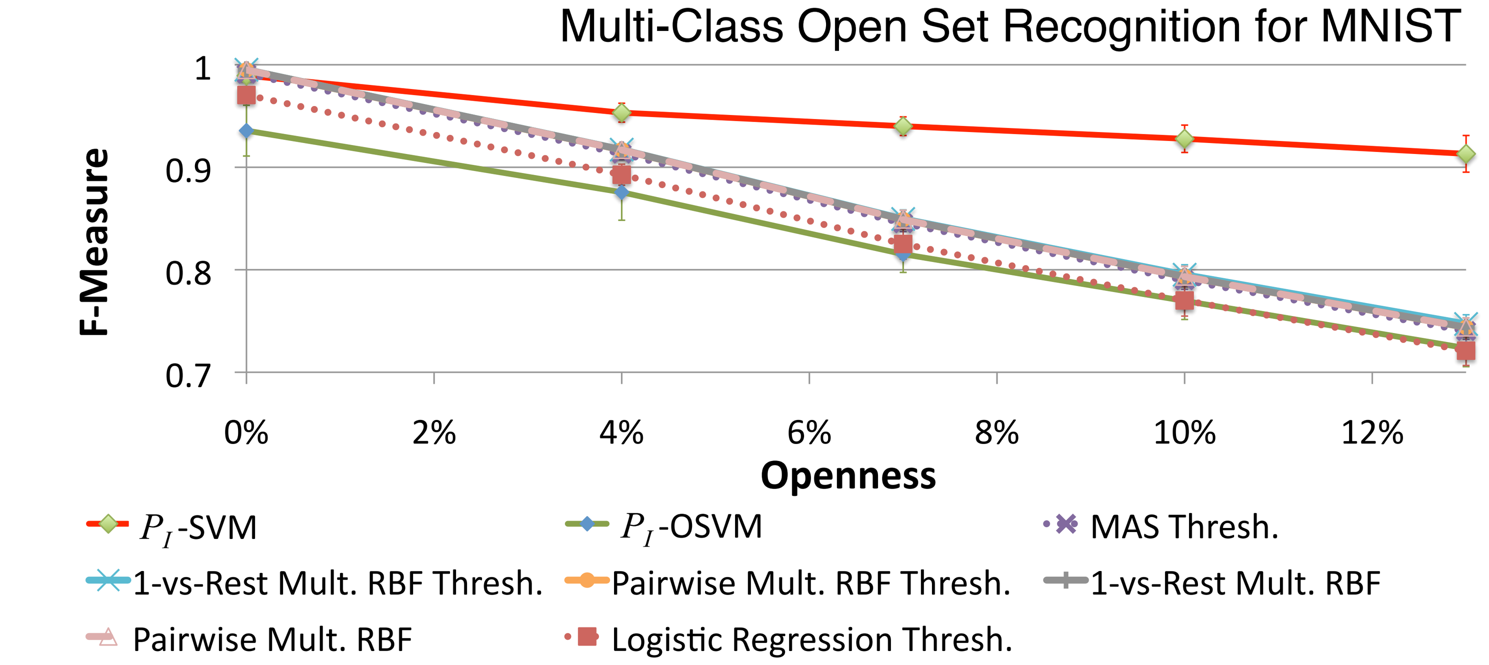

Software & References

- The P_I-SVM Code: https://github.com/ljain2/libsvm-openset
- The Meta-Recognition Library: http://www.metarecognition.com

1. W.J. Scheirer, A. Rocha, A. Sapkota, and T.E. Boult, "Towards Open Set Recognition," *T-PAMI*, 36(7): 1757-1772, Jul. 2013.
2. W.J. Scheirer, N. Kumar, P.N. Belhumeur, and T.E. Boult, "Multi-Attribute Spaces: Calibration for Attribute Fusion and Similarity Search," *CVPR*, 2012.
3. W.J. Scheirer, A. Rocha, R. Micheals and T.E. Boult, "Meta-Recognition: The Theory and Practice of Recognition Score Analysis," *T-PAMI*, 33(8): 1689-1695, Aug. 2011.

Acknowledgements

This work was supported in part by ONR MURI Award N00014-08-1-0638 and NSF IIS-1320956.