Towards Open World Recognition

Abhijit Bendale, Terrance Boult
University of Colorado at Colorado Springs

With the of advent rich classification models and high computational power visual recognition systems have found many operational applications. Recognition in the real world poses multiple challenges that are not apparent in controlled lab environments. The datasets are dynamic and novel categories must be continuously detected and then added. At prediction time, a trained system has to deal with myriad unseen categories. Operational systems require minimal downtime, even to learn. To handle these operational issues, we present the problem of **Open World Recognition** and formally define it. We prove that thresholding sums of monotonically decreasing functions of distances in linearly transformed feature space can balance "open space risk" and empirical risk. Our theory extends existing algorithms for open world recognition. We present a protocol for evaluation of open world recognition systems. We present the Nearest Non-Outlier (NNO) algorithm that evolves model efficiently, adding object categories incrementally while detecting outliers and managing open space risk.

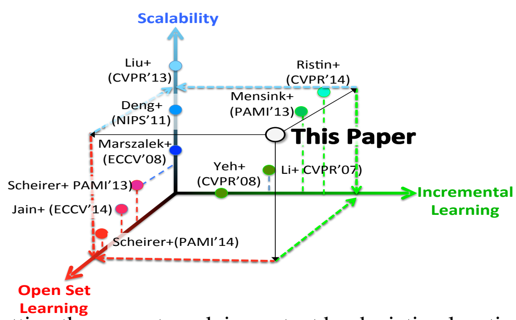

Figure 1: Putting the current work in context by depicting locations of prior work with respect to three axes of the major issues for open world recognition: open set learning, incremental learning and scalability. In this work, we present a system that is scalable, can handle open set recognition and can learn new categories incrementally without having to retrain the system every time a new category arrives.

The first contribution of this paper is a formal definition of the problem of open world recognition, which extends the existing definition of open set recognition which was defined for a static notion of set. In order to solve open world recognition, the system needs to be robust to unknown classes, but also be able to move through the stages and knowledge progression.

Open World Recognition: A solution to open world recognition is a tuple $[F, \varphi, \nu, L, I]$ with:

1. A **multi-class open set recognition function** $F(x): \mathbb{R}^d \mapsto \mathbb{N}$ using a vector function $\varphi(x)$ of i per-class measurable recognition functions $f_i(x)$, also using a **novelty detector** $\nu(\varphi): \mathbb{R}^i \mapsto [0,1]$. We require the per class recognition functions $f_i(x) \in \mathcal{H}: \mathbb{R}^d \mapsto \mathbb{R}$ for $i \in \mathcal{K}_t$ to be open set recognition functions that manage open space risk as Eq.4 [1]. The novelty detector $\nu(\varphi): \mathbb{R}^i \mapsto [0,1]$ determines if results from vector of recognition functions is from an unknown (0) class.

2. A labeling process $L(x): \mathbb{R}^d \mapsto \mathbb{N}^+$ applied to novel unknown data U_t from time t, yielding labeled data $D_t = \{(y_j, x_j)\}$ where $y_j = L(x_j) \forall x_j \in U_t$. Assume the labeling finds m new classes, then the set of known classes becomes $\mathcal{K}_{t+1} = \mathcal{K}_t \cup \{i+1, \ldots i+m\}$.

3. An incremental learning function $I_t(\varphi; D_t): \mathcal{H}^i \mapsto \mathcal{H}^{i+m}$ to scalably learn and add new measurable functions $f_{i+1}(x) \ldots f_{i+m}(x)$, each of which manages open space risk, to the vector φ of measurable recognition functions.

This is an extended abstract. The full paper is available at the Computer Vision Foundation webpage.

Ideally, all of these steps should be automated, but herein we presume supervised learning with labels obtained by human labelling.

Second contribution of the work is a recognition system that can continuously learn new object categories in an open world model. We build on the concept of a Compact Abating Probability (CAP) model [3] but generalize the model showing that any non-negative combination of abating functions, e.g., a convex combination of decreasing functions of distance, can be thresholded to have zero open space risk. We further show (see theorem 1 and 2 in [1]) we can work in linearly transformed spaces, including projection onto subspaces, and still manage open space risk and NCM type algorithms manage open space risk. In particular, we show how to extend Nearest Class Mean type algorithms (NCM) [2], to a Nearest Non-Outlier (NNO) algorithm that can balance open space risk and accuracy.

Nearest Non-Outlier (NNO) Algorithm: To adapt NCM for open world recognition, we introduce Nearest Non-Outlier (NNO) which uses a measurable recognition function consistent with Theorems 1 and 2 in [1]. Let NNO represent its internal model as a vector of means $\mathcal{M} = [\mu_1, \ldots \mu_k]$. Let $W \in \mathbb{R}^{d \times m}$ be the linear transformation dimensional reduction weight matrix learned by the process described in [2?]. Then given τ, let

$$\hat{f}_i(x) = \frac{\Gamma(\frac{m}{2}+1)}{\pi^{\frac{m}{2}} \tau^m} \left(1 - \frac{1}{\tau} \|W^\top x - W^\top \mu_i\|\right) \quad (1)$$

be our measurable recognition function with $\hat{f}_i(x) > 0$ giving the probability of being in class i, where Γ is the standard gamma function which occurs in the volume of a m-dimensional ball. Intuitively, the probability is a tent-like function in the sphere and the first fraction in eqn 1 comes from volume of m-sphere and ensures that the probability integrates to 1.

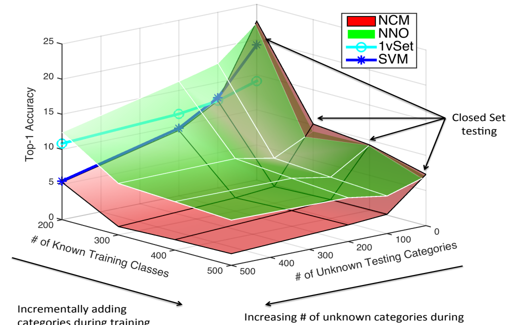

Figure 2: Open World learning on data from ILSVRC'10 challenge. Top-1 accuracy is plotted as a function of known classes in the system and unknown classes used during testing NNO performs at par with NCM in closed set testing (marked with arrows in above figure) as categories are added incrementally to the system. The proposed Nearest Non-Outlier (NNO) approach of handling unknown categories based on extending NCM with Compact Abating Probabilities remains robust in both circumstances: as more number of categories are added to the system and as the system is tested with more unknown categories.

Finally, we present a protocol for evaluation for open world recognition, and use this protocol to show our NNO algorithm perform significantly better on open world recognition evaluation using ILSVRC'10 challenge.

[1] A Bendale and T Boult. Towards open world recognition. *CVPR*, 2015.

[2] Thomas Mensink, Jakob Verbeek, Florent Perronnin, and Gabriela Csurka. Distance-based image classification: Generalizing to new classes at near-zero cost. *IEEE TPAMI*, 2013.

[3] W.J. Scheirer, L.P. Jain, and T.E. Boult. Probability models for open set recognition. *IEEE TPAMI*, 2014.