

© 2017 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Pre-print of article that will appear in T-PAMI.

# The Extreme Value Machine

Ethan M. Rudd, Lalit P. Jain,  
Walter J. Scheirer, *Senior Member, IEEE*,  
and Terrance E. Boult, *IEEE Fellow*

**Abstract**—It is often desirable to be able to recognize when inputs to a recognition function learned in a supervised manner correspond to classes unseen at training time. With this ability, new class labels could be assigned to these inputs by a human operator, allowing them to be incorporated into the recognition function — ideally under an efficient incremental update mechanism. While good algorithms that assume inputs from a fixed set of classes exist, *e.g.*, artificial neural networks and kernel machines, it is not immediately obvious how to extend them to perform incremental learning in the presence of unknown query classes. Existing algorithms take little to no distributional information into account when learning recognition functions and lack a strong theoretical foundation. We address this gap by formulating a novel, theoretically sound classifier — the Extreme Value Machine (EVM). The EVM has a well-grounded interpretation derived from statistical Extreme Value Theory (EVT), and is the first classifier to be able to perform nonlinear kernel-free variable bandwidth incremental learning. Compared to other classifiers in the same deep network derived feature space, the EVM is accurate and efficient on an established benchmark partition of the ImageNet dataset.

**Index Terms**—Machine Learning, Supervised Classification, Open Set Recognition, Open World Recognition, Statistical Extreme Value Theory

## I. INTRODUCTION

Recognition problems which evolve over time require classifiers that can incorporate novel classes of data. What are the ways to approach this problem? One is to periodically retrain classifiers. However, in situations that are time or resource constrained, periodic retraining is impractical. Another possibility is an online classifier that incorporates an efficient incremental update mechanism. While methods have been proposed to solve the incremental learning problem, they are computationally expensive [1]–[4], or provide little to no characterization of the statistical distribution of the data [5]–[8]. The former trait is problematic because it is contrary to the fundamental motivation for using incremental learning — that of an efficient update system — while the latter trait places limitations on the quality of inference.

There is also a more fundamental problem in current incremental learning strategies. When the recognition system encounters a novel class, that class should be incorporated into the learning process at subsequent increments. But in order to do so, the recognition system needs to identify novel classes in the first place. For this type of *open set* problem in which unknown classes appear at query time, we cannot rely on a *closed set* classifier, even if it supports incremental learning, because it implicitly assumes that all query data is well represented by the training set.

E. M. Rudd, L. P. Jain, and T. E. Boult are with the Department of Computer Science, University of Colorado Colorado Springs, Colorado Springs, CO, 80918. W. J. Scheirer is with the Department of Computer Science and Engineering, University of Notre Dame, Notre Dame, IN, 46556. Corresponding Author’s E-mail: erudd@vast.uccs.edu

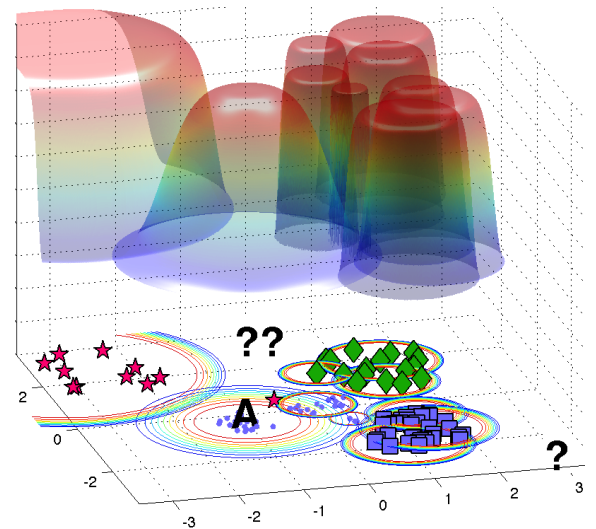


Fig. 1. A solution from the proposed EVM algorithm trained on four classes: dots, diamonds, squares, and stars. The colors in the isocontour rings show a  $\Psi$ -model (probability of sample inclusion) for each *extreme vector* (EV) chosen by the algorithm, with red near 1 and blue near .005. Via kernel-free non-linear modeling, the EVM supports open set recognition and can reject the three “?” inputs that lie beyond the support of the training set as “unknown.” Each  $\Psi$ -model has its own independent shape and scale parameters learnt from the data, and supports a soft-margin. For example, the  $\Psi$ -model for the blue dots corresponding to extreme vector A has a more gradual fall off, due to the effect of the outlier star during training.

Closed set classifiers have been developed for approximating the Bayesian optimal posterior probability,  $P(C_l|x'; C_1, C_2, \dots, C_M), l \in \{1, \dots, M\}$ , for a fixed set of classes, where  $x'$  is an input sample,  $l$  is the index of class  $C_l$  (a particular known class), and  $M$  is the number of known classes. When  $\Omega$  unknown classes appear at query time, however, the Bayesian optimal posterior becomes  $P(C_{\tilde{l}}|x'; C_1, C_2, \dots, C_M, U_{M+1}, \dots, U_{M+\Omega}), \tilde{l} \in \{1, \dots, M + \Omega\}$ , a distribution that we cannot model because classes  $U_{M+1}$  through  $U_{\Omega}$  are unknown. Making closed set assumptions in training leads to regions of unbounded support for an open set problem because a sample  $x'$  from an unknown class  $U_{\tilde{l}}$  will be misclassified as a known class  $C_l$ . For classifiers that assess confidence in terms of signed distance from a decision boundary, or some calibration thereof, this misclassification will occur with high confidence if  $x'$  is far from any known data — a result that is very misleading. Scheirer et al. [9] termed this problem *open space risk*.

More formally, let  $f$  be a measurable recognition function for known class  $C$ ,  $\mathcal{O}$  be the open space, and  $S_o$  be a ball of radius  $r_o$  that includes all of the known positive training examples  $x \in C$  as well as the open space  $\mathcal{O}$ . Open space risk  $R_{\mathcal{O}}(f)$  for class  $C$  can be defined as  $R_{\mathcal{O}}(f) = \frac{\int_{\mathcal{O}} f_C(x) dx}{\int_{S_o} f_C(x) dx}$ , where open space risk is considered to be the relative measure of positively labeled open space compared to the overall measure of positively labeled space. In this probabilistic formulation, the objective of *open set recognition* is to exercise a rejection option [10] for queries that lie beyond the reasonable support of known data, thus mitigating this risk.

Open set recognition [9], [11], [12], and more generally novelty/outlier detection [13], [14] are well established areas in their own right, but much less research has been conducted on how to treat unknown samples in an incremental context, which is the focus of this work. When an open set recognition system labels a sample as unknown, it suggests that the model was not trained with data from the class corresponding to that sample. In response, the classifier’s decision boundaries should be updated so that the system can incorporate this new class information for future decision making. But there is a caveat: full retraining is not always feasible, depending on timing constraints and the availability of computational resources.

Recent work [5] extended the open set recognition problem to include the incremental learning of new classes in a regime dubbed *open world recognition*, which is the problem we are most concerned with in this paper. An effective open-world recognition system must perform four tasks: detecting unknowns, choosing which points to label for addition to the model, labelling the points, and updating the model. An algorithm, nearest non-outlier (NNO), was proposed as a demonstration of these elements — the first of its kind. Unfortunately, NNO lacks strong theoretical grounding. The algorithm uses thresholded distances from the nearest class mean as its decision function, and otherwise ignores distributional information. Weak classifiers are a persistent problem for this task: it is not immediately obvious how one might extend class boundary models from classical machine learning theory (*e.g.*, neural networks and kernel machines) to incorporate both incremental learning and open set constraints. A new formulation is required.

In this article we address the construction of a compact representation of open world decision boundaries based on the distribution of the training data. Obtaining this representation is difficult because training points that do not contribute to a decision boundary at one point in time may be extremely relevant in defining a decision boundary later on, and retraining on all points is infeasible at large scales. Moreover, by the definition of the open world problem, the hypothesis space will be under-sampled, so in many cases linearity of the decision boundaries cannot be guaranteed and the data bandwidth is unknown. So how does one obtain a compact statistical model without discarding potentially relevant points — especially in regions where the data bandwidth is unknown? To this end, we introduce the Extreme Value Machine (EVM), a model which we derive from statistical extreme value theory (EVT).

EVT dictates the functional form for the radial probability of inclusion of a point with respect to the class of another. By selecting the points and distributions that best summarize each class, *i.e.*, are least redundant with respect to one another, we arrive at a compact probabilistic representation of each class’s decision boundary, characterized in terms of its *extreme vectors* (EV), which provides an abating bound on open space risk. This is depicted in schematic form in Fig. 1. When new data arrives, these EVs can be efficiently updated. The EVM is a scalable nonlinear classifier, with radial inclusion functions that are in some respects similar to RBF kernels, but unlike RBF kernels assume variable bandwidths and skew that are data derived and grounded in EVT.

## II. RELATED WORK

With respect to classifiers that mitigate open space risk at classification time, the 1-vs-Set machine [9] approaches the problem of open set recognition by replacing the half-space of a binary linear classifier by bounding the positive data with two hyperplanes. An algorithm similar to the 1-vs-Set machine was described by Cevikalp and Triggs [15] for object detection, where a binary classifier with a slab is combined with a nonlinear SVDD classifier for just the positive class. In later work, Scheirer et al. introduced the W-SVM for multi-class open set recognition problems using nonlinear kernels, with provable guarantees of open space risk reduction [12]. These nonlinear models were more accurate, but also more costly to compute and store. For the more expansive problem of open world recognition, Bendale and Boulton modified the Nearest Class Mean [16] algorithm by limiting open space risk for model combinations and transformed spaces, resulting in the NNO algorithm introduced in Sec. I, which we will use as a baseline for comparison in Sec. V.

Other approaches exist for related problems involving unknown class data such as multi-class novelty detection [17], domain adaptation [18], and zero-shot classification [19]. However, these problems need not be addressed by a classifier that is incrementally updated over time with class-specific feature data. More related is the problem life-long learning, where a classifier receives tasks and is able to adapt its model in order to perform well on new task instances. Pentina and Ben-David [20] lay out a cogent theoretical framework for SVM-based life-long learning tasks, but leave the door open to specific implementations that embody it. Along these lines, Royer and Lampert [21] describe a method for classifier adaptation that is effective when inherent dependencies are present in the test data. This works for fine-grained recognition scenarios, but does not address unknown classes that are well separated in visual appearance from the known and other unknown classes. The problem most related to our work is rare class discovery, for which Haines and Xiang have proposed an active learning method that jointly addresses the tasks of class discovery and classification [22]. We consider their classification algorithm in Sec. V, even though we do not make distinctions between common and rare unknown classes.

There is growing interest in statistical extreme value theory for visual recognition. With the observation that the tails of any distance or similarity score distribution must always follow an EVT distribution [23], highly accurate probabilistic calibration models became possible, leading to strong empirical results for multi-biometric fusion [24], describable visual attributes [25], and visual inspection tasks [26]. EVT models have also been applied to feature point matching, where the Rayleigh distribution was used for efficient guided sampling for homography estimation [27], and the notion of extreme value sample consensus was used in conjunction with RANSAC for similar means [28]. Work in machine learning has shown that EVT is a suitable model for open set recognition problems, where one- [11] and two-sided calibration models [12], [29] of decision boundaries lead to better generalization. However, these are *post hoc* approaches that do not apply EVT at training time.

### III. THEORETICAL FOUNDATION

As discussed in Sec. I and as illustrated in Fig. 1, each class in the EVM's training set is represented by a set of extreme vectors, where each vector is associated with a radial inclusion function modeling the *Probability of Sample Inclusion* (PSI or  $\Psi$ ). Here we derive the functional form for  $\Psi$  from EVT; this functional form is not just a mathematically convenient choice — it is statistically guaranteed to be the limiting distribution of relative proximity between data points under the minor assumptions of continuity and smoothness.

The EVM formulation developed herein stems from the concept of *margin distributions*. This idea is not new; various definitions and uses of margin distributions have been explored [30]–[34], involving techniques such as maximizing the mean or median margin, taking a weighted combination margin, or optimizing the margin mean and variance. Leveraging the margin distribution itself can provide better error bounds than those offered by a soft-margin SVM classifier, which in some cases translates into reduced experimental error. We model  $\Psi$  in terms of the distribution of sample half-distances relative to a reference point, extending margin distribution theory from a per-class formulation [30]–[34] to a sample-wise formulation. The model is fit on the distribution of margins — half distances to the nearest negative samples — for each positive reference point. From this distribution, we derive a radial inclusion function which carves out a posterior probability of association with respect to the reference point. This radial inclusion function falls toward zero at the margin.

#### A. Probability of Sample Inclusion

To formalize the  $\Psi$ -model, let  $x \in \mathcal{X}$  be training samples in a feature space  $\mathcal{X}$ . Let  $y_i \in \mathcal{C} \in \mathbb{N}$  be the class label for  $x_i \in \mathcal{X}$ . Consider, for now, only a single positive instance  $x_i$  for some class with label  $y_i$ . Given  $x_i$ , the maximum margin distance would be given by half the distance to the closest training sample from a different class. However, the closest point is just one sample and we should consider the potential maximum margins under different samplings. We define margin distribution as the distribution of the margin distances of the observed data. Thus, given  $x_i$  and  $x_j$ , where  $\forall j, y_j \neq y_i$ , consider the margin distance to the decision boundary that would be estimated for the pair  $(x_i, x_j)$  if  $x_j$  were the closest instance. The margin estimates are thus  $m_{ij} = \|x_i - x_j\|/2$  for the  $\tau$  closest points. Considering these  $\tau$  nearest points to the margin, our question then becomes: what is the distributional form of the margin distances?

To estimate this distribution, we turn to the Fisher-Tippett Theorem [35] also known as the Extreme Value Theorem<sup>1</sup>. Just as the Central Limit Theorem dictates that the random variables generated from certain stochastic processes follow Gaussian distributions, EVT dictates that given a well-behaved overall distribution of values, *e.g.*, a distribution that is continuous and has an inverse, the distribution of the maximum or minimum values can assume only limited forms. To find the appropriate form, let us first recall:

<sup>1</sup>There are other types of extreme value theorems, *e.g.*, the second extreme value theorem, the Pickands-Balkema-de Haan Theorem [36], addresses probabilities conditioned on the process exceeding a sufficiently high threshold.

**Theorem 1** (Fisher-Tippett Theorem [37]). *Let  $(v_1, v_2, \dots)$  be a sequence of i.i.d samples. Let  $\zeta_n = \max\{v_1, \dots, v_n\}$ . If a sequence of pairs of real numbers  $(a_n, b_n)$  exists such that each  $a_n > 0$  and  $\lim_{z \rightarrow \infty} P\left(\frac{\zeta_n - b_n}{a_n} \leq z\right) = F(z)$  then if  $F$  is a non-degenerate distribution function, it belongs to the Gumbel, the Fréchet or the Reversed Weibull family.*

In other words for any sequence  $(a_n, b_n)$  of shifts and normalizations of the samples such that the probability of the maximum value converges, it converges to one of three distributions<sup>2</sup>. From Theorem 1, we can derive the following:

**Theorem 2** (Margin Distribution Theorem). *Assume we are given a positive sample  $x_i$  and sufficiently many negative samples  $x_j$  drawn from well-defined class distributions, yielding pairwise margin estimates  $m_{ij}$ . Assume a continuous non-degenerate margin distribution exists. Then the distribution for the minimal values of the margin distance for  $x_i$  is given by a Weibull distribution.*

*Proof.* Since Theorem 1 applies to maxima, we transform the variables via  $z = -m_{ij}$  and consider the maximum set of values  $-m_{ij}$ . The assumption of sufficient samples and a well-defined set of margin distances converging to a non-degenerate margin implies that Theorem 1 applies. Let  $\phi$  be the associated distribution of the maxima of  $\zeta_n$ . Combining Theorem 1 with knowledge that the data are bounded ( $-m_{ij} < 0$ ) means that  $\phi$  converges to a reversed Weibull, as it is the EVT distribution that is bounded from above [37]. Changing the variables back ( $m_{ij} = -z$ ) means that the minimum distance to the boundary must be a Weibull distribution.  $\square$

Theorem 2 holds for any point  $x_i$ , with each point estimating its own distribution of distance to the margin yielding:

**Corollary 1** ( $\Psi$  Density Function). *Given the conditions for the Margin Distribution Theorem, the probability that  $x'$  is included in the boundary estimated by  $x_i$  is given by:*

$$\Psi(x_i, x'; \kappa_i, \lambda_i) = \exp\left(-\left(\frac{\|x_i - x'\|}{\lambda_i}\right)^{\kappa_i}\right) \quad (1)$$

where  $\|x_i - x'\|$  is the distance of  $x'$  from sample  $x_i$ , and  $\kappa_i, \lambda_i$  are Weibull shape and scale parameters respectively obtained from fitting to the smallest  $m_{ij}$ .

*Proof.* The Weibull cumulative distribution function (CDF)  $F_W(\|x' - x_i\|; \kappa_i, \lambda_i)$  provides the probability that the margin is at or below a given value, but we seek the probability that  $x'$  does not exceed the margin yielding the inverse:  $1 - F_W(\|x' - x_i\|; \kappa_i, \lambda_i)$  (cf. Eq. 1).  $\square$

The  $\Psi$ -model defines a radial inclusion function that is an EVT rejection model where the probability of inclusion corresponds to the probability that the sample *does not* lie well into or beyond the negative margin. While  $\Psi$  is designed to have zero probability around the margin, half-way to the negative data, the model still supports a soft margin because the EVT estimation uses  $\tau$  points and hence may cover space with both positive and negative points. However, this does not force the model to include *any* positive training samples within its probability of inclusion.

<sup>2</sup>A more thorough introductory overview of EVT can be found in [38].

## B. Decision Function

The probability that a query point  $x'$  is associated with class  $C_l$  is  $\hat{P}(C_l|x') = \operatorname{argmax}_{\{i:y_i=C_l\}} \Psi(x_i, x'; \kappa_i, \lambda_i)$ . Given  $\hat{P}$ , we compute the open set multi-class recognition result for  $x'$ . Let threshold  $\delta$  on the probability define the boundary between the set of known classes  $\mathcal{C}$  and unsupported open space [9] so that the final classification decision is given by:

$$y^* = \begin{cases} \operatorname{argmax}_{l \in \{1, \dots, M\}} \hat{P}(C_l|x') & \text{if } \hat{P}(C_l|x') \geq \delta \\ \text{"unknown"} & \text{Otherwise.} \end{cases} \quad (2)$$

A principled approach to selecting  $\delta$  is to optimize the tradeoff between closed set accuracy and rejection of unknown classes via *cross-class validation* [5] on the training set, leaving out a subset of classes as "unknown" at each fold. A slight generalization of the decision function in Eq. 2 is to average over the  $k$ -largest probabilities for each class. For all experiments in Sec. V we selected a  $k$  value via a hyperparameter search on non-test data, choosing  $k$  from  $\{1, \dots, 10\}$ . In practice, we found that a choice of  $k > 1$  yields only slight performance gains of 1-2% in accuracy.

## IV. EVM FORMULATION

With the  $\Psi$ -models, we can develop an algorithm that is not only advantageous for open world recognition, but is also useful for limiting trained model size and obtaining favorable scaling characteristics. The pseudocode for this algorithm is provided in the supplemental material for this paper. Corresponding source code will be made available after publication.

### A. Model Reduction

Keeping all  $\Psi$ -models and associated data points results in larger models and longer classification times as dataset sizes increase, which is undesirable in both incremental and resource constrained scenarios. The success of sparse classification algorithms in other problem domains (*e.g.*, SVM) suggests that we can strategically discard many redundant  $\langle x_i, \Psi(x_i, x', \kappa_i, \lambda_i) \rangle$  pairs within a class  $C_l$  of  $N_l$  training points with minimal degradation in classification performance. Intuitively, if many points that characterize the class in question are sufficiently close to one another compared to points from negative classes, then we expect redundancy in their  $\Psi$  responses. By thresholding on a minimum redundancy probability, we can select a subset of points that characterize the class. While many strategies can be used for this selection, we wish to select the minimum number of points required to cover the class. We can formulate this strategy as a minimization problem: Let  $x_i$  be a point in the class of interest and  $\Psi(x_i, x', \kappa_i, \lambda_i)$  be its corresponding model. Without loss of generality, let  $x_j$  be another point in the same class with model  $\Psi(x_j, x', \kappa_j, \lambda_j)$ . Let  $\varsigma$  be the probability threshold above which to designate redundancy of the pair  $\langle x_j, \Psi(x_j, x', \kappa_j, \lambda_j) \rangle$  with respect to  $\langle x_i, \Psi(x_i, x', \kappa_i, \lambda_i) \rangle$ , such that if  $\Psi_i(x_i, x_j, \kappa_i, \lambda_i) \geq \varsigma$  then  $\langle x_j, \Psi(x_j, x', \kappa_j, \lambda_j) \rangle$  is redundant with respect to  $\langle x_i, \Psi(x_i, x', \kappa_i, \lambda_i) \rangle$ . Finally, let  $I(\cdot)$  be an indicator function such that

$$\begin{cases} I(x_i) = 1 & \text{if } \langle x_i, \Psi(x_i, x', \kappa_i, \lambda_i) \rangle \text{ kept} \\ I(x_i) = 0 & \text{otherwise.} \end{cases} \quad (3)$$

If  $x_i$  and  $\Psi(x_i, x', \kappa_i, \lambda_i)$  are retained, they become *extreme vectors* defining the final model. We can then express our optimization strategy in terms of this objective function:

$$\text{minimize } \sum_{i=1}^{N_l} I(x_i) \quad \text{subject to} \quad (4)$$

$$\forall j \exists i | I(x_i) \Psi(x_i, x_j, \kappa_i, \lambda_i) \geq \varsigma. \quad (5)$$

The constraint (Eq. 5) requires that every  $\langle x_i, \Psi(x_i, x', \kappa_i, \lambda_i) \rangle$  pair be an EV or be covered by at least one other pair. Note that the implicit binary constraint in the range of  $I(\cdot)$  makes the optimization an integer linear programming problem. The formulation in Eqs. 4 and 5 is a special case of Karp's Set Cover problem. We can see this by defining a coverage set of indices  $s_i \equiv \{j \in \{1, \dots, N_l\} | \Psi_i(x_i, x_j, \kappa_i, \lambda_i) \geq \varsigma\}$  for each  $\langle x_i, \Psi_i(x_i, x', \kappa_i, \lambda_i) \rangle$  pair and a universe  $U = \{1, \dots, N_l\}$ . The objective of the Set Cover problem is then to select the minimum number of sets that contains all elements of  $U$ . While Set Cover is NP-hard, we employ the greedy approximation described in [39] that offers a polynomial time  $(1 + \ln(N_l))$  approximate solution (*cf.* Theorem 2 in [39]). This algorithm offers the smallest error bound for any polynomial time approximation. The greedy approximation entails selecting the sets of highest cardinality at each step. The upper bound in approximation error is  $(1 - o(1)) \cdot \ln(N_l)$ , where  $N_l$  is the cardinality of the universe of set elements.

Note that with the model fitting, an outlier is generally covered by a point from another class (see Fig. 1), and such outlier points are also unlikely to cover many other points. Thus, outliers are added to the coverage set very late, if at all. This is not an *ad hoc* assumption; it is an outcome of the process of minimizing the number of points that cover all examples. Like the inherent softness of the margin, this is an inherent part of the model-reduction approach that follows from the EVT-modeling.

### B. Incremental Learning

Once EVs have been acquired for the EVM, the model can be updated with a new batch of data by fitting  $\Psi$ -models for the new data using both the current EVs and all points in the new batch of data. The new EVs are obtained by running model reduction over both the old EVs and the new training points. While new points can be added individually, adding data in batches will result in more meaningful fits because batches represent a richer distributional sample at each increment. This means that newly added training points may or may not become EVs and new classes can also impact previously learned models and EVs. The efficient model reduction technique discussed in Sec. IV-A allows the EVM to limit the size of its models either probabilistically via an explicit selection of  $\varsigma$  or by a specific maximum number of EVs in a max- $k$  cover greedy approach [40]. This allows the EVM to scale to many different incremental problems via different modes of operation. In Sec. V-B, for example, we choose a static  $\varsigma$  and perform model reduction using this threshold at each training increment for classes to which data get added. While this limits the growth in model size, the number of EVs still

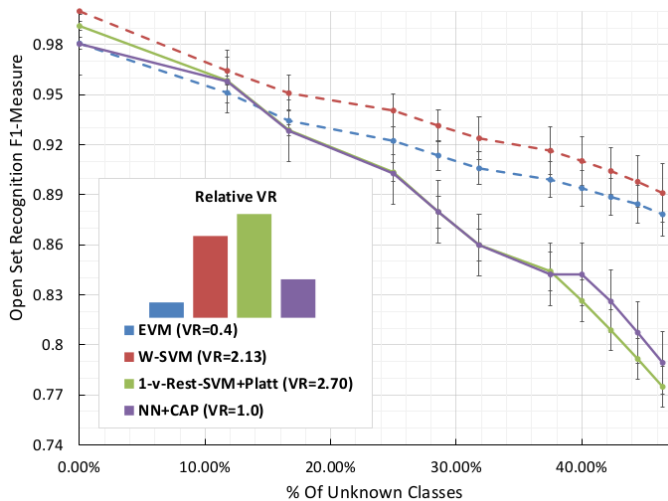


Fig. 2. *Multi-class open set recognition performance on OLETTER.* The  $x$ -axis represents the percentage of classes in the test set that were unseen during training. Error bars show standard deviation. The EVM is comparable to the existing state-of-the-art W-SVM [12] in F1-Measure, but at a substantial savings in training efficiency and scalability, as reflected in the vector ratio (VR). The EVM’s VR is an order of magnitude smaller than the two SVM-based models. Both EVM and W-SVM algorithms have favorable performance degradation characteristics as the problem becomes more open. The two other probabilistically calibrated algorithms degrade far more rapidly. Hyperparameters of  $\tau = 75$  (tail size) and  $k = 4$  (number of EVs to average over) were used in this evaluation, and were selected using the same training set cross validation technique as the W-SVM.

increases over time. Bounded optimization would specify a maximum per-class size or maximum total size, recalculating  $\zeta$  at each increment. Alternatively, maximum model sizes could be pre-specified with model reduction performed only when the maximum size is violated. Thus, the EVM is not only an incremental classifier, but it is an incremental classifier whose size can be controlled at each learning increment.

## V. EXPERIMENTAL EVALUATION

We conducted several evaluations in which we compared the EVM to other open set and open world classifiers on published benchmarks, including the state-of-the-art for both problems. To ensure valid comparisons and consistency with previous research where applicable, we report the results of these evaluations using the same evaluation measures and thresholds that were used in the original benchmark settings. For all of the experiments in this section we selected a  $\zeta$  value of 0.5 based on the probabilistic assumption that points with greater than 50% probability of being covered by others in that class are redundant with respect to the model.

### A. Multi-class Open Set Recognition on Letter

To establish the ability of the the EVM to identify unknown classes, we first looked at the problem of multi-class open set recognition. Scheirer et al. developed the OLETTER protocol [12] to evaluate classifiers in an open set recognition context. It consists of randomly selecting 15 distinct labels from the Letter dataset [41] as known classes during training and adding unknown classes by incrementally including subsets of

the remaining 11 labels during testing. This process is repeated over 20 folds to calculate averages and error bars. For consistency with [12], we report results in terms of F1-Measure (over precision and recall), and dynamically set a threshold on open space  $\delta = \frac{1}{2} \left( 1 - \sqrt{(2 \times C_T / (C_R + C_E))} \right)$ , where  $C_R$  is the number of classes to be recognized (common to training and testing),  $C_T$  is the number of classes used in training, and  $C_E$  is the number of classes used in evaluation (testing).

For the OLETTER evaluation, we used Euclidean distance to compute the margins for our EVM. We evaluate model compactness in terms of the *vector ratio*, defined as  $VR = \frac{\# \text{ points retained by model}}{\text{total } \# \text{ training points}}$ . The VR is a scaled form of the support vector ratio introduced by Vapnik to provide an approximation of generalization error [42], [43]. This allows us to compare the scalability of different nonlinear models.

Fig. 2 shows results for all of the evaluated algorithms, including the open set-specific W-SVM [12], which is currently the best performing algorithm in the literature for this problem. Results for Nearest Neighbor (NN) classifiers with CAP probability estimation [12], and 1-vs-Rest RBF SVMs with Platt’s probability estimator [44] are also shown. Other calibrated models assessed in [12] performed significantly worse and are not shown. We selected RBF parameters for the SVMs via 5-fold cross validation on the training dataset, using a grid of  $C = 2^{-5}, 2^{-3}, \dots, 2^{15}$  and  $\gamma = 2^{-15}, 2^{-13}, \dots, 2^3$ , consistent with [45].

The EVM performs comparably to the W-SVM, and outperforms all other algorithms. The W-SVM is certainly a viable algorithm for this dataset, but its slight advantage comes at a greater cost than the EVM, requiring *two* trained SVM models (per-class: one 1-class and one binary) for its operation. The vector ratio for this experiment is computed for models trained on all 26 classes. For the evaluation in Fig. 2, the EVM’s vector ratio is an order of magnitude smaller than that of any of the SVM models, indicating that for the chosen  $\zeta$  (0.5), fewer than half of the training data points were included as EVs. The number of support vectors in the SVM models is greater than the number of points in the training set. This is due to redundancy in support vectors kept in the multi-class regime. Although Platt-Calibrated SVM model processing and storage costs can be reduced by caching duplicated support vectors, the computational savings is less feasible for the W-SVM, since it uses different SVM models with multiple kernels. This is because different RBF kernels require different calculations even if they are centered on the same point. Also, for the EVM, and unlike the W-SVM, we can easily obtain a lower vector ratio while minimizing any degradation in accuracy. We analyze this tradeoff in Sec. VI. Finally, we would like to mention that, apart from the EVM, none of the classifiers whose performance is depicted in Fig. 2 support incremental learning, so they cannot be applied to open world problems.

### B. Open World Recognition on ImageNet

Open world recognition, the problem we are most concerned with, consists of three distinct phases [5]: one initial training phase, followed by alternating phases of open set recognition and updates to incorporate newly labeled data into the model. We evaluated open world recognition performance on



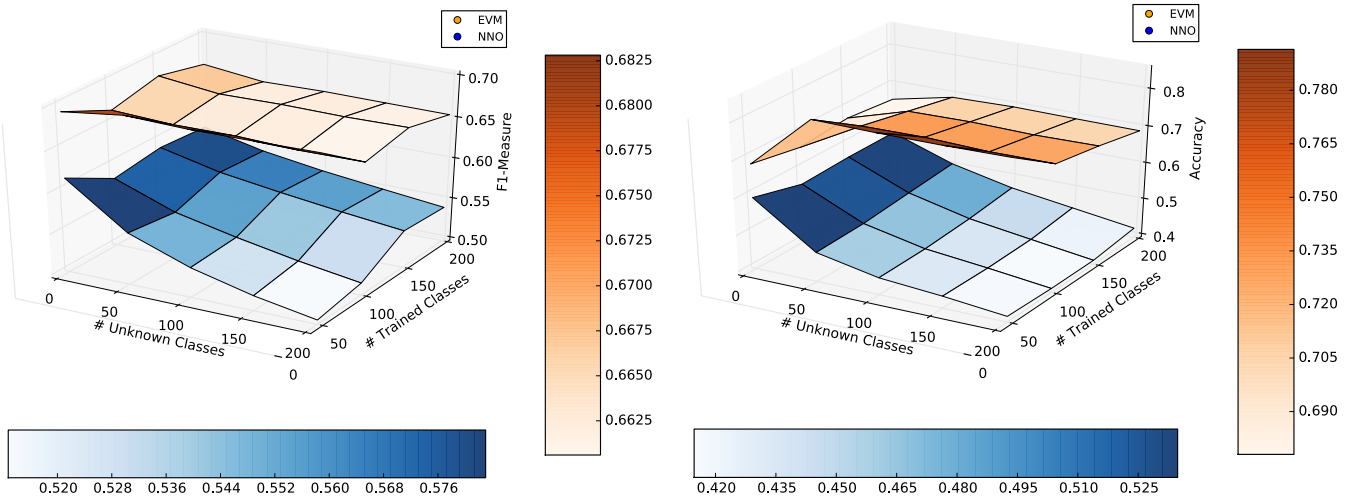


Fig. 3. Open world performance of the EVM and NNO algorithms on the open world ImageNet benchmark in terms of both F1-measure (left) and accuracy (right). Initial training considering 50 classes was performed, then classes were incrementally added in groups of 50. The EVM dramatically outperforms NNO at all points evaluated and the divergence of the respective surfaces suggests superior scalability with respect to both the number of additional training classes and the number of additional unknown classes. Hyperparameter values of  $k = 6$  and  $\tau = 33998$  were selected using the cross validation procedure discussed in the text.

a benchmark protocol that uses an open world partition of ImageNet introduced in [5]. The protocol consists of an initial training phase on 50 classes, performing classification with 0, 50, 100, 150, and 200 unknown classes in the test set. At each increment another group of 50 classes is added, and classification is again performed on test instances from known classes, with samples from 0, 50, 100, 150, and 200 additional unknown classes. One manner in which we depart from the protocol in [5] is that instead of using dense SIFT features, which are no longer state-of-the-art, we use a 4,096-dimensional deep feature space representation derived from the fc7 layer of AlexNet [46]. This is a much better representation for today’s visual recognition tasks.

We compared the EVM’s computational performance with the state-of-the-art NNO algorithm [5] on this benchmark, as well as the incremental KDE classification algorithm employed by Haines and Xiang in [22] that had heretofore not been tested on this benchmark, or any particularly large-scale benchmark, in terms of either number of samples or dimensionality. While both the EVM and the NNO algorithms completed in a matter of hours, the incremental KDE classifier had enrolled fewer than 40 samples after 24 hours. Assuming a constant rate of enrollment, it would take approximately 18 years for the benchmark training experiment to complete, leading us to conclude that the approach, at least as implemented, is not scalable, and therefore report recognition results for just the EVM and NNO algorithms.

For training and testing, we used the ILSVRC2014 training and validation partitions respectively. For consistency with [5], we selected an NNO rejection threshold via 3-fold cross-class validation on the first 50 classes from the training set. This is also how we selected  $\delta$  for the EVM, searching over a range of  $[0.05, 0.1, \dots, 0.3]$ . Hyperparameters  $k$  and  $\tau$  were obtained via a Bayesian search using the hyperopt library [47]. The Bayesian search conducted with hyperopt consisted of 50

iterations of three-fold cross-validation over the first 50 classes of the training set. The hyperparameter ranges consisted of 1-10 for  $k$  and 100-32,000 for  $\tau$ . Noting that 3-fold cross validation reduces the size of the training set by 1/3, and assuming a rough proportionality on these hyperparameters to training set size (and increment batch size), we multiplied the selected values by 1.5 and rounded to the nearest integer value to arrive at hyperparameter selections used for training.

During our ImageNet experiments, we found that Euclidean distance led to poor performance when computing margins for the EVM. This is consistent with the previous finding that Euclidean distance does not work well when comparing deep features of individual samples [48]. We therefore turned to cosine similarity, which is a commonly used measure of divergence between samples in a deep feature space. NNO still performed reasonably well when using Euclidean distance; which we suspect is because it compares samples with respect to class means. However, when cosine similarity was used, the classifier rejected nearly all test samples as unknown. Thus, we report the best results for each classifier, using cosine similarity for the EVM and Euclidean distance for NNO. These results are shown in Fig. 3.

The EVM consistently and dramatically outperforms NNO, both in terms of accuracy and F1-Measure. Readers might wonder why the EVM’s accuracy increases as the number of classes are added; this is because the EVM is a good rejector and is able to tightly bound class hypotheses by their support. However, it does so while simultaneously maintaining reasonable classification performance. While NNO models the deviation from class means according to a support bound, such a rigid and over-simplified model of each class yields a poorer performance trend than the EVM delivers, hence the overall divergence of the surface plots in Fig.3. Finally, we note that with the chosen value  $\zeta = 0.5$ , the EVM provides significant reductions in model size over using all points (cf. Table I).

TABLE I  
NUMBERS OF EXTREME VECTORS, CUMULATIVE NUMBERS OF TRAINING POINTS USED, AND VECTOR RATIOS AFTER EACH BATCH IS ADDED.

Batch	0-50	50-100	100-150	150-200
# EVs	16309	35213	52489	74845
# Points	64817	129395	194217	255224
Vector Ratio	0.25	0.27	0.27	0.29

## VI. PRACTICAL CONSIDERATIONS FOR THE EVM

The model reduction strategy that we employed in Sec. V of selecting a threshold  $\zeta$  and running the Set Cover algorithm is a simple way to increase efficiency at classification time and achieve a compact representation of the training set. While our experiments in Sec. V used a probabilistically motivated choice of  $\zeta = 0.5$ , what constitutes a good redundancy threshold in practice is often tied to the computational, storage, or time budget of the problem at hand. The model must therefore be sized accordingly so that it does not exceed the maximum budget, while still maintaining as good performance as possible. We refer to this as the *budgeted optimization* problem in which the objective is to obtain the most representative model achievable that meets but does not exceed the budget. We can perform this selection via a binary search for  $\zeta$ , for which the optimization, given a target budget, most closely returns the requested number of EVs. Since the greedy optimization selects EVs in order of their coverage, we can easily retain only the most important of these EVs. This allows EVM classifiers to be approximately portable across many device types of heterogeneous computational capabilities.

We performed an evaluation of this technique on the closed set Letter dataset, using all points in the training set ( $\zeta = 1.0$ ), which yielded a base accuracy of 96%. Using 50% of the data ( $\zeta = 0.492$ ) or 40% of the data ( $\zeta = 0.186$ ) yielded accuracies comparable to using all points. Reducing to 10% of the training points ( $\zeta = 0.008$ ) yielded an accuracy of 92%. This suggests that budgeted optimization is quite effective for classifier compression/portability, and that  $\zeta$  can assume a very wide range with minimal impact on classification performance.

With respect to computational efficiency, much of the EVM’s training procedure can be performed independently on a class-by-class basis, making the algorithm well suited for a cluster or GPU implementation. Each statistical fit, made via Maximum Likelihood Estimation (MLE), is fast and constant in time complexity, due to a cap on the maximum number of iterations. Model reduction is  $O(N_l^2)$ , where  $N_l$  is the number of points in class  $C_l$ . The complexity of each tail retrieval, given  $N$  training points, can be reduced from  $O(N \log N)$  to  $O(\tau \log N)$  by introducing space partitioning data structures, e.g.,  $k$ - $d$  trees [49]. However, employing them may impose constraints on the types of measures used to compute the margin — the common requirement that the distance function be a strict metric on the hypothesis space precludes quasi-metrics such as cosine similarity.

While the EVM may appear to be highly parameterized, **the per point Weibull scale and shape parameters are purely data derived and are automatically learnt during training.** They are a function of the MLE optimization process, which

we must distinguish from hyperparameters selected prior to fitting. The only hyperparameters are  $\tau$ ,  $k$ , and  $\zeta$ . As we have previously discussed, barring hard model size constraints (e.g., in budgeted optimization), a wide range of cover probability  $\zeta$  values work well, including the probabilistically driven choice of 0.5 used in the experiments in Sec. V. With respect to the value  $k$ , the number of models to average, we found that searching over  $\{1, \dots, 10\}$  on a validation set yielded slightly better results on the test set (with further decreased performance for  $k > 10$ ), although performance variations over this parameter range typically accounted for less than 2% in accuracy or  $F1$ -measure respectively. Thus, while  $\zeta$  and  $k$  might have a slight impact on performance, the vast majority of performance variation can be attributed to only a single hyperparameter, namely  $\tau$ .

Unfortunately, EVT provides no principled means of selecting the tail size  $\tau$ . The theory only dictates the family of distribution functions that will apply and proves convergence for an unspecified  $\tau$ . We use cross-class validation for selecting  $\tau$  that yields optimal cross-validation accuracy with missing classes. For the selection of  $\tau$  obtained on the OLETTER dataset through cross-class validation, the value  $\tau = 75$  accounts for only 0.5% of the 15K OLETTER training set and that small a fraction represents plausible extrema. However, the results of applying the cross-class validation methodology discussed in Sec. V-B for our ImageNet experiments might raise questions since the selected tail size ( $\tau = 33998$ ) consists of approximately half the number of samples in the training set. This result is surprising and counter-intuitive because one would not ordinarily think of “half of the data” as being extreme values. But looking at a fraction of data is probably the wrong way to determine the boundary or extreme points. While the EVM’s  $\Psi$ -models are fit on 1D margins, the high dimensionality of the feature space (4,096) translates into many more directions yielding “boundary points” than for the relatively low-dimensional (16) OLETTER dataset. Normalizing for feature dimensionality and the number of classes we find both examples are similar. Dividing the tail size used for the first batch of ImageNet by the dimensionality of the feature space and number of classes yields 0.17 points per dimension per class. Doing the same for OLETTER, we obtain 4.69 points per dimension and an average of 0.17 points per dimension per class.

## VII. CONCLUSION

Perhaps the most important conclusion of this work is that the EVM is able to do kernel-free nonlinear classification. Interestingly, the EVM shares some relationships with radial basis functions. When  $\kappa = 2$ , the functional form of Eq. 1 is the same as a Gaussian RBF, and when  $\kappa = 1$  it is the same as an exponential or Laplacian RBF. While these  $\kappa$  values can occur in practice,  $\kappa$  assumes a much broader range of values, which are generally larger. Furthermore for  $\kappa > 2$ , Eq. 1 is not a Mercer kernel. Alternatively, if one approximates Eq. 1 by a weighted sum of Gaussians (or Laplacians) we have two different ways of viewing a Gaussian (or Laplacian) RBF kernel as an approximation of a  $\Psi$ -model. While the  $\Psi$ -model



parameters vary in scale and shape with the bandwidth and density of the data set, in a Gaussian approximation the number of kernel elements and/or the accuracy of approximation must vary spatially. The EVM requires the fewest points for the margin distribution and its  $\Psi$ -model. **For the EVM, we do not make an *ad hoc* assumption of a kernel trick nor a *post hoc* assumption of a particular kernel function; the functional form of the  $\Psi$ -model is a direct result of EVT being used to model input space distance distributions.**

The Weibull fitting ensures that a small number of mislabeled points or other outliers will not cause the margin estimated from the Weibull to be at that location. If the fitting includes more distant points, the  $\Psi$ -model will broaden in scale / shape providing a naturally derived theory for the “softness” in its margin definition. However, the overall optimization with Set Cover currently lacks a parameter to adjust the risk tradeoff between positive and negative classes. Future directions of research may include directly extending the EVM by obtaining a better parameterized soft-margin during Set Cover, perhaps by adding weights to balance soft-margin errors and formulating the problem in terms of linear programming. Another potential extension would be to incorporate margin weights in a loss function in an SVM-style optimization algorithm.

#### REFERENCES

- [1] K. Crammer, O. Dekel, J. Keshet, S. Shalev-Shwartz, and Y. Singer, “Online passive-aggressive algorithms,” *JMLR*, vol. 7, pp. 551–585, 2006.
- [2] T. Yeh and T. Darrell, “Dynamic visual category learning,” in *IEEE CVPR*, 2008.
- [3] P. Laskov, C. Gehl, S. Krüger, and K.-R. Müller, “Incremental support vector learning: Analysis, implementation and applications,” *JMLR*, vol. 7, pp. 1909–1936, 2006.
- [4] L.-J. Li and L. Fei-Fei, “Optimol: automatic online picture collection via incremental model learning,” *IJCV*, vol. 88, no. 2, pp. 147–168, 2010.
- [5] A. Bendale and T. Boulton, “Towards open world recognition,” in *IEEE CVPR*, 2015.
- [6] A. Kapoor, S. Baker, S. Basu, and E. Horvitz, “Memory constrained face recognition,” in *IEEE CVPR*, 2012.
- [7] T. Mensink, J. Verbeek, F. Perronnin, and G. Csurka, “Metric learning for large scale image classification: Generalizing to new classes at near-zero cost,” in *ECCV*, 2012.
- [8] M. Ristin, M. Guillaumin, J. Gall, and L. Gool, “Incremental learning of ncm forests for large-scale image classification,” in *IEEE CVPR*, 2014.
- [9] W. J. Scheirer, A. Rocha, A. Sapkota, and T. E. Boulton, “Towards open set recognition,” *IEEE T-PAMI*, vol. 36, July 2013.
- [10] C. M. Bishop, *Pattern Recognition and Machine Learning*. Springer Science+Business Media, LLC, 2006.
- [11] L. P. Jain, W. J. Scheirer, and T. E. Boulton, “Multi-class open set recognition using probability of inclusion,” in *ECCV*, September 2014.
- [12] W. Scheirer, L. Jain, and T. Boulton, “Probability models for open set recognition,” *IEEE T-PAMI*, vol. 36, pp. 2317–2324, 2014.
- [13] M. Markou and S. Singh, “Novelty detection: a reviewpart 1: statistical approaches,” *Signal Processing*, vol. 83, no. 12, pp. 2481–2497, 2003.
- [14] V. J. Hodge and J. Austin, “A survey of outlier detection methodologies,” *Artificial Intelligence Review*, vol. 22, no. 2, pp. 85–126, 2004.
- [15] H. Cevikalp and B. Triggs, “Efficient object detection using cascades of nearest convex model classifiers,” in *IEEE CVPR*, 2012.
- [16] T. Mensink, J. Verbeek, F. Perronnin, and G. Csurka, “Distance-based image classification: Generalizing to new classes at near-zero cost,” *IEEE T-PAMI*, vol. 35, no. 11, pp. 2624–2637, 2013.
- [17] P. Bodesheim, A. Freytag, E. Rodner, and J. Denzler, “Local novelty detection in multi-class recognition problems,” in *IEEE WACV*, 2015.
- [18] R. Gopalan, R. Li, V. M. Patel, and R. Chellappa, “Domain adaptation for visual recognition,” *Foundations and Trends in Computer Graphics and Vision*, vol. 8, no. 4, pp. 285–378, 2015.
- [19] C. H. Lampert, H. Nickisch, and S. Harmeling, “Attribute-based classification for zero-shot visual object categorization,” *IEEE T-PAMI*, vol. 36, no. 3, pp. 453–465, 2014.
- [20] A. Pentina and S. Ben-David, “Multi-task and lifelong learning of kernels,” in *Int. Conf. on Algorithmic Learning Theory*, 2015.
- [21] A. Royer and C. H. Lampert, “Classifier adaptation at prediction time,” in *IEEE CVPR*, 2015.
- [22] T. Haines and T. Xiang, “Active rare class discovery and classification using dirichlet processes,” *IJCV*, vol. 106, no. 3, pp. 315–331, 2014.
- [23] W. J. Scheirer, A. Rocha, R. J. Micheals, and T. E. Boulton, “Meta-recognition: The theory and practice of recognition score analysis,” *IEEE T-PAMI*, vol. 33, no. 8, pp. 1689–1695, 2011.
- [24] W. J. Scheirer, A. Rocha, R. Michaels, and T. E. Boulton, “Robust fusion: Extreme value theory for recognition score normalization,” in *ECCV*, 2010.
- [25] W. J. Scheirer, N. Kumar, P. N. Belhumeur, and T. E. Boulton, “Multi-attribute spaces: Calibration for attribute fusion and similarity search,” in *IEEE CVPR*, 2012.
- [26] X. Gibert, V. Patel, and R. Chellappa, “Sequential score adaptation with extreme value theory for robust railway track inspection,” in *IEEE ICCV Workshops*, 2015.
- [27] V. Fragoso and M. Turk, “SWIGS: A swift guided sampling method,” in *IEEE CVPR*, Portland, Oregon, June 2013.
- [28] V. Fragoso, P. Sen, S. Rodriguez, and M. Turk, “EVSAC: Accelerating hypotheses generation by modeling matching scores with extreme value theory,” in *IEEE ICCV*, December 2013.
- [29] H. Zhang and V. Patel, “Sparse representation-based open set recognition,” *IEEE T-PAMI*, 2017.
- [30] J. Shawe-Taylor and N. Cristianini, “Margin distribution and soft margin,” *Advances in Large Margin Classifiers*, pp. 349–358, 2000.
- [31] A. Garg and D. Roth, “Margin distribution and learning algorithms,” in *ICML*, 2003.
- [32] L. Reyzin and R. Schapire, “How boosting the margin can also boost classifier complexity,” in *ICML*, 2006.
- [33] F. Aioli, G. Da San Martino, and A. Sperduti, “A kernel method for the optimization of the margin distribution,” in *ICANN*, 2008.
- [34] K. Pelckmans, J. Suykens, and B. Moor, “A risk minimization principle for a class of parzen estimators,” in *NIPS*, 2008.
- [35] R. Fisher and L. Tippett, “Limiting forms of the frequency distribution of the largest or smallest member of a sample,” in *Mathematical Proc. of the Cambridge Philosophical Society*, vol. 24, no. 2. Cambridge University Press, 1928, pp. 180–190.
- [36] J. Pickands, “Statistical Inference Using Extreme Order Statistics,” *The Annals of Statistics*, vol. 3, no. 1, pp. 119–131, 1975.
- [37] S. Kotz and S. Nadarajah, *Extreme Value Distributions: Theory and Applications*. World Sci. Pub. Co., 2001.
- [38] S. Coles, *An Introduction to Statistical Modeling of Extreme Values*. Springer, 2001.
- [39] P. Slavik, “A tight analysis of the greedy algorithm for set cover,” in *ACM symposium on theory of computing*, 1996.
- [40] B. Mirzasoleiman, A. Karbasi, R. Sarkar, and A. Krause, “Distributed submodular maximization: Identifying representative elements in massive data,” in *NIPS*, 2013.
- [41] P. W. Frey and D. J. Slate, “Letter recognition using Holland-style adaptive classifiers,” *Machine Learning*, vol. 6, no. 2, pp. 161–182, 1991.
- [42] V. Vapnik, *Statistical learning theory*. Wiley New York, 1998, vol. 1.
- [43] G. Valentini and T. G. Dietterich, “Bias-variance analysis of support vector machines for the development of SVM-based ensemble methods,” *JMLR*, vol. 5, pp. 725–775, 2004.
- [44] J. Platt, “Probabilistic outputs for support vector machines and comparison to regularization likelihood methods,” in *Advances in Large Margin Classifiers*, A. Smola, P. Bartlett, B. Schoelkopf, and D. Schuurmans, Eds., 2000, pp. 61–74.
- [45] C.-W. Hsu, C.-C. Chang, C.-J. Lin *et al.*, “A practical guide to support vector classification,” 2003.
- [46] A. Krizhevsky, I. Sutskever, and G. E. Hinton, “Imagenet classification with deep convolutional neural networks,” in *NIPS*, 2012.
- [47] J. Bergstra, D. Yamins, and D. D. Cox, “Hyperopt: A python library for optimizing the hyperparameters of machine learning algorithms,” in *SciPy*, 2013.
- [48] C. C. Aggarwal, A. Hinneburg, and D. A. Keim, “On the surprising behavior of distance metrics in high dimensional space,” in *Int. Conf. on Database Theory*. Springer, 2001, pp. 420–434.
- [49] J. L. Bentley, “Multidimensional binary search trees used for associative searching,” *Comm. of the ACM*, vol. 18, no. 9, pp. 509–517, 1975.