

## Automatic Open-World Reliability Assessment

Mohsen Jafarzadeh    Touqeer Ahmad    Akshay Raj Dhamija  
 Chunchun Li    Steve Cruz    Terrance E. Boult\*  
 University of Colorado, Colorado Springs  
 Colorado Springs, Colorado 80918, USA


{mjafarzadeh, tahmad, adhamija, cli, scruz, tboult}@vast.uccs.edu

### Abstract

Image classification in the open-world must handle out-of-distribution (OOD) images. Systems should ideally reject OOD images, or they will map atop of known classes and reduce reliability. Using open-set classifiers that can reject OOD inputs can help. However, optimal accuracy of open-set classifiers depend on the frequency of OOD data. Thus, for either standard or open-set classifiers, it is important to be able to determine when the world changes and increasing OOD inputs will result in reduced system reliability. However, during operations, we cannot directly assess accuracy as there are no labels. Thus, the reliability assessment of these classifiers must be done by human operators, made more complex because networks are not 100% accurate, so some failures are to be expected. To automate this process, herein, we formalize the open-world recognition reliability problem and propose multiple automatic reliability assessment policies to address this new problem using only the distribution of reported scores/probability data. The distributional algorithms can be applied to both classic classifiers with SoftMax as well as the open-world Extreme Value Machine (EVM) to provide automated reliability assessment. We show that all of the new algorithms significantly outperform detection using the mean of SoftMax.


### 1 Introduction

Many online businesses, especially shopping websites and social media, created a platform that users can upload an image, and it classifies the images into predefined classes, potentially open-set classifiers, which in addition to known classes in the training data set, have an out-of-distribution (OOD) class, also known as novel, garbage, unknown, reject, or background class, that the system otherwise ignores. The OOD class can be considered as a "background" separate class or as the origin of feature space [11]. Mapping



(a) SoftMax for 6 batches

(b) EVM for 6 batches



(c) Scores over different data partitions

**Fig 1:** Percentage of images that yield the given score/probability for either SoftMax or EVM. In (a) and (b), we show the percentage for six different batches of 100 images. Reliability is about to determine when the system's accuracy will be degraded by OOD samples. Can you determine which of the six batches (A–F) are "normal" and which have increased OOD data? As a hint, three are normal; the others have 3%, 5%, and 20% OOD data. See the text for answers. Open-set classifiers deal with OOD inputs rejecting inputs below a threshold on score/probability. Plot (c) shows the cumulative percentage of images above threshold for knowns and below threshold for unknowns, i.e., the percentage of data correctly handled at that threshold. Assessing a change in reliability, i.e., detecting a shift in the distribution of knowns and unknowns, is a non-trivial problem because of the overlap of scores on knowns and OOD images. This paper formalizes and explores solutions to the open-world reliability problem.

too many images to reject class or misclassifying them reduces users' satisfaction. But since there are no labels during operation, there is no obvious way to detect that the input distribution has changed. Therefore, online businesses use human laborers to monitor image classification quality and to fine-tune the system when the quality drops. Mon-

\*Supported in part by DARPA SAIL-ON Contract # HR001120C0055

itoring using human labor is expensive. In other applications, e.g., robotic or vehicular navigation, reduced reliability from increasing OOD inputs must be handled automatically as the processing cycle is too fast for human oversight. Fig. 1 shows score distributions and why thresholding can be problematic, and why detecting changes in the mix of known and OOD is difficult. Did you correctly guess that A, B, and C are the normal batches while D, E, and F have 3%, 5%, and 20% OOD data, respectively?

As we move toward the path of operational reliability, we start by adding robustness to close-set classifiers. We can further improve reliability by transitioning to open-set classifiers or even transition to open-world classifiers that incrementally add the OOD samples [5]. While there has been a modest amount of research looking at various definitions of classifier reliability determination (see Section 2) and defining open-set and open-world recognition algorithms, no prior work has addressed the problem of reliability when there are changes in the frequency of OOD inputs.

In this work, we take operational reliability one step further by defining and presenting solutions to the problem of automatically assessing classifier reliability in an open-set/open-world setting. Such assessments are necessary to determine when it may be time to fine-tune and incrementally learn new classes; however, the choice of thresholds and processes for incremental learning is domain-specific and beyond the scope of this paper.

Extreme value theory (EVT) is a branch of statistics that studies the behavior of extreme events, the tails of probability distributions [8, 1, 6]. EVT estimates the probability of events that are more extreme than any of the already observed ones. EVT is an extrapolation from observed samples to unobserved samples. The extreme value machine (EVM) [26, 16] uses EVT to compute the radial probability of inclusion of a point with respect to the class of another. Therefore, EVM has been successfully used for instance-level OOD detection.

While the EVM is an inherently open-set classifier, it is less widely used, and hence it is also natural to ask if we can detect when there is a shift in the distribution for a standard closed-set classifier and hence the need for retraining. In this paper, we show distribution-based algorithms, including Gaussian-model with Kullback–Leibler (KL) divergence applied to SoftMax scores or EVM scores, are effective in detection. We also develop a novel detection algorithm for EVM, which makes Bernoulli distributional assumptions and find it works equally as well. Thus, the proposed reliability algorithms can be used for both traditional "closed-set" and more recent "open-set" classifiers for open-world reliability assessment.

### Our Contributions

- Formalizing open-world reliability assessment.
- Proposing three new policies to automate open-world

reliability assessment of closed-set (SoftMax) and open-set image classifiers (EVM).

- Developing a test-bed to evaluate the performance of reliability assessment policies.
- Comparing proposed policies with simple approaches, showing a significant improvement over the base algorithm of tracking the mean of SoftMax scores.

## 2 Related work

The robustness of image classifiers has been studied for many years, with a known trade-off between the accuracy and robustness of image classifiers [28]. The effect of data transformation in robustness is demonstrated in [3]. Stability training is another way of increasing robustness [34]. Although, by improving robustness, an image classifier can increase reliability to variation of known classes; in practice, robustness degrades when the ratio of OOD images to other images increases.

When presented with an out-of-distribution input, most image classifiers map OOD to a known class. Recently introduced open-set classifiers [2, 26, 32, 24, 4, 13, 19], have developed improved approaches to reasonably detect if an input is from a novel class (a.k.a unknown or out-of-distribution) and map it to a reject or novel class label. A family of quasi-linear polyhedral conic discriminant is proposed in [7] to overcome label imbalance. An iterative learning framework for training image classifiers with open-set noisy labels is proposed in [31]. Networks can be trained for joint classification and reconstruction [32]. Class conditioned auto-encoders proposed may improve the accuracy of known while rejecting unknowns [24]. Policies for a guided exploration of the unknown in the open-world are proposed in [21]. The robustness of open-world learning has been studied in [27]. Some techniques, such as [26], allow incremental updating to support open-world learning of these new unknowns but provide no guidance on when the system's reliability has been reduced to the point of needing incremental updates. The quality of these open-set image classifiers is directly related to the ratio of known to unknown images. By decreasing this ratio, the open-set image classifiers should be fine-tuned to have better reliabilities on the data. Thus, businesses should invest in the reliability assessment policy.

There have been various studies of the reliability of classifiers [20, 35, 9] that looked at close-set classifiers' accuracy or error. These do not apply in the open-world setting. Still others considered the largest few scores (extreme values) in the close-set setting [36, 29, 15] weak approximation of using extreme value for open-set classification as in [2]. In [22], reliability is defined as the difference between the classification confidence values. Matiz and Barner defined reliability as confidence values of predictions [23]. In a published paper [18], authors showed that confidence is

not equal to reliability because image classifiers may put images in the wrong class with high confidence. Any reliability measure on a single input can be viewed as a type of classifier with rejection, which is a subset of open-set recognition algorithms.

Looking at sets of data, either in batches or moving windows, it is necessary to detect the change in the distribution associated with increasing OOD samples. In [25], authors defined reliability as an inverse to its standard deviation, i.e.,  $R = \sigma^{-1}$ . Zhang defined reliability as the sum of rejection rate and accuracy [33]. This definition is not acceptable because if the image classifier puts all images in the reject class, the rejection rate is 1.0, and consequently, the reliability is 1.0. In [18], they used a false positive rate of points where no desirable correct predictions are lost. Their definition of reliability also fails for the same reason, if an image classifier says every image is negative (reject class), then it has a reliability of 1.0.

Considering the data as a sequence in time, one can view the reliability assessment problem as a specialized type of distributional change detection. In the change detection literature, there is a wide range of algorithms, and an important algorithm is KL divergence [10].

Thus, while there has been a modest amount of research looking at various definitions of classifier reliability determination, no prior work addresses it from the point of view of open-set classifiers when there are changes in the frequency of OOD inputs.

### 3 Problem Formalization & Evaluation

British standards institution defined reliability as “the characteristic of an item expressed by the probability that it will perform a required function under stated conditions for a stated period of time” [12]. In this paper, we assume a well-trained network, such that per-class errors are relatively consistent across classes, i.e., the network is trained such that their errors are not a function of the classes. With that assumption and the standard definition, networks should be reasonably reliable to changes in sampling frequency between classes in their training distributions, so distributional shifts between such classes are not important. However, since OOD data is not something we can use in training, we cannot assume the uniformity of errors on OOD data and networks. Classifiers become unreliable if the frequency of samples from unseen classes change, and they increase their frequency of rejection. Thus, we seek to formalize reliability via changes in the distribution with respect to OOD samples.

#### Definition 1 Open-world Reliability Assessment

Let us define  $\mathcal{T}$  for the mixture of known and unknown classes seen in training, as well as  $\mathcal{U}$ , classes unseen in training. Let  $x_1 \dots x_N$ , be samples drawn from distribution

$\mathcal{D}_1$ , where  $x \in \mathcal{D}_1 \rightarrow x \in \mathcal{T}$ , i.e.  $\mathcal{D}_1$  is either consistent with training or some operational validation set and hence is by definition reliable. Let  $x_{N+1} \dots x_m$ , be samples drawn from another distribution  $\mathcal{D}_2$ . Let  $P_{\mathcal{U}}(x; \mathcal{D})$  be the probability, given distribution  $D$ , that  $x \in \mathcal{U}$ , i.e., the probability  $x$  is a class unseen in training. Finally, let  $\mathcal{M}$  be a dissimilarity measure between two probability distributions.

*Open-world Reliability Assessment* is determining if there is a change in the distribution of samples from the unknown or OOD classes, i.e.

$$\mathcal{M}(P_{\mathcal{U}}(x; \mathcal{D}_2); P_{\mathcal{U}}(x; \mathcal{D}_1)) > \delta \quad (1)$$

for some user-defined  $\delta$ . When the above equation holds, we say that  $\mathcal{D}_2$  has a change in reliability.

The choice of  $\mathcal{M}$  and  $\delta$  may be domain-dependent as the cost of errors may be influenced by the size and frequency of the distributional shifts. In this paper, we explore multiple models. As defined, this does not say if reliability is increasing or decreasing, just that it changed. However, with some measures, one can determine the direction of the change.

From an experimental point of view, we can introduce different distributions  $\mathcal{D}_2$  and measure detection rates based on the ground truth label of when it was introduced.

Given a sequence of samples, detection of the distributional change from reliable to unreliable is often an important operational criterion. We define reliability localization error as the absolute error in the reported location of reliability change, i.e., if the detection is reported at time  $n_r$  and ground truth is  $n_g$ , for a given trial absolute error is

$$e_l = |n_g - n_r| \quad (2)$$

and can consider the mean absolute error over many trials. We define *On-Time detection* as reporting the change with zero absolute error for a given trial, and early and late detection have a non-zero error with the appropriate sign.

To implement such an assessment, we need some way to estimate the probability of something being out of distribution. One simple and common approach is using classifier confidence, i.e., SoftMax, as in [22]. Better OOD detection approaches are based on the more formal open-set classifiers that have been developed, and for this work, we consider the model from the extreme value machine (EVM) [26, 16]. While EVM may be better for individual OOD detection, we show the proposed reliability assessment approaches works equally well for both and expect it will work well with building on top of any good per-instance estimate of the probability of input being out-of-distribution.

#### Evaluation Protocol

Similar to testing open-set classifiers, testing of open-world reliability assessment needs a mixture of in-distribution

classes that was used during training as well as out-of-distribution data. For measuring open-set evaluation, we recommend a large number of classes to ensure good features; a small number of classes produce more degenerate features from overfitting. For this evaluation, we used all 1000 classes of ImageNet 2012 for training and validation as known classes. Then, to create a test set, we used combinations of 1000 classes of ImageNet 2012 validation set as known classes and 166 non-overlapping classes of ImageNet 2010 training set as unknown classes. We note that while some earlier work, such as [2, 26], used 2012/2010 splits, they claimed 2010 has 360 non-overlapping class. While it may be true by class name, our analysis showed many super-class/subclass relations or renamed classes, e.g., ax and hatchet. To determine non-overlapping classes, we first processed 2010 classes with a 2012 classifier to determine which classes had more than 10% of its test images classified within the same 2012 class. We then had a human check if the classes were the same. The final list has 166 classes and is in our GitHub repo at <https://github.com/ROBOTICSENGINEER/Automatic-Open-World-Reliability-Assessment>.

Testing uses different percentages of data from the 2012 validation test data as knowns and from the 2010 classes as the unknowns or out-of-distribution data. To evaluate the policies, we created 24 configuration tests. Each configuration has 10,000 tests. Each test consists of 4000 images. The first 2000 images are mostly knowns, 20 of them unknowns (1%). The ratio of unknown images increased in the second 2000 images, ranging from 2% to 25% by 1% increments. We used a batch size of 100.

## 4 Background on EVM

Because we use the EVM as a core of multiple new algorithms, we briefly review its properties, so this paper is self-contained. The EVM is a distance-based kernel-free non-linear classifier that uses Weibull families distribution to compute the radial probability of inclusion of a point with respect to the class of another. Weibull families distribution is defined as

$$W(x; \mu, \sigma, \xi) = \begin{cases} e^{-(1+\xi(\frac{x-\mu}{\sigma}))^\xi} & , x < \mu - \frac{\sigma}{\xi} \\ 1 & , x \geq \mu - \frac{\sigma}{\xi} \end{cases} \quad (3)$$

where  $\mu \in \mathbb{R}$ ,  $\sigma \in \mathbb{R}^+$ , and  $\xi \in \mathbb{R}^-$  are locations, scale, and shape parameters [8, 1, 6]. EVM provides a compact probabilistic representation of each class's decision boundary, characterized in terms of its extreme vectors. Each extreme vector has a family of Weibull distribution.s Probability of a point to each class is defined as the maximum probability of point belonging to each extreme vector of the class.

$$\hat{P}(C_l|x) = \max_k W_{l,k}(x; \mu_{l,k}, \sigma_{l,k}, \xi_{l,k}) \quad (4)$$

where  $W_{l,k}(x)$  is Weibull probability of  $x$  corresponding to  $k$  extreme vector in class  $l$ . If we show unknown label

with 0, the predicted class label  $\hat{y} \in \mathbb{W}$  can be computed as following.

$$m(x) = \max_l \hat{P}(C_l|x) \quad (5)$$

$$z(x) = \arg \max_l \hat{P}(C_l|x) \quad (6)$$

$$q(x; \tau) = \text{Heaviside}(m(x) - \tau) \quad (7)$$

$$\hat{y}(x; \tau) = q(x; \tau) z(x) \quad (8)$$

where  $\tau \in \mathbb{R}^+$  is a threshold that can be optimized on some validation set. If the threshold is too large, EVM tends to put known images into the reject class, i.e., high false rejection rate. If the selected threshold is too small, EVM classifies unknown images to the known classes, i.e., high false classification rate. Therefore, this threshold acts as a slider in the trade-off between false rejection rate and false classification rate.

## 5 Algorithms for Open-world Reliability Assessment

We process batches, or windows of data, and attempt to detect the first batch where the distribution has shifted. Given the definition, the natural family of algorithms is to estimate the probability of input being unknown and then estimate properties of the distributions and compare them.

**Mean of SoftMax** Thresholding the SoftMax score of neural networks is well-known and the simplest method for out of distribution rejection of single instances [14]. Tracking the mean of recognition scores is a natural way to consider assessing if the input distribution has changed. To design a reliability assessor, we use this, which is the simple estimate of a distribution, its mean. Using the observation that if the distribution changes to include more OOD samples, one would expect the mean SoftMax score to also be reduced. In this method, we collect the maximum value of SoftMax for each image in the batch and save it in a buffer, where buffer size is equal to batch size. Then, we compute the mean of the buffer. Finally, reliability can be assessed by comparing the mean with a calibrated mean value, e.g., a mean calibrated to optimize performance on a validation set with some fixed mixing ratio. Algorithm 1 in supplemental material summarizes this method.

### Kullback-Leibler Divergence of Truncated Gaussian model of SoftMax Scores

The mean of the maximum SoftMax only considers the first moment of the distribution. Therefore, if the distribution of a batch of images changes while the mean of distribution remains constant or changes a little, the first method fails. A more advanced method for change detection is using Kullback–Leibler divergence with a more advanced distributional model of the data.

The Kullback-Leibler divergence is one of the most fundamental measures in information theory, which measures

the relative entropy

$$\text{KL}(P\|Q) = \int_{-\infty}^{+\infty} p(x) \log\left(\frac{p(x)}{q(x)}\right) dx \quad (9)$$

where  $x$  is maximum SoftMax,  $p(x)$  is probability density function of testing batch, and  $q(x)$  is probability density function of training data set.

Making the classic Gaussian assumption for the distribution, i.e. letting  $p(x) \sim \mathcal{N}(\mu, \sigma^2)$  and  $q(x) \sim \mathcal{N}(m, s^2)$ , we can derive the KL divergence measure as:

$$\text{KL}(P\|Q) = \log\left(\frac{s}{\sigma}\right) + \frac{\sigma^2 + (\mu - m)^2}{2s^2} - \frac{1}{2} \quad (10)$$

Algorithm 2 in the supplemental material summarizes this method, which is a direct application of KL divergence, as a measure between distribution of SoftMax values. Looking at the distribution in Fig. 1(a), one might note that the distribution is not at all symmetric and is bounded above by 1, which is the mode. Hence, it is not well modeled by a Gaussian with simple mean and standard deviation of the raw data. The important insight here is to consider this a truncated Gaussian and use a moment-matching approach for approximation, such as [17].

**KL divergence of Truncated Gaussian-model of EVM scores** It is, of course, natural to also consider combining the information theory-based KL divergence method with EVM-based probability estimates on each batch. Algorithm 4 in the supplement provides the details. Again, the distribution of scores shown in Fig. 1(b) is bounded, and this time with two disparate peaks, they seem even less Gaussian-like. This is a novel combination of KL divergence and EVM scores and is included in the evaluation to help assess the information gain from EVM probabilities over SoftMax with exactly the same algorithm.

**Fusion of KL divergence models** KL divergence of SoftMax and EVM are both viable methods with different error characteristics. Thus, it is also natural to ask if there is a way to fuse these algorithms. With KL divergence, an easy way to do this is to generalize to a bi-variate distributional model. We do this with a 2D Gaussian model, using a full 2x2 co-variance model, and Algorithm 5 summarizes this method. Plots in the supplemental material show its performance, which is not significantly different. Thus, given the added cost we don't recommend it.

**Open-world novelty detection with EVM** There are many ways to consider EVM probabilities, and given that the shape of data in Fig. 1 is not apparently Gaussian, we sought to develop a method with a different distributional assumption. We develop a novel algorithm, we call open-world novelty detection (OND), which is designed to use

only examples that have significant novelty, i.e., high probability of being OOD samples. Not only do we build from the extreme-value theory probabilities from EVM, but rather than looking at the mean or a Gaussian model of all data, we consider the mean of only those samples whose values probability of an input being from an unknown class exceed a threshold. We then combine that with a hypothesis test, which assumes a Bernoulli distribution of known and unknown inputs. This allows us to not just detect a distributional change but also if the change is improving or reducing reliability.

From (4), we expect

$$\begin{cases} 1 \geq m(x_{\text{known}}) \geq \tau \\ \tau > m(x_{\text{unknown}}) \geq 0 \end{cases} \quad (11)$$

By subtracting from 1,

$$\begin{cases} 1 - \tau \geq 1 - m(x_{\text{known}}) \geq 0 \\ 1 \geq 1 - m(x_{\text{unknown}}) > 1 - \tau \end{cases} \quad (12)$$

Let's define

$$\begin{cases} u(x) := 1 - m(x) \\ \delta := 1 - \tau \end{cases} \quad (13)$$

Then,

$$\begin{cases} \delta \geq u(x_{\text{known}}) \geq 0 \\ 1 \geq u(x_{\text{unknown}}) > \delta \end{cases} \quad (14)$$

By subtracting from  $\Delta \in \mathbb{R}^+$

$$\begin{cases} \delta - \Delta \geq u(x_{\text{known}}) - \Delta \geq -\Delta \\ 1 - \Delta \geq u(x_{\text{unknown}}) - \Delta > \delta - \Delta \end{cases} \quad (15)$$

If  $1 \geq \Delta > \delta > 0$

$$\begin{cases} \max\{0, u(x_{\text{known}}) - \Delta\} = 0 \\ 1 - \Delta \geq \max\{0, u(x_{\text{unknown}}) - \Delta\} > 0 \end{cases} \quad (16)$$

In batch mode with  $N$  images

$$\begin{cases} \sum \max\{0, u(x_{\text{known}}) - \Delta\} = 0 \\ N(1 - \Delta) \geq \sum \max\{0, u(x_{\text{unknown}}) - \Delta\} > 0 \end{cases} \quad (17)$$

For a given domain, one might assume the mixture of known and OOD classes can be modeled by a Bernoulli distribution with probability of  $\rho$ .

$$f(x) = \begin{cases} \rho & , x \in \text{OOD} \\ 1 - \rho & , x \in \text{Known classes} \end{cases} \quad (18)$$

By applying (18) in (17)

$$\rho N(1 - \Delta) \geq \sum \max\{0, u(x) - \Delta\} > 0 \quad (19)$$

over time  $\rho$  will change. If  $\rho$  is constant or decreases, image classifier remains reliable. However, if it increases, the image classifier should be updated.

We can find the error of this hypothesis as

$$\varepsilon := \max\{0, \left(\frac{1}{N} \sum \max\{0, u(x) - \Delta\}\right) - \rho(1 - \Delta)\} \quad (20)$$

Finally, we propose the OND policy for automatic reliability assessment via thresholding on  $\varepsilon$ , see Algorithm 3 in the supplemental material. Note this algorithm has added parameters  $\Delta$  and  $\rho$ , which need to be estimated on training/validation data. Although not evaluated herein, this algorithm has the operational advantage that it can actually be trivially tuned to target expected OOD ratios by changing  $\rho$ .

## 6 Experimental Results


For all experiments in this paper, we use an EfficientNet-B3 [30] network that was trained to classify the ImageNet 2012 training data, with 79.3% Top-1 accuracy on the 2012 validation set. Then, we extract features for images in the training set and train EVM with these extracted features. In the evaluation, we consider the overall detection rate, on-time detection, and mean absolute error (see Eq 2) as our primary metrics. We also show a metric called total accuracy. We start by considering the many batches in each test sequence and scoring each batch as 1.0 if it is correctly labeled as Reliable before the change or correctly labeled Unreliable after the change, and scoring 0.0 if incorrectly labeled. Total accuracy is the number of batches correctly labeled divided by the total number of batches.

In Fig. 2 we use window size of 100 and select each algorithms threshold to maximize its true detection rate (sum of on-time and late detection) when the percentage of unknown (out-of-distribution) data is at 2%. Operationally, users might determine the minimum variation they want to consider and the measure that is most important to them, then optimize that metric over the validation set to choose the parameter. For example, some users might prefer to maximize on-time detection or minimize false detection or mean absolute error.


Given these parameters, the performance of the various algorithms as the percentage of unknown data is increased, shown in Fig. 2. For this setting, All of the distributional algorithms, the KL divergent of EVM or SoftMax (Sec 5) as well as the OND algorithm 3, have similar good performance. The baseline mean of SoftMax-based Algorithm 1 has by far the lowest detection rate, total accuracy, and on-time.

In the remaining plots in the paper, we are doing ablation studies. Such ablation studies allow us to understand the deeper aspects of performance. In Fig. 3, we use batch/window size of 20 and 25 and select each algorithms threshold to limit false detections to at most 1% when the percentage of unknown (out-of-distribution) data is at 2%. With smaller windows the difference between KL EVM and KL SoftMax become significant, while KL EVM and OND are similar showing that the difference is the superior EVM features for detection of OOD data.


Next let us consider peak performance, which we shall see as a function of the percentage of OOD/Unknown data,



(a) Total Detection Percentage



(b) Percentage of on-time




(c) Mean Absolute Error

**Fig 2:** Performance of proposed policies when the threshold is selected to maximize the sum of on-time + late detection rates validation test with 2% unknown.

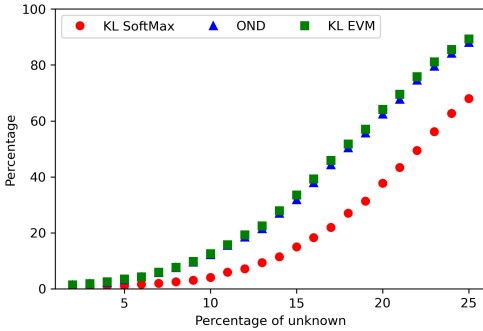
which is unknowable. This ablation study is to determine the impact and sensitivity of thresholds, Fig. 4 demonstrates the effect of threshold on performance Algorithm 5 in different percentages of unknown.

## 7 Discussion

Fig. 1 (on the first page) shows the distribution of the magnitude of the maximum probability of SoftMax and EVM on known validation and OOD validation sets. There are few known images with low predicted probability, and for SoftMax, many unknowns with high probability. Therefore, for any given threshold, a significant portion of the unknown images will be classified as known. Thus, with neither of



(a) True Detection, window size 20



(b) True Detection, window size 25

**Fig 3:** True detection performance with different sliding window sizes for the proposed policies when the threshold is selected to limit false detection to at most 1%. We show true detection rate with is sum of on-time and late detection.

these base models, is it possible to design a reliability assessor without failures. This also makes it clear why very simple naive algorithms, such as thresholding on SoftMax and retraining when images with low scores are detected, are useless – because the network is not perfect, and even training images will fail such a simple test, and the system will always be training. This is part of what makes this problem challenging and interesting – the algorithms must account for the imperfect nature of the core classification algorithm as well as handling the open-set nature of the problem while detecting the change in OOD.

The current EVM has three important parameters: cover threshold, tail size, and distance multiplier. We use 0.7 and 33998 for cover threshold and tail size, the same as the first EVM paper [26]. The EVM paper did not discuss a distance multiplier, seeming to use an implicit 0.5. We try 0.40, 0.45, 0.5, 0.75, and 1.0 for the distance multiplier. Among them, 0.45 demonstrates the best separation between known validation and unknown validation sets.

After training EVM, the threshold  $\Delta$  must be selected for KL EVM Algorithm 4. Because the cover threshold is 0.7, the  $\Delta$  must be greater than 0.3 and less than 0.7. This data can be optimized based on the validation data set. However, with such a small region of allowed values, we

decide to choose the middle, i.e., 0.5, to avoid bias. There is not any consensus or common sense to determine thresholds exactly, especially given the inherent trade-off between false detection and failure. As the percentage of failure decreases, the false detection increases and vice versa. Studying all possible methods to tune the thresholds is beyond the scope of this paper. One approach is to select a threshold slightly above the signal to avoid early detection. This approach in practice can be found by creating a validation set and updating it in a timely manner according to the need and scene of the images in a real test.

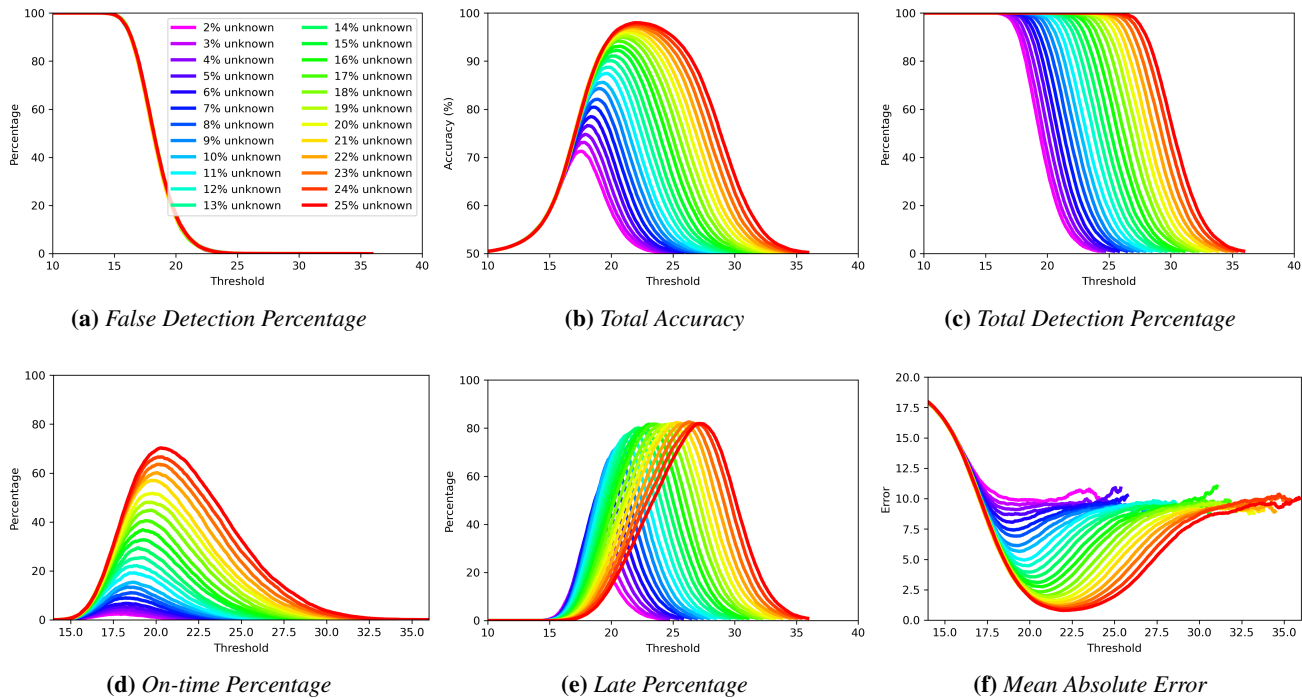
We note that our KL divergence models are assuming a truncated Gaussian distribution, which is only an approximation to the shapes of the distributions shown in Fig. 1, and the quality of the approximation will be impacted by mixing and by window sizes. While only an approximation, the KL divergence does much better than the simple detection of a shift in the mean. The OND algorithm makes very different assumptions using a Bernoulli model. The nearly equal performance of KL divergence assuming truncated Gaussian and OND assuming Bernoulli suggest the overall detection is relatively stable and not strongly dependent on either set of assumptions.

A good algorithm must detect the change to unreliability even with a small input distribution shift of only a few percent and it does. However, our experiments show that even with a large distribution shift, i.e., 25% OOD data, algorithms fail to always detect the shift on-time. The main reason is that features of unknown images can overlap with known images, and, therefore algorithms cannot sense the shift. In a few splits, Kullback–Leibler divergence of SoftMax (baseline) works somewhat better than algorithms that use EVM information. However, when the shift is larger, the algorithms that use EVM information outperform the algorithms that use only SoftMax. For lower percentages of unknown KL EVM does slightly better in total detection, and for larger percentages OND has slightly lower mean absolute error of detection, but neither are very significant differences.

The sensitivity and accuracy of all presented methods increases as batch size increases. However, the batch size is proportional to the inverse of speed. If the batch size goes to infinity, the speed of detection goes to zero, which is not practical. Thus, in the trade-off between accuracy and speed, we should find a conformation to reduce the batch size while the distribution approximately is a Gaussian. For the ImageNet data, we found that a batch size of 100 is a good balance for both and EVM is reasonably stable down to windows of 20, but KL on SoftMax requires larger batches.

We note that all experiments herein were with ImageNet classifiers which have 1000 classes. For smaller problems, there may be a greater issue of OOD classes folding on-





**Fig 4:** Ablation study on performance of Bivariate KL Fusion Algorithm Section 5 when the threshold changes. As the percentage of unknown increases, the breadth of good algorithm performance increases, and the peak performance of moves slightly to the right (to higher thresholds).

top of known classes(see [11]), and hence the difference between SoftMax and EVM based detection may be more apparent for problems with smaller number of classes.

## 8 Conclusions

Both standard closed-set and more recent open-set classifiers face an open-world, and errors can increase when there is a change in the distribution of the frequency of unknown inputs. While open-set image classifiers can map images to a set of predefined classes and simultaneously reject images that do not belong to the known classes, they cannot predict their own overall reliability when there is a shift in the OOD distribution of the world. Currently, one must allocate human resources to monitor the quality of closed or open-set image classifiers. A proper automatic reliability assessment policy can reduce the cost by decreasing human operators and simultaneously increasing user satisfaction.

In this paper, we formalized the open-world reliability assessment problem and proposed multiple automatic reliability assessment policies (Algorithms 2, 3, 4, and 5) that use only score distributional data. Algorithm 3 uses the maximum of the mean of thresholded maximum EVM class probability to determine reliability. This algorithm ignores the higher order moments of distribution for simplicity. Algorithm 4 uses Gaussian Model of score distributions and Kullback–Leibler divergence of maximum EVM class probabilities. We also evaluate, Algorithm 2, that uses the

Kullback–Leibler divergence on a Gaussian Model of SoftMax scores, and show it also provides an effective open-world reliability assessment.

We used the mean of SoftMax as the baseline, and the new algorithms are significantly better. We used ImageNet 2012 for known images and non-overlapping ImageNet 2010 for unknown images. As in any detection problem, there is an inherent trade-off between false or early-detection and true detection rates. If we tune the threshold, so there are zero-false detection, then even the best algorithm fails to detect in 5.68% of tests. This is because features of unknown images can overlap on known images, and therefore, algorithms cannot sense the shift. In future works, we will investigate extreme value theory rather than Gaussian assumptions to design a better automatic reliability assessment policy.

## References

- [1] Jan Beirlant, Yuri Goegebeur, Johan Segers, and Jozef L. Teugels. *Statistics of extremes: theory and applications*. John Wiley & Sons, 2006.
- [2] Abhijit Bendale and Terrance E Boulton. Towards open set deep networks. In *IEEE CVPR*, pages 1563–1572, 2016.
- [3] Arjun Nitin Bhagoji, Daniel Cullina, Chawin Sitawarin, and Prateek Mittal. Enhancing robustness of machine learning systems via data transformations. In *2018 52nd Annual Conference on Information Sciences and Systems (CISS)*, pages 1–5. IEEE, 2018.



- [4] Supritam Bhattacharjee, Devraj Mandal, and Soma Biswas. Multi-class novelty detection using mix-up technique. In *The IEEE Winter Conference on Applications of Computer Vision*, pages 1400–1409, 2020.
- [5] Terrance E Boulton, Steve Cruz, Akshay Raj Dhamija, M Gunther, James Henrydoss, and Walter J Scheirer. Learning and the unknown: Surveying steps toward open world recognition. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 33, pages 9801–9807, 2019.
- [6] Enrique Castillo. *Extreme value theory in engineering*. Elsevier, 2012.
- [7] Hakan Cevikalp and Halil Saglam. Polyhedral conic classifiers for computer vision applications and open set recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2019.
- [8] Stuart Coles, Joanna Bawa, Lesley Trenner, and Pat Dorazio. *An introduction to statistical modeling of extreme values*, volume 208. Springer, 2001.
- [9] Antitza Dantcheva, Arun Singh, Petros Elia, and Jean-Luc Dugelay. Search pruning in video surveillance systems: Efficiency-reliability tradeoff. In *2011 IEEE International Conference on Computer Vision Workshops (ICCV Workshops)*, pages 1356–1363. IEEE, 2011.
- [10] Tamraparni Dasu, Shankar Krishnan, Dongyu Lin, Suresh Venkatasubramanian, and Kevin Yi. Change (detection) you can believe in: Finding distributional shifts in data streams. In *International Symposium on Intelligent Data Analysis*, pages 21–34. Springer, 2009.
- [11] Akshay Raj Dhamija, Manuel Günther, and Terrance Boulton. Reducing network agnostophobia. In *Advances in Neural Information Processing Systems*, pages 9157–9168, 2018.
- [12] GWA Dummer, R Winton, and Mike Tooley. *An elementary guide to reliability*. Elsevier, 1997.
- [13] Chuanxing Geng and Songcan Chen. Collective decision for open set recognition. *IEEE Transactions on Knowledge and Data Engineering*, 2020.
- [14] Chuanxing Geng, Sheng-jun Huang, and Songcan Chen. Recent advances in open set recognition: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2020.
- [15] Lykele Hazelhoff, Ivo Creusen, and Peter HN de With. Robust classification system with reliability prediction for semi-automatic traffic-sign inventory systems. In *2013 IEEE Workshop on Applications of Computer Vision (WACV)*, pages 125–132. IEEE, 2013.
- [16] James Henrydoss, Steve Cruz, Ethan M Rudd, Manuel Gunther, and Terrance E Boulton. Incremental open set intrusion recognition using extreme value machine. In *2017 16th IEEE International Conference on Machine Learning and Applications (ICMLA)*, pages 1089–1093. IEEE, 2017.
- [17] Malte Kuss and Carl Edward Rasmussen. Assessing approximate inference for binary gaussian process classification. *Journal of machine learning research*, 6(Oct):1679–1704, 2005.
- [18] Salar Latifi, Babak Zamirai, and Scott Mahlke. Polygraphmr: Enhancing the reliability and dependability of cnns. In *2020 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 99–112. IEEE, 2020.
- [19] Qingming Leng, Mang Ye, and Qi Tian. A survey of open-world person re-identification. *IEEE Transactions on Circuits and Systems for Video Technology*, 30(4):1092–1108, 2019.
- [20] Yi Liu and Yuan F Zheng. One-against-all multi-class svm classification using reliability measures. In *Proceedings. 2005 IEEE International Joint Conference on Neural Networks, 2005.*, volume 2, pages 849–854. IEEE, 2005.
- [21] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2537–2546, 2019.
- [22] Ling Ma, Xiabi Liu, Li Song, Chunwu Zhou, Xinming Zhao, and Yanfeng Zhao. A new classifier fusion method based on historical and on-line classification reliability for recognizing common ct imaging signs of lung diseases. *Computerized Medical Imaging and Graphics*, 40:39–48, 2015.
- [23] Sergio Matiz and Kenneth E Barner. Inductive conformal predictor for convolutional neural networks: Applications to active learning for image classification. *Pattern Recognition*, 90:172–182, 2019.
- [24] Poojan Oza and Vishal M Patel. C2ae: Class conditioned auto-encoder for open-set recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2307–2316, 2019.
- [25] Santiago Romaní, Pilar Sobrevilla, and Eduard Montseny. On the reliability degree of hue and saturation values of a pixel for color image classification. In *The 14th IEEE International Conference on Fuzzy Systems, 2005. FUZZ'05.*, pages 306–311. IEEE, 2005.
- [26] Ethan M Rudd, Lalit P Jain, Walter J Scheirer, and Terrance E Boulton. The extreme value machine. *IEEE transactions on pattern analysis and machine intelligence*, 40(3):762–768, 2017.
- [27] Vikash Sehwal, Arjun Nitin Bhagoji, Liwei Song, Chawin Sitawarin, Daniel Cullina, Mung Chiang, and Prateek Mittal. Analyzing the robustness of open-world machine learning. In *Proceedings of the 12th ACM Workshop on Artificial Intelligence and Security*, pages 105–116, 2019.
- [28] Dong Su, Huan Zhang, Hongge Chen, Jinfeng Yi, Pin-Yu Chen, and Yupeng Gao. Is robustness the cost of accuracy?—a comprehensive study on the robustness of 18 deep image classification models. In *Proceedings of the European Conference on Computer Vision (ECCV)*, pages 631–648, 2018.
- [29] Zhenan Sun, Hui Zhang, Tieniu Tan, and Jianyu Wang. Iris image classification based on hierarchical visual codebook. *IEEE Transactions on pattern analysis and machine intelligence*, 36(6):1120–1133, 2013.
- [30] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114, 2019.
- [31] Yisen Wang, Weiyang Liu, Xingjun Ma, James Bailey, Hongyuan Zha, Le Song, and Shu-Tao Xia. Iterative learning with open-set noisy labels. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 8688–8696, 2018.

- [32] Ryota Yoshihashi, Wen Shao, Rei Kawakami, Shaodi You, Makoto Iida, and Takeshi Naemura. Classification-reconstruction learning for open-set recognition. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 4016–4025, 2019.
- [33] Bailing Zhang. Reliable classification of vehicle types based on cascade classifier ensembles. *IEEE Transactions on intelligent transportation systems*, 14(1):322–332, 2012.
- [34] Stephan Zheng, Yang Song, Thomas Leung, and Ian Goodfellow. Improving the robustness of deep neural networks via stability training. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4480–4488, 2016.
- [35] Weibao Zou, Yan Li, and Arthur Tang. Effects of the number of hidden nodes used in a structured-based neural network on the reliability of image classification. *Neural Computing and Applications*, 18(3):249–260, 2009.
- [36] Weiwen Zou and Pong C Yuen. Discriminability and reliability indexes: two new measures to enhance multi-image face recognition. *Pattern Recognition*, 43(10):3483–3493, 2010.