

Generative OpenMax for Multi-Class Open Set Classification

Zongyuan Ge*¹
zongyuan@au1.ibm.com

Sergey Demyanov*¹
sergey@demyanov.net

Zetao Chen²
chenze@ethz.ch

Rahil Garnavi¹
rahilgar@au1.ibm.com

¹ IBM Research
204 Lygon Street
Melbourne, Australia

² Vision for Robotics Lab
ETH Zurich
Zurich, Switzerland

Abstract

We present a conceptually new and flexible method for multi-class open set classification. Unlike previous methods where unknown classes are inferred with respect to the feature or decision distance to the known classes, our approach is able to provide explicit modelling and decision score for unknown classes. The proposed method, called Generative OpenMax (G-OpenMax), extends OpenMax by employing generative adversarial networks (GANs) for novel category image synthesis. We validate the proposed method on two datasets of handwritten digits and characters, resulting in superior results over previous deep learning based method OpenMax. Moreover, G-OpenMax provides a way to visualize samples representing the unknown classes from open space. Our simple and effective approach could serve as a new direction to tackle the challenging multi-class open set classification problem.

1 Introduction

The computer vision community has progressively improved object classification performance in recent years. Driven by recent research in machine learning, the best performing system can achieve impressive results on both small and large scale multi-class classification tasks [8, 9, 13, 21, 28]. Most commonly, these recognition tasks are defined under a closed-set setting, i.e. all testing samples belong to one of the known classes in the training set. However, that is not the ideal setting to train a machine targeting real world applications. The most comprehensive image dataset these days, ImageNet [9] only contains 1000 classes. Considering only the number of animal species on the earth, the number of classes would exceed 1.5 million [16] animal species on the earth, not to mention other general object categories.

To tackle this challenge, open set classification topic has recently attracted much attention [2, 3, 11, 24, 25]. Open set problem assumes that for those test objects who do not belong to any known classes (i.e. classes the model has been seen before and trained on), the

classifier must correctly identify them as unknown class, as opposed to false classification to one of the known classes. Multi-class open set classification is challenging because it requires the correct probability estimation of all known classes, together with simultaneous precise predicting of unknown classes. Given this, [24] defines open set classification as a problem of balancing known space (specialization) and unknown open space (generalization) of the model (see Fig. 1). Formalization for open space risk is considered as the relative measure of open space compared to the overall measure space:

$$R_{open} = \frac{Unknown\ Space}{(Known\ Space + Unknown\ Open\ Space)} \quad (1)$$

Several recent solutions to multi-class open set problem convert it into a score calibration task [10, 13, 17]. Those methods model the known class distribution using a parametric model from which the posterior probabilities of a test sample are computed. The unknown class probabilities are then estimated through statistical modelling, specially Extreme Value Theory (EVT) [16]. However, those models do not assume any prior information about open space. In essence, the “open space risk” is only estimated in decision space rather in pixel space.

Our proposed method, called Generative OpenMax (G-OpenMax), extends OpenMax [9] by providing explicit probability estimation over unknown categories. This is done by using generative adversarial networks (GANs) [8] to generate synthetic samples from unknown classes. The synthetic samples are generated from mixture distributions of known classes in latent space, which leads to plausible representation with respect to the known classes domain. Explicit representation of unknown classes enables the classifier to locate the decision margin with the knowledge of both known and unknown samples, thus resulting in better balance of open space risk and known space for the multi-class open set model. The contributions of this paper can be summarized as:

1. We propose a new method G-OpenMax to tackle the challenging multi-class open set classification. Our proposed method enables non-pseudo probability estimation and visualization over both known and unknown classes.
2. We modify the training procedure of conditional GAN and mix several prior known classes distributions to generate plausible and domain adapted synthetic unknown samples.
3. Comprehensive experimental analysis have been conducted to evaluate the effectiveness and limitations of the proposed method on small scale (10 classes) and large scale (95 classes) openness problems.

2 Background

Openset set: Early attempts to solve open set classification task involved adapted closed-set classifiers. Scheirer et al. [24] proposed “1-vs-Set Machine” which detects an unknown class by exploiting a decision space from the marginal distances of a binary SVM classifier. Later Scheirer et al. [25] proposed compact abating probability (CAP) model which extended “1-vs-Set Machine” to non-linear W-SVM for multi-class open set scenario. Meanwhile, there were a few works exploring non-close-set methods for open set tasks. Jain et al. [10] introduced P₁-SVM to leverage Statistical Extreme Value Theory (EVT) and multi-class SVM

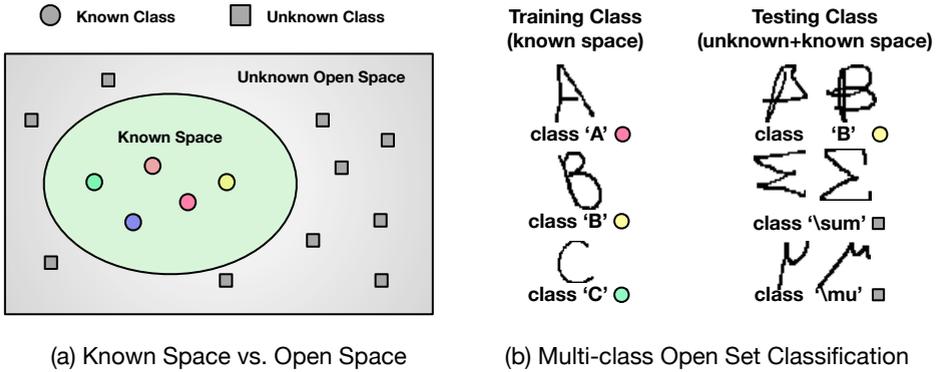


Figure 1: Figure (a) illustrates the relationship between the known space and unknown open space. Green area defines known space containing various known classes (coloured circles). Gray area represents open space and consists of unknown classes (grey rectangular). Figure (b) illustrates the open set setting by showing classes from train/test phase, where classes such as “ μ ” and “ Σ ” do not belong to the training set, hence unknown to the classifier during training (images are from HASYv2 dataset [24]).

scores to robustly estimate the posterior probability of known and unknown classes. Bendale et al. [4] offered the definition of “open world recognition” and proposed Nearest Non-Outlier (NNO) algorithm to actively detect and learn new classes to the model. Recently, driven by the effectiveness of deep networks, Bendale and Boulton [5] adapted deep learning framework for open set recognition. Combining the concept of penultimate layer with meta recognition [23], they proposed OpenMax which enables basic deep network to reject either “fooling” or unknown open set classes.

Generative Adversarial Learning: There are various approaches for generative models. The most prominent are variational autoencoder (VAE) [12] and generative adversarial networks (GAN) [1], among which, GAN is a powerful model of learning arbitrarily complex data distributions. GAN is formulated on a game-theoretic using the philosophy of competition between two networks for training an image synthesis model. Recent work shows GAN is able to produce highly plausible images on both gray-scale and RGB images [20]. Many improvements have been proposed to enhance the overall performance for training GANs [6, 22]. GAN has a wide range of applications including, sketch/text to image synthesis [4], image to image translation [10], image editing [19], super resolution [15] and video prediction [8], etc. To the best of our knowledge, this is the **first** time generative model like GAN has been used for open set classification task.

3 Methodology

Generative OpenMax (G-OpenMax) is conceptually simple. While OpenMax estimates the pseudo probability of unknown class by aggregating calibrated scores from known classes, our proposed G-OpenMax, which is an intuitive solution, directly estimates the probability of unknown class. This is done by using synthetic images as an extra training label apart from known labels. Therefore, the main challenge is the following: how to generate plausible

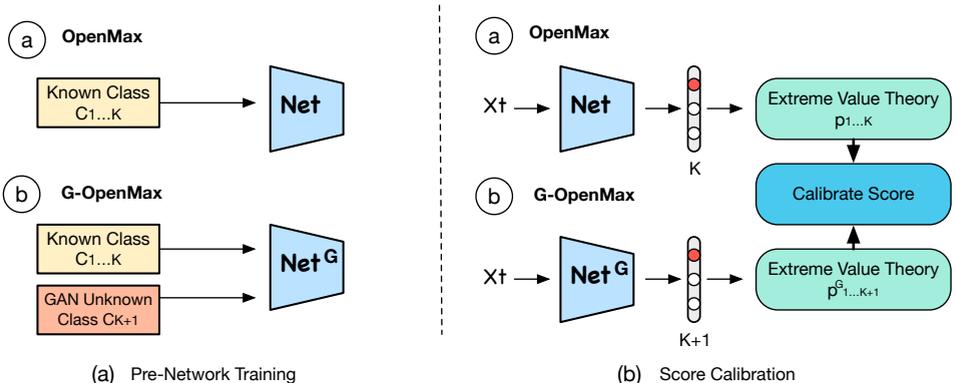


Figure 2: Figure (a) illustrates the pre-training process of Net and Net^G . GAN-based synthetic images are used as an extra training label c_{K+1} apart from known labels $[c_{1,...,K}]$. Figure (b) explains the difference between score calibration in OpenMax and G-OpenMax.

samples such that they meet the criteria of being distinct from the known classes, while they are well represent for the open space. In Sec. 3.2 and Sec. 3.3, we introduce the key elements of G-OpenMax, including how to synthesis unknown samples and post-processing to select suitable cases for open set classifier training.

3.1 OpenMax

We begin by reviewing the OpenMax open set classifier [3]. It explores the effectiveness of penultimate layer from the deep network to limit open space risk. Like multi-class open set classifier W-SVM [24], OpenMax holds the property of compact abating probability (CAP) model. The learning progress of OpenMax can be summarized into two stages. The first stage leverages EVT as a base theory and fit the known classes' post recognition activations to a Weibull distribution. To do so, a base network Net (AlexNet [13] is used in the original paper) is trained as a penultimate activation layer extractor. Mean activation vector (MAV) $MAV = [\mu_{i,...,K}]$ for each class $C = [c_{i,...,K}]$ is computed based on the penultimate activation layer of each sample. Weibull model with hyper-parameters $p_{c_i} = (t_{c_i}, \lambda_{c_i}, k_{c_i})$ for each class is returned. The second stage, which is essential in OpenMax, recalibrates each activation of penultimate layer by estimating Weibull CDF probability on the distance between sample x_i and known class' MAV $[\mu_{i,...,K}]$, which servers as the core of the rejection estimation. At last, pseudo-probability of unknown class is estimated from known class' activation scores to support explicit rejection.

3.2 Generative OpenMax (G-OpenMax)

Here we describe the overall inference procedure of the proposed method G-OpenMax in Alg. 1. The main difference between vanilla OpenMax and G-OpenMax is shown in Fig. 2. G-OpenMax adopts the same two-stage procedure as OpenMax. However, in parallel to training the classifier and Weibull model at the first stage by using known classes $c_{i,...,K}$ only, we train the network Net^G with $K + 1$ classes (see Fig. 2 (a)), where the extra images come from our generator G to represent the unknown class c_{K+1} . Consequently, the activation

Algorithm 1: Generative OpenMax algorithm

input : Number of top-score predictions to calibrate, α
Pre-set threshold ε

output: known class probability $Prob_{1,\dots,K}$ + unknown class probability $Prob_{K+1}$

- 1 **Train** Net from known classes $c_{i,\dots,K}$;
- 2 **Train** the generative model G ;
- 3 **Unknown sample selection** using Net ;
- 4 **Train** Net^G with known class and unknown class samples;
- 5 **Computing** activation vector $v^G(x) = v_1(x), \dots, v_{K+1}(x)$ from Net^G ;
- 6 **Fit** Weibull models $p^G = [p_{1,\dots,K+1}^G]$;
- 7 $argsort(v^G(x))$;
- 8 **for** $c=1,\dots,\alpha$ **do**
 - 9 $w_{s(c)}(x) = 1 - \frac{\alpha-c}{\alpha} e^{-\left(\frac{\|v^G(x)-t_{s(c)}\|}{\lambda_{s(c)}}\right)^{k_{s(c)}}}$;
- 10 **end**
- 11 Recalibrate activation vector: $v(x)' = v^G(x) \odot w(x)$;
- 12 $Prob = Softmax(v(x)')$;
- 13 **if** $\max_c Prob(y = c|x) < \varepsilon$ **then**
 - 14 $pred = unknownclass$;
 - 15 **else**
 - 16 $pred = argmax_c P(y = c|x)$;
 - 17 **end**

vector v^G in G-OpenMax has one extra dimension output to represent the $K + 1$ class which is regarded as unknown. It is noticed that the penultimate layer provides dependent information about how classes are related with each other. Therefore, the extra dimension is able to supply information about MAV distance between known and unknown classes in feature space ($\mu_{i,\dots,K}$ with respect to μ_{K+1}). At the second stage, G-OpenMax follows the same procedure as OpenMax using Weibull CDF probability for score calibration. However, G-OpenMax provides explicit probability estimation of the unknown class (line 6 in Alg. 1).

3.3 Understanding Generation of Synthetic Samples

While in theory the open space should be infinitely large, an ideal open set classifier is able to recognize unknown classes come from any distributions. In practice, many recent open set works [2, 3, 10, 24, 25] assume that the classes from unknown open space will share some common properties with known classes. To elaborate, if the system is trained on a set of English or Chinese characters, an open set class will probably be another character from the same group. An RGB based colour image is not expected to evaluate the characters trained system. This is a fair assumption to make a first step solution toward the challenging open set problem. Thus, in this paper we assume that *open space classes belong to a subspace of the original space*, which includes known classes. In order to preserve it, we also do not consider an unlikely scenario when the test set contains objects from other datasets. The advantage of building algorithm based on this assumption is that we would turn an open space classification problem into a standard closed set problem by sampling objects from known

subspace. A natural way to do this is linear interpolation (or linear combination) between objects, belonging to the subspace. However, in the original pixel space this subspace is highly nonlinear, so linear interpolation would not produce any plausible results. For instance, interpolating between photos of two birds will not result in a realistic photo of another bird class. However in the latent space, class information is separated from other object characteristics, so the class can be modified independently without affecting them [19]. Since all dataset classes are encoded by one of the basis one-hot vectors, we can generate new classes based on the linear combination of those known classes in latent space. Moreover, the generated mixed-class samples share a number of common properties with known class images, all of those produced images will have high probability in a given data distribution. Which makes those generated samples good candidates to represent a reasonable open space under our assumption. Formally, we can assume that existing classes are encoded by one-hot basis vectors $\mathbf{b} \in \mathbb{R}^N$, which define the class subspace in the latent space. In this case, we can generate a class mixture vector $\mathbf{m} \in \mathbb{R}^N$ using any distribution P_{norm} (e.g. Normal distribution) by sampling m_1, \dots, m_{N-1} from P_{norm} , and assigning $m_N = 1 - \sum_{i=1}^{N-1} m_i$, such that $\sum_{i=1}^N m_i = 1$ as for the basis vectors $\mathbf{b}_i \in \mathbb{R}^N$. The dimensionality of class subspace space (number of classes) is typically much smaller than the dimensionality of the original space (number of pixels), so the generated open set images will be located in a small nonlinear subspace of it as desired.

We employ a modified conditional GAN [18] to train the generator and then synthesize unknown classes out of it. Note that the algorithm is not GAN-specific, it also generalizes to any generative models, e.g. VAE [20]. A typical GAN is composed of a generator G and a discriminator D . While G learns to transform random noise to samples similar to those from the dataset, D learns to differentiate them. In a case of conditional GAN, random noise is fed to G together with a one-hot vector $c \in c_{1,\dots,K}$, which represents a desired class. In this case, discriminator D also learns faster if the input image is supplied together with the class it belongs to. Thus, the optimization of a conditional GAN with class labels can be formulated as:

$$\min_{\phi} \max_{\theta} = E_{x,c \sim P_{data}} [\log D_{\theta}(x, c)] + E_{z \sim P_z, c \sim P_c} [\log(1 - D_{\theta}(G_{\phi}(z, c), c))] \quad (2)$$

where the generator inputs z and c are the latent variables drawn from their prior distribution $P(z)$ and $P(c)$. Here ϕ and θ denote trainable parameters for G_{ϕ} and D_{θ} , respectively.

Selection of Generated Samples: One important step is that we need to select open space samples which do not belong to the subspace of known classes. For each synthetic sample generated from a mixture of class distribution, the class with the highest value is treated as the ground-truth label. We feed all the generated samples $x_{1,\dots,N}^G$ into the pre-trained classifier Net , all incorrectly predicted samples are selected as candidates for c_{k+1} to train the NET^G .

3.4 Implementation Details

We set hyper-parameters following existing DCGAN work [20]; the GAN is optimized with Adam optimizer ($\beta_1 = 0.5, \beta_2 = 0.999, \epsilon = 10^{-8}$) using a learning rate of $2e-4$. The architectures of both classifier and generator/discriminator of the conditional GAN are provided in Table 1 to 3. The generator G and discriminator D takes data x as initial input, and at each linear layer thereafter, the latent representation c ($N = 50$) is transformed using a learned linear transformation to the hidden layer dimension and added to the non-linearity input. Moreover, to improve the model convergence time, we use soft-labels with 0.9 for training [22].

Table 1: Classifier

Layer-type	Kernel	Stride/Pad	Filters	Activation	Norm
Convolution	(3,3)	1/1	32	ReLu	-
Max Pooling	(2,2)	2/0	-	-	-
Convolution	(3,3)	1/1	64	ReLu	-
Max Pooling	(2,2)	2/0	-	-	-
Fully-Connected	-	-	1024	Relu	-
Dropout	0.5 ratio	-	-	-	-
Softmax Classifier	-	-	K/K+1	-	-

Table 2: Generator

Layer-type	Kernel	Stride/Pad	Filters	Activation	Norm
Concatenation of z and c/b					
Fully-Connected	-	-	1024	Relu	-
Fully-Connected	-	-	512	Relu	-
Fully-Convolution	(4,4)	2/2	128	Relu	BN
Fully-Convolution	Image Size	2/2	1	Relu	BN
Output-Layer	-	-	-	-	Tanh

Table 3: Discriminator

Layer-type	Kernel	Stride/Pad	Filters	Activation	Norm
Convolution	(3,3)	1/1	64	ReLu	-
Max Pooling	(2,2)	2/0	-	-	-
Convolution	(3,3)	1/1	64	ReLu	BN
Max Pooling	(2,2)	2/0	-	-	-
Fully-Connected	-	-	256	Relu	-
Fully-Connected	-	-	1	Sigmoid	-

4 Experiments

We perform a thorough comparison of G-OpenMax with OpenMax and its variants on two datasets: MNIST [24] and HASYv2 [29]. MNIST is a digits dataset of 10 classes from 0 to 9 composed of 70,000 images. HASYv2 is the latest dataset of handwritten symbols published in early 2017. It is a more complicated dataset than MNIST in terms of the number of classes (369 classes) and domain of characters (including Arabic and Latin characters). However, HASYv2 is not a balanced dataset with respect to the number of samples per class. In this paper we use only the classes from HASYv2 with no less than 500 samples to ensure enough samples to train a reasonable generator G . Standard data sets for open set cross-validation is performed for analysis. We use 60% samples of each class for training¹, and 20% samples are held out for validation. The last 20% are used for testing. All results are reported as average value from three-fold cross-validation.

4.1 Main Results

Following [25], we plot and compare “openness” vs F-measure to a few baseline methods in multi-class open set classification. The “openness” value defines the number of known classes in training and the total number of classes to be recognized in testing.

$$openness = 1 - \sqrt{\frac{2N_{Train}}{N_R + N_{Test}}} \quad (3)$$

where N_{Train} defines the number of training classes, N_{Test} denotes the number of testing classes and N_R presents the number of classes to be recognized. The setup on MNIST is a

¹Given that training a generator G with 60 classes from HASYv2 is difficult to converge, we randomly choose 20 classes from the training pool to train one generator in each iteration. Later on all generated images from different iterations are used together as synthetic samples.

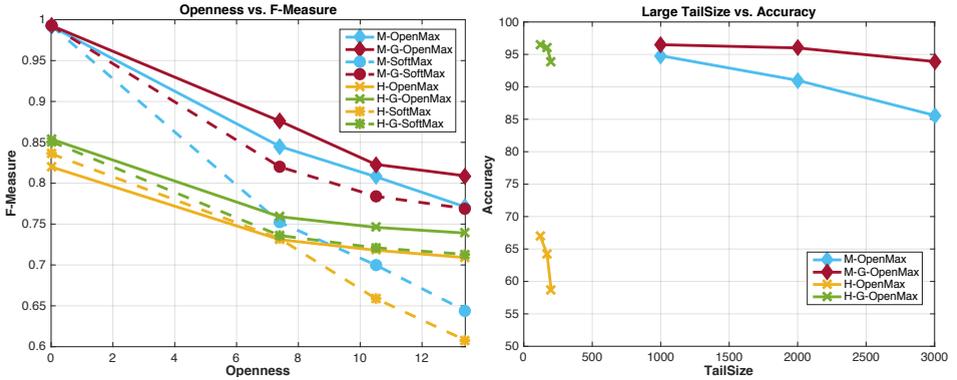


Figure 3: **Left:** F-measure for multi-class open set recognition on MNIST (**M**) and HASYv2(**H**). G-OpenMax maintains high F-measure scores as the openness grows, while vanilla Softmax degrades quickly on MNIST. G-SoftMax is better compared to its baseline SoftMax on both datasets and achieves approximately the same performance as OpenMax on MNIST with openness equal to 13.04. **Right:** This graph shows known classes recognition accuracy when using large tail sizes for Weibull model fitting. We observe that G-OpenMax consistently outperforms OpenMax, which leads to the conclusion that G-OpenMax is robust on known samples detection even with large tail size.

replication of the experiment performed in [10], where 6 classes are hold out in training and the rest 4 classes are used to varying “openness”. For HASYv2 dataset, we use 60 classes for training and the rest 35 classes are held out for testing as unknown classes. The results of our method (G-OpenMax) are compared with: 1) **SoftMax:** Softmax with probability thresholding as unknown class detector. 2) **OpenMax:** OpenMax described in Sec. 3.1. 3) **G-SoftMax:** Like Softmax, but G-SoftMax employs synthetic samples as unknown class training samples.

Fig. 3 left shows open set multi-class recognition F-measure performance for MNIST and HASYv2. F-measure for SoftMax degrades rapidly on both test scenarios. On MNIST we see that SoftMax trained with GAN based synthetic samples (G-SoftMax), without score calibration, is competitive against OpenMax. Furthermore, G-Softmax is much better at tolerating increasing openness than normal SoftMax. On HASYv2, which has larger number of classes than MNIST, the performance gain of employing GAN samples is expected to be higher than MNIST due to rich variety of generated samples. However, the relatively poor quality of generated mixture samples (see Sec. 4.3) in HASYv2 hinders the training of G-OpenMax and G-Softmax. Overall, the proposed G-OpenMax based approach provides consistent improvement regardless of datasets.

4.2 Ablation Experiments

We run a number of ablation experiments to analyze robustness of different models. The openness evaluated from all experiments in this subsection is 13.04. To reduce the complexity of running relevant experiments, we use fixed α in those experiments ².

Tail Sizes for Score Calibration: Fig. 4 left presents the results of varying tail size and

²We use $\alpha = 2$ for MNIST and $\alpha = 4$ for HASYv2.

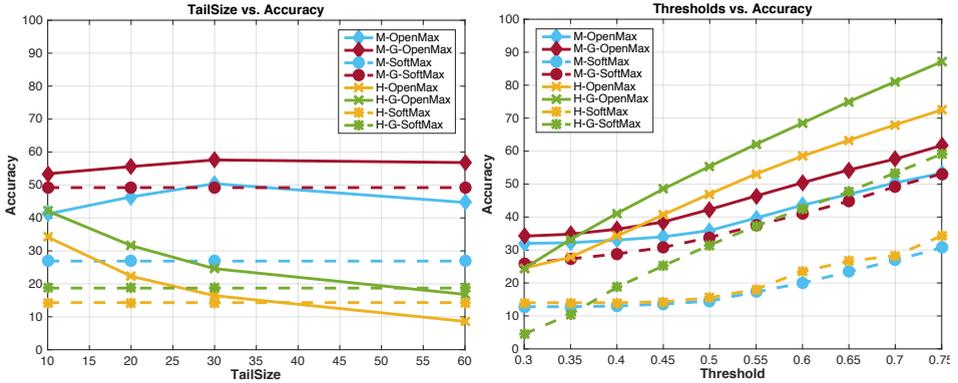


Figure 4: **Left:** This graph shows accuracy on unknown class performance with respect to tail size for open set multi-class recognition. **Right:** Given different threshold, we report accuracy on unknown class test samples from different methods on MNIST and HASYv2 dataset.

evaluating the unknown class accuracy with its optimal threshold. We observe that performance on MNIST continues to improve upto tail size of 60. However, large tail size does not work well on HASYv2 dataset. This may attribute to the difficulty of generating meaningful samples from HASYv2 dataset since the complexity and large variations of known classes.

Large Tail Size vs. Performance: In this part, we tried training EVT model (OpenMax and G-OpenMax) with extremely large tail size. Although large tail size tends to increase the EVT model rejections for unknown samples, it also increases rejection for known classes [4]. According to our observations from Fig. 3 left, G-OpenMax is able to maintain the same overall performance even with large tail size.

Threshold vs. Performance: We evaluated performance of different models on unknown test samples with optimal tail size³. Our experiments results (see Fig. 3 right) highlights that the methods trained with GAN samples consistently outperforms methods without employing GAN samples. G-OpenMax is nearly (+10%) absolute performance improvement over OpenMax with optimal threshold.

4.3 Qualitative Analysis

It is helpful to visualize the mixture samples generated by GAN to get an insight about the classes used as training samples for G-SoftMax and G-OpenMax. Fig.5 (a) and (b) demonstrate mixture examples generated from MNIST, HASYv2 and ImageNet12 respectively. Mixture samples from MNIST show some distinct attributes of digits, such as stroke continuity and orientation etc. Mixture samples from HASYv2 are sometimes hard for human to understand (second row last sample). The main reason may attributes to the complexity of the HASYv2 dataset since it consists of English alphabets, Greek letters etc.

A natural extension of our method is to perform experiments on the natural image dataset. We trained the generator G on ImageNet12 and made several observations: 1) Several classes blend into unrealistic type of object in terms of colour and global structure. 2) Mixture model tend to collapse into a single class appearance output. However, we performed the

³tail sizes are selected according to Fig. 3 left

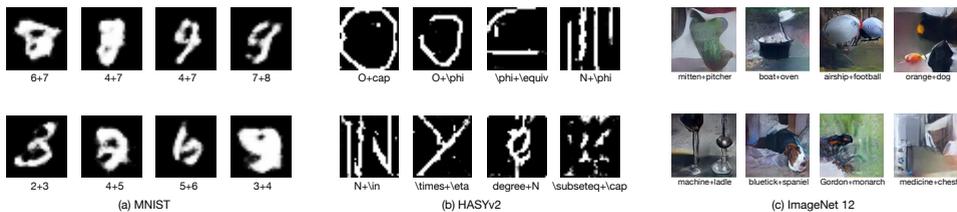


Figure 5: Figure (a) illustrates the synthesis samples of unknown classes. Training classes include: 2-7 from MNIST dataset. It is quite surprising to see that the generator is able to generate outputs like digits 8 and 9 (variants from digits class 3 and 4). Figure (b) shows mixture samples from HASyV2 dataset. Random dots are occasionally being seen from the background. Figure (c) illustrates mixture samples learned from natural image samples. Highest two class activations are displayed at the bottom of the image.

same experiment as in [3] and figured out there are no obvious performance improvement by using G-OpenMax over the natural image setting. A more general reason might be that the generated images are not plausible with respect to the training classes in order to be good candidates to represent unknown classes from open space. More specifically, unlike the MNIST and HASyV2 dataset, ImageNet classes enjoy large variety, therefore not many common features are encoded in the learned latent space. Since our method is based on disentangling class information from other object features, the results it produces in this case do not differ much from interpolation in the pixel space.

5 Conclusion

In this paper, we presented a novel method for multi-class open set classification. The proposed method G-OpenMax extends OpenMax by applying generative adversarial network and is able to provide explicit probability estimation and visualization of unknown classes. Based on the reasonable assumption for open space modelling, which assumes that all classes share common features. Results demonstrate that when this assumption holds, results are indeed better. Moreover, we have visualised the objects and demonstrated that they look plausible with respect to training classes. However, there still remains some unanswered questions to be addressed such as how the complexity of mixture classes affects the performance of G-OpenMax? How to solve the open set classification for natural images? In the future work, we will explore those directions.⁴

References

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein gan. *arXiv preprint arXiv:1701.07875*, 2017.
- [2] Abhijit Bendale and Terrance Boult. Towards open world recognition. In *CVPR*, 2015.
- [3] Abhijit Bendale and Terrance E Boult. Towards open set deep networks. In *CVPR*, 2016.

⁴*Equal contribution of this paper.

- [4] Antonia Creswell and Anil Anthony Bharath. Adversarial training for sketch retrieval. In *ECCV*. Springer, 2016.
- [5] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.
- [6] Chelsea Finn, Ian Goodfellow, and Sergey Levine. Unsupervised learning for physical interaction through video prediction. In *NIPS*, pages 64–72, 2016.
- [7] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *NIPS*, 2014.
- [8] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. in *CVPR*, 2016.
- [9] Kaiming He, Georgia Gkioxari, Piotr Dollár, and Ross Girshick. Mask r-cnn. *arXiv preprint arXiv:1703.06870*, 2017.
- [10] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. *arXiv preprint arXiv:1611.07004*, 2016.
- [11] Lalit P Jain, Walter J Scheirer, and Terrance E Boult. Multi-class open set recognition using probability of inclusion. In *ECCV*, 2014.
- [12] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.
- [13] Alex Krizhevsky, Ilya Sutskever, and Geoff Hinton. Imagenet classification with deep convolutional neural networks. In *NIPS*, 2012.
- [14] Yann LeCun, Corinna Cortes, and Christopher JC Burges. The mnist database of handwritten digits, 1998.
- [15] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. *arXiv preprint arXiv:1609.04802*, 2016.
- [16] Robert M May. How many species are there on earth? *Science*, 241(4872):1441, 1988.
- [17] Alexandru Niculescu-Mizil and Rich Caruana. Predicting good probabilities with supervised learning. In *ICML*, pages 625–632. ACM, 2005.
- [18] Augustus Odena, Christopher Olah, and Jonathon Shlens. Conditional image synthesis with auxiliary classifier gans. *arXiv preprint arXiv:1610.09585*, 2016.
- [19] Guim Perarnau, Joost van de Weijer, Bogdan Raducanu, and Jose M Álvarez. Invertible conditional gans for image editing. *arXiv preprint arXiv:1611.06355*, 2016.
- [20] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *ICLR*, 2016.

-
- [21] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, et al. Imagenet large scale visual recognition challenge. *IJCV*, 115(3):211–252, 2015.
- [22] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved techniques for training gans. In *NIPS*, 2016.
- [23] Walter J Scheirer, Anderson Rocha, Ross J Micheals, and Terrance E Boulton. Meta-recognition: The theory and practice of recognition score analysis. *PAMI*, 33(8):1689–1695, 2011.
- [24] Walter J Scheirer, Anderson de Rezende Rocha, Archana Sapkota, and Terrance E Boulton. Toward open set recognition. *PAMI*, 35(7):1757–1772, 2013.
- [25] Walter J Scheirer, Lalit P Jain, and Terrance E Boulton. Probability models for open set recognition. *PAMI*, 36(11):2317–2324, 2014.
- [26] Richard L Smith. Extreme value theory. *Handbook of applicable mathematics*, 7: 437–471, 1990.
- [27] Alexander J Smola. *Advances in large margin classifiers*. MIT press, 2000.
- [28] Christian Szegedy, Sergey Ioffe, Vincent Vanhoucke, and Alex Alemi. Inception-v4, inception-resnet and the impact of residual connections on learning. *arXiv preprint arXiv:1602.07261*, 2016.
- [29] Martin Thoma. The hasyv2 dataset. *arXiv preprint arXiv:1701.08380*, 2017.