

Open-Set Recognition Using Intra-Class Splitting

Patrick Schlachter, Yiwen Liao and Bin Yang
 Institute of Signal Processing and System Theory
 University of Stuttgart, Germany

Abstract—This paper proposes a method to use deep neural networks as end-to-end open-set classifiers. It is based on intra-class data splitting. In open-set recognition, only samples from a limited number of known classes are available for training. During inference, an open-set classifier must reject samples from unknown classes while correctly classifying samples from known classes. The proposed method splits given data into typical and atypical normal subsets by using a closed-set classifier. This enables to model the abnormal classes by atypical normal samples. Accordingly, the open-set recognition problem is reformulated into a traditional classification problem. In addition, a closed-set regularization is proposed to guarantee a high closed-set classification performance. Intensive experiments on five well-known image datasets showed the effectiveness of the proposed method which outperformed the baselines and achieved a distinct improvement over the state-of-the-art methods.

Index Terms—Open-Set Recognition, Intra-Class Splitting, Deep Learning

I. INTRODUCTION

In recent years, machine learning and deep learning have achieved a huge success in classification [1]. However, most approaches share the assumption that each sample during inference belongs to one of a fixed number of known classes. In other words, these models are trained and evaluated under a closed-set (CS) condition. Unfortunately, such a closed-set environment is ideal and not common in practice. Indeed, many real applications are subject to an open-set (OS) condition as shown in Fig. 1, meaning that some test samples belong to classes that are unknown during training, so-called “unknown unknowns” [2], [3]. For example, in the field of medical image classification, some test images may indicate a certain kind of disease which is unknown in advance. Such images should not be classified as any of the known classes but as belonging to a new abnormal class.

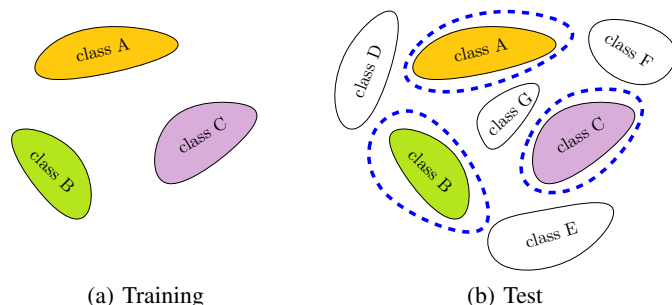


Fig. 1: Open-set recognition: (a) Some known classes (class A, B and C) during training. (b) The unknown classes (class D, E, F and G) should be rejected while the known classes must be discriminated from each other by the learned decision boundaries (blue dashed frames).

Although open-set recognition is a common scenario in practice, it lacks of attention in the past, because it is much harder to solve than closed-set problems. Conventional methods to OS problems are variants of the support vector machine (SVM) such as the 1-vs-set SVM [4] or W-SVM [5]. However, as shown in [5], they are sensitive to the thresholds for rejecting abnormal samples and therefore need abnormal samples to find a proper threshold during training, which is often not possible in practice. Moreover, these methods can only achieve a good performance with extracted features based on expert knowledge, which require a search and have a limited transferable performance. Thus, classical SVM-based methods only achieve a limited performance on complex datasets such as natural images [5].

In contrast to conventional shallow models, deep neural networks such as VGG-16 [6], Inception [7] or ResNet [8] achieved state-of-the-art performance in classification and recognition. Moreover, a generative adversarial network (GAN) is able to generate more realistic images than ever [9]. Intuitively, a modern approach to deal with open-set problems is to generate fake images based on a deep GAN and use them to model the abnormal class. Consequently, an open-set problem is reformulated to a closed-set classification problem [10]. However, GAN based methods bear the following challenges. First, the assumption that generated fake images can represent the unseen abnormal samples is not solid, because it is still an open question whether a GAN structure can really approximate the true data distribution [11]. Second, GANs tend to generate images indistinguishable from the majority of the training dataset. This is not desirable for open-set problems because discriminating those images from the original images leads to a poor closed-set accuracy. Finally, unknown abnormal samples cannot be compactly defined without any prior information as in [10]. Hence, defining a reasonable objective to train a GAN-based OS classifier remains challenging.

Although modeling the unknown abnormal samples by generated fake images has some disadvantages, transforming an OS problem into a classification problem by introducing one additional class for all unknown classes still has a high potential for OS problems. A more natural idea to model abnormal samples is to use a certain part of the given normal samples. Schlachter et al. proposed an intra-class data splitting method which splits the given normal dataset into typical and atypical normal subsets and uses the latter to model the unknown abnormal class [12]. However, this method was originally designed for one-class classification. Therefore, it does not consider the inter-class information in OS problems, meaning the relations among the several known classes as shown in Fig. 2.

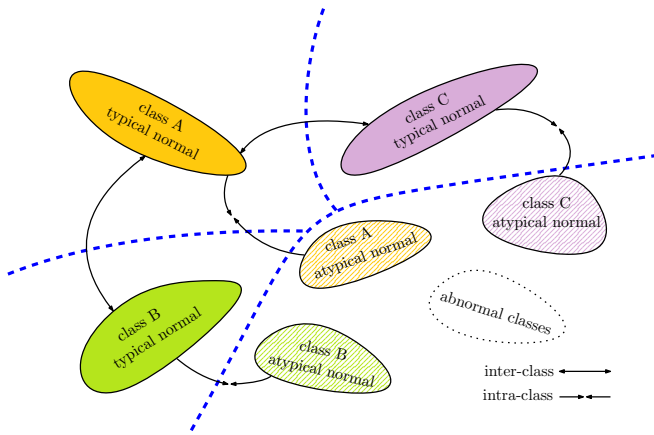


Fig. 2: The basic idea of the proposed method: The given normal dataset (class A, B and C) is split into typical and atypical normal subsets. The atypical normal subsets of all known classes are used to model the unknown abnormal classes and these subsets share the same new label during training. By discriminating the atypical normal samples from the typical normal ones, the trained classifier is expected to reject the abnormal classes as well.

In this paper, we propose a novel deep learning method for open-set recognition problems based on improved intra-class splitting of data. In particular, a given N -class normal dataset is split into typical and atypical normal subsets. Then, the atypical normal samples are used to model the unknown abnormal data. Correspondingly, an OS problem is transformed into an $(N + 1)$ -class classification problem. In order to maintain a high closed-set classification ability, a novel closed-set regularized deep neural network is designed for this $(N + 1)$ -class classification.

Compared to prior work towards open-set problems, our work has three main contributions:

- It is the first work using a small part of the given normal classes to model the unknown abnormal class. Accordingly, only the given normal samples are used during training without generating new fake samples. Therefore, no strong assumptions about the unknown abnormal samples are required. This is helpful for real-world open-set recognition scenarios.
- An improved intra-class splitting method is adapted to open-set recognition problems, which exploits the inter-class information among the given normal classes. The ameliorated splitting method uses the metric of class probability instead of the structural similarity index (SSIM) [12]. Hence, the new splitting method is more general and follows the human understanding as shown in Fig 4.
- We propose a closed-set regularized deep neural network which realizes a high closed-set accuracy while having the ability of rejecting unknown abnormal samples.

II. PROPOSED METHOD

A. Basic idea

The proposed method reformulates an original N -class open-set recognition problem into an $(N + 1)$ -class classification problem. This reformulation is realized by modeling the unknown abnormal class by atypical normal samples obtained through an improved intra-class data splitting. Formally, a given set of samples \mathcal{X}_i , where i indicates any of the known N classes, is split into typical and atypical normal subsets $\mathcal{X}_{i,\text{typical}}$ and $\mathcal{X}_{i,\text{atypical}}$ as illustrated in Fig. 2. Then, the atypical normal subsets of all N classes are considered as one additional class by assigning them a new label during training.

Based on splitting given normal data, a deep neural network (DNN) with $N + 1$ output neurons can then serve as an open-set classifier. However, the atypical normal samples are actually normal samples and their new labels differ from the ground truth. Hence, a naive neural network with $N + 1$ outputs will result in a low closed-set accuracy, because the atypical normal samples are incorrectly predicted. To prevent this situation, we propose a closed-set regularization subnetwork which forces the atypical normal samples to be correctly classified during training. Fig. 3 visualizes the resulting architecture.

During inference, only the deep neural network without the closed-set regularization layer is used as an end-to-end classifier for open-set recognition which outputs a predicted label \hat{y}_i for each input x_i .

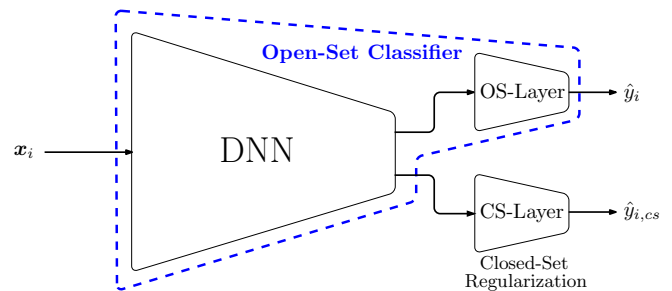


Fig. 3: The architecture of the proposed method. The open-set classifier is based on an arbitrary deep neural network (DNN) with $N + 1$ outputs. The closed-set regularization is modeled by a one-layer network with N outputs in this work.

B. Improved Intra-Class Data Splitting

The original intra-class splitting method [12] trains an autoencoder and uses the reconstruction error of samples as a similarity score to split a given normal dataset. In particular, the samples with lower reconstruction errors are considered as typical normal, whereas the samples with higher reconstruction errors are atypical normal. This method works well for one-class classification problems. However, directly applying this autoencoder-based intra-class splitting method to OS problems is not optimal, because the available inter-class information is not utilized. More precisely, the discriminations among known classes are overlooked during splitting.

In order to take full advantage of the inter-class information in OS problems, we use a multi-class classifier instead of an




Fig. 4: Examples for typical and atypical samples according to intra-class data splitting.

autoencoder for intra-class data splitting. Concretely, an N -class classifier is trained with the given N -class normal data. Once the classifier is trained, the incorrectly predicted samples and the correctly predicted samples with a low probability are selected as atypical normal samples. Thereby, probabilities correspond to the linear activations (logits) of the last layer in a regular deep neural network.

In general, this improved intra-class data splitting method is formulated as follows. Let $f(\cdot)$ indicate the mapping of an N -class neural network. \mathbf{x} denotes a sample from the training dataset. Therefore, the predicted class probabilities under the learned mapping are $\hat{\mathbf{y}}_{\text{prob}} = f(\mathbf{x})$ with $\hat{\mathbf{y}}_{\text{prob}} \in \mathbb{R}^{N \times 1}$. Correspondingly, $\hat{\mathbf{y}}$ is the resulting class prediction in one-hot coding. Furthermore, let $\mathbf{y} \in \mathbb{R}^{N \times 1}$ be the ground truth in one-hot coding and \odot be the element-wise product. Consequently, the score for intra-class splitting is denoted as

$$\text{score} = (\hat{\mathbf{y}}_{\text{prob}} \odot \hat{\mathbf{y}} \odot \mathbf{y})^\top \cdot \mathbf{1}, \quad (1)$$

where $\text{score} \in \mathbb{R}$ and $\mathbf{1} \in \mathbb{R}^{N \times 1}$ is a vector of ones. According to a predefined ratio ρ , the $\rho\%$ samples with the lowest scores are considered as atypical normal samples. The remaining samples are considered to be typical normal.

Therefore, the improved intra-class splitting method has two advantages:

- The improved method is more general than the autoencoder-based one. Indeed, the original intra-class splitting was limited to image datasets due to utilizing SSIM as a similarity metric to split given normal data. In this work, the class probabilities are used as the metric to accomplish the intra-class data splitting. This is more general and can be extended to all kinds of datasets such as time series signals or extracted features.
- The inter-class information is taken into account. By training a multi-class classifier, only samples having low probability scores are selected as atypical normal samples. This splitting procedure matches the human understanding as shown in Fig. 4.

C. Closed-Set Regularization

The neural network for $(N + 1)$ -class classification is an arbitrary regular deep neural network (DNN) with an additional subnetwork acting as a closed-set regularization as shown in Fig. 3.

In this work, the additional closed-set regularization subnetwork only consists of one layer. Hence, the proposed architecture has two separate output layers: the OS-layer and the CS-layer. The OS-layer has $N + 1$ neurons for open-set predictions. In contrast, the CS-layer has N neurons and

serves as a closed-set regularization layer. In particular, the OS-layer reserves an output neuron for the unknown abnormal class which is modeled by the atypical normal samples during training. On the other hand, in order to maintain a high closed-set accuracy as explained above, the CS-layer works as a regularization to force the atypical normal samples to be correctly classified to their own classes.

Consequently, the objective of the $(N + 1)$ -class neural network is learning to classify the training samples into the $N + 1$ classes under the constraint that all training samples are still able to be classified into the given known N classes using a simple one-layer subnetwork, see below.

D. Loss Functions

The objective of the entire network is transformed into a joint optimization problem and consists of two individual loss terms for the OS-layer and CS-layer as

$$\mathcal{L} = \mathcal{L}_{\text{os}} + \gamma \cdot \mathcal{L}_{\text{cs}}, \quad (2)$$

where \mathcal{L}_{os} is the loss function for the OS-layer and \mathcal{L}_{cs} is the loss function for the CS-layer. γ is a hyperparameter to tune the ratio between these two terms.

Let B be the minibatch size during training. Moreover, $\mathbb{1}_{y_i \in y^{(n)}}$ is an indicator function which returns 1 if a given sample \mathbf{x}_i with a scalar label y_i belongs to the class $y^{(n)}$ and otherwise returns 0. Based on these notations, the two loss terms are introduced as:

a) *OS-Loss*: The open-set problem is transformed into an $(N + 1)$ -class classification problem due to the intra-class splitting. Therefore, the OS-loss is a simple $(N + 1)$ -class categorical cross-entropy loss

$$\mathcal{L}_{\text{os}} = -\frac{1}{B} \sum_{i=1}^B \sum_{n=1}^{N_{\text{os}}} \mathbb{1}_{y_i \in y^{(n)}} \log[P(\hat{y}_i \in y^{(n)})], \quad (3)$$

where $N_{\text{os}} = N + 1$ and $P(\hat{y}_i \in y^{(n)})$ denotes the predicted probability that sample \mathbf{x}_i belongs to the class $y^{(n)}$, meaning the value of the n -th element of the output vector of the network.

b) *CS-Loss*: The closed-set regularization loss is an N -class categorical cross entropy loss

$$\mathcal{L}_{\text{cs}} = -\frac{1}{B} \sum_{i=1}^B \sum_{n=1}^{N_{\text{cs}}} \mathbb{1}_{y_i \in y^{(n)}} \log[P(\hat{y}_i \in y^{(n)})], \quad (4)$$

where \mathcal{L}_{cs} shares the same notation as \mathcal{L}_{os} and $N_{\text{cs}} = N$ is the number of the given known classes.

III. EXPERIMENTS

A. Setup

As a basic experiment, the proposed method was first evaluated on MNIST [13], SVHN [14] and CIFAR-10 [15]. Each dataset has 10 classes and 6 of them were randomly selected as the known classes during training. Certainly, the test set consisted of the known 6 classes and the left 4 unknown classes. We repeated this basic experiment with 5 different seeds, i.e. 5 random combinations of the known classes. In order to evaluate the robustness of the proposed method over different openness, the second experiment was

to train the model with 6 known classes in CIFAR-10 and test it with different numbers of unknown classes from CIFAR-100 [15] and Tiny ImageNet [16]. Finally, as ρ is the key hyperparameter in our method, the sensitivity to ρ was further evaluated with the same settings as the basic experiment except that 6 different ρ were tested.

Balanced accuracy [17] was selected as the primary metric for evaluation, because it allows a fair comparison of balanced and imbalanced datasets which can both occur in OS problems. In this work, the unknown or abnormal classes are considered as negative while the known classes are considered as positive. Consequently, the balanced accuracy for OS problem is defined as

$$BACC = \frac{1}{2} \cdot \left(\frac{TP}{TP + FN} + \frac{TN}{FP + TN} \right), \quad (5)$$

where TP is the number of “true positives”. In contrast to binary problems, TP represents those samples which are correctly classified as one of the known classes and not only samples that are correctly classified as positive.¹

We selected the following four baseline models including one state-of-the-art method based on the generation of counterfactual images from the literature:

- 1) WSVM: Weibull support vector machine with the default settings of the libsvm-openset package [18].
- 2) OCSVM: An N -class network was trained for closed-set prediction, while a separate one-class SVM [19] with $\nu = 0.1$ was trained on the training dataset for rejecting abnormal samples. The final results were the multiplication of the predictions from both classifiers.
- 3) GAN: $(N+1)$ -class neural network using fake images as abnormal class which are generated by a regular GAN.
- 4) CF: Same settings with counterfactual image generation method in [10].

Furthermore, two variants of the proposed method were evaluated to judge the effectiveness of closed-set regularization and improved intra-class data splitting:

- 5) NN-ics: An $(N + 1)$ -class neural network combined with the intra-class data splitting method but without any closed-set regularization layers.
- 6) AE-ics: Same settings with the proposed method except that an autoencoder was used for intra-class data splitting as in [12].

Note that the baselines OCSVM, GAN, NN-ics and AE-ics shared the same architecture with the proposed method for a fair comparison and they were implemented by scikit-learn [20] and TensorFlow [21].

The proposed method used a modified VGG-16 [6] as a backbone with residual blocks [8] to reduce the number of network parameters. L2-regularization was used for each convolutional layer with a decay of 10^{-3} . γ was equal to 1 for the entire loss function. The splitting ratio ρ was selected as 10 for MNIST and 20 for SVHN and CIFAR-10. Finally, the batch size was 32 and the model was trained for 50 epochs.

¹Correspondingly, FN are “false positives”, TN are “true negatives” and FP are “false positives”.

B. Basic Experiments

The basic experimental results are listed in Table I. The proposed method outperformed other baselines including the state-of-the-art methods in all conducted experiments.

CIFAR-10, as a natural image dataset, is challenging in OS problems. The conventional shallow model WSVM or even the state-of-the-art method CF only reached a balanced accuracy of about 50%. In comparison, our method achieved a balanced accuracy of more than 71%, which corresponds to an improvement of 39% over the other considered methods.

Considering the less difficult image datasets, MNIST and SVHN, both shallow and deep models showed a good performance. However, the proposed method still had the best performance with about 8% higher balanced accuracy.

Interestingly, there was a huge gap between the performances of a regular GAN and that of CF. This showed the difficulty in designing a correct objective for generating fake samples to represent the unknown abnormal data, because there is no prior information of the unknowns during training. In contrast, our method only uses a part of the training dataset to model the abnormal samples which does not require any prior information.

Eventually, the proposed method with an improved intra-class splitting achieved a better performance than the baseline using autoencoder-based splitting as expected.

TABLE I: Balanced Accuracy (standard deviation) in %.

Dataset	WSVM	OCSVM	GAN	CF	AE-ics	Ours
MNIST	84.6 (± 3.5)	64.5 (± 5.0)	55.6 (± 2.6)	87.5 (± 2.0)	82.7 (± 2.5)	94.3 (± 0.4)
SVHN	75.2 (± 3.3)	49.2 (± 0.6)	48.4 (± 1.1)	76.2 (± 4.6)	72.2 (± 3.3)	82.8 (± 0.5)
CIFAR-10	46.5 (± 4.1)	50.0 (± 3.0)	43.5 (± 4.0)	51.2 (± 0.7)	50.2 (± 4.8)	71.2 (± 2.1)

C. Performance with Different Openness

Following [10], the openness of the OS problem is defined as

$$openness = 1 - \sqrt{\frac{N}{M}}, \quad (6)$$

where N denotes the number of the known normal classes during training and M is the total number of the encountered classes during testing, i.e. $M = N + N'$ with N' denoting the number of unknown abnormal classes during testing. To test the proposed method’s robustness to different openness, we used the following settings. First, $N = 6$ for all cases. Moreover, the number of abnormal classes N' during testing was chosen as 20, 30, 50, 100 and 200. In the former four cases, we trained the model on CIFAR-10 and tested it on CIFAR-100. In the last case, we trained the model on CIFAR-10 and tested it on Tiny ImageNet. The corresponding results are listed in Table II. Here we only compare our method with the state-of-the-art method CF from the basic experiment and the variant NN-ics of the proposed method.

Regarding challenging natural image datasets such as CIFAR-100 and Tiny ImageNet, the proposed method outperformed CF in all cases with high balanced accuracies.

Note that the variant of the proposed method NN-ics, i.e. a naive neural network with intra-class data splitting method, already performed better than CF in all considered cases.

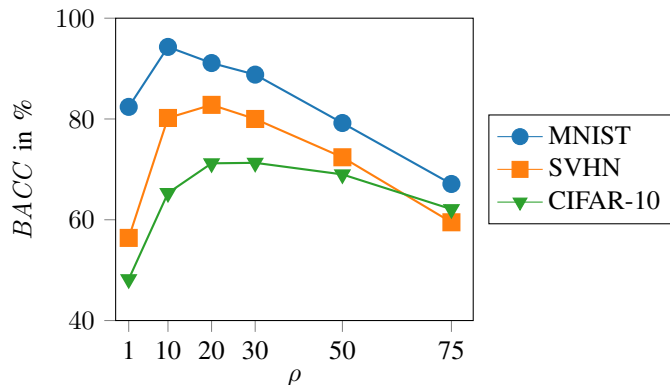


Fig. 5: The balanced accuracy vs. different splitting ratios ρ .

Interestingly, all three methods showed a stable performance with different openness.

D. Sensitivity to the splitting ratios

The splitting ratio ρ is a crucial hyperparameter for the proposed method. Fig. 5 shows the performance regarding different ratios. As expected, both a very low ratio and a very high ratio lead to a worse performance than proper ratios. A very low ratio, e.g. $\rho = 1$, means that only a small part of the training data is used to represent abnormal data. Therefore, the training procedure is highly imbalanced and the trained model cannot gain adequate gradient information for the additional abnormal class during training. Consequently, the model can poorly identify abnormal samples during testing. On the other hand, a large ratio, e.g. $\rho = 75$, causes too many normal samples to be incorrectly predicted as abnormal which results in a low closed-set accuracy.

From another perspective, the optimal ρ is also an indicator for the homogeneity of a dataset. For instance, MNIST is more homogenous and hence requires less atypical samples than SVHN and CIFAR-10 which results in a smaller value for the optimal ρ .

Although ρ plays an important role, our method is not very sensitive to this ρ in a wide range. For example, as illustrated in Fig. 5, the proposed method has a stable performance on SVHN with $\rho \in [10, 30]$.

IV. CONCLUSION

We proposed a novel deep learning method for open-set recognition. By using intra-class data splitting, it allows to introduce a categorical cross-entropy loss and additional closed-set regularizations. As a result, the proposed method allows end-to-end training of regular deep neural networks for open-set recognition. Our method was evaluated in a large amount of experiments with natural images. It showed a distinct improvement over state-of-the-art methods towards open-set recognition in average. Future work may integrate GANs into the proposed method to generate more atypical normal samples which further increase its performance and robustness. Furthermore, more realistic datasets such as fingerprints or face images will be used to evaluate the proposed method. Finally, we will also evaluate the proposed method on non-image datasets such as radar signals.

TABLE II: Balanced Accuracy vs. Openness in %.

Method	$N' = 20$	$N' = 30$	$N' = 50$	$N' = 100$	$N' = 200$
CF	52.0 (± 2.1)	53.2 (± 2.0)	52.1 (± 3.0)	52.6 (± 3.1)	52.5 (± 1.0)
NN-ics	57.1 (± 2.3)	57.8 (± 2.0)	57.7 (± 1.9)	58.2 (± 1.8)	58.4 (± 2.5)
Ours	69.5 (± 1.9)	70.0 (± 1.9)	70.3 (± 1.6)	70.8 (± 1.5)	70.1 (± 1.7)

REFERENCES

- [1] Y. LeCun, Y. Bengio, and G. Hinton, "Deep learning," *nature*, vol. 521, no. 7553, p. 436, 2015.
- [2] A. Bendale and T. E. Boulton, "Towards open set deep networks," *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 1563–1572, 2016.
- [3] C. Geng, S. Huang, and S. Chen, "Recent advances in open set recognition: A survey," vol. abs/1811.08581, 2018. [Online]. Available: <http://arxiv.org/abs/1811.08581>
- [4] W. J. Scheirer, A. Rocha, A. Sapkota, and T. E. Boulton, "Towards open set recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)*, vol. 35, July 2013.
- [5] W. J. Scheirer, L. P. Jain, and T. E. Boulton, "Probability models for open set recognition," *IEEE Transactions on Pattern Analysis and Machine Intelligence (T-PAMI)*, vol. 36, November 2014.
- [6] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," *CoRR*, vol. abs/1409.1556, 2014.
- [7] C. Szegedy, W. Liu, Y. Jia, P. Sermanet, S. Reed, D. Anguelov, D. Erhan, V. Vanhoucke, and A. Rabinovich, "Going deeper with convolutions," in *Proceedings of the IEEE conference on computer vision and pattern recognition*, 2015, pp. 1–9.
- [8] K. He, X. Zhang, S. Ren, and J. Sun, "Deep residual learning for image recognition," *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 770–778, 2016.
- [9] I. Goodfellow, J. Pouget-Abadie, M. Mirza, B. Xu, D. Warde-Farley, S. Ozair, A. Courville, and Y. Bengio, "Generative adversarial nets," in *Advances in Neural Information Processing Systems 27*, Z. Ghahramani, M. Welling, C. Cortes, N. D. Lawrence, and K. Q. Weinberger, Eds. Curran Associates, Inc., 2014, pp. 2672–2680. [Online]. Available: <http://papers.nips.cc/paper/5423-generative-adversarial-nets.pdf>
- [10] L. Neal, M. Olson, X. Fern, W.-K. Wong, and F. Li, "Open set learning with counterfactual images," in *The European Conference on Computer Vision (ECCV)*, September 2018.
- [11] E. Nalisnick, A. Matsukawa, Y. W. Teh, D. Gorur, and B. Lakshminarayanan, "Do deep generative models know what they don't know?" *arXiv preprint arXiv:1810.09136*, 2018.
- [12] P. Schlachter, Y. Liao, and B. Yang, "Deep one-class classification using data splitting," in *2019 IEEE Data Science Workshop (DSW)*, June 2019.
- [13] Y. LeCun, L. Bottou, Y. Bengio, and P. Haffner, "Gradient-based learning applied to document recognition," *Proceedings of the IEEE*, vol. 86, no. 11, pp. 2278–2324, 1998.
- [14] Y. Netzer, T. Wang, A. Coates, A. Bissacco, B. Wu, and A. Y. Ng, "Reading digits in natural images with unsupervised feature learning," 2011.
- [15] A. Krizhevsky, "Learning multiple layers of features from tiny images," Citeseer, Tech. Rep., 2009.
- [16] O. Russakovsky, J. Deng, H. Su, J. Krause, S. Satheesh, S. Ma, Z. Huang, A. Karpathy, A. Khosla, M. Bernstein, A. C. Berg, and L. Fei-Fei, "ImageNet Large Scale Visual Recognition Challenge," *International Journal of Computer Vision (IJCV)*, vol. 115, no. 3, pp. 211–252, 2015.
- [17] K. H. Brodersen, C. S. Ong, K. E. Stephan, and J. M. Buhmann, "The balanced accuracy and its posterior distribution," in *Pattern recognition (ICPR), 2010 20th international conference on*. IEEE, 2010, pp. 3121–3124.
- [18] W. J. Scheirer, A. Rocha, R. Michaels, and T. E. Boulton, "Meta-recognition: The theory and practice of recognition score analysis," *IEEE Transactions on Pattern Analysis and Machine Intelligence (PAMI)*, vol. 33, pp. 1689–1695, 2011.
- [19] B. Schölkopf, J. C. Platt, J. Shawe-Taylor, A. J. Smola, and R. C. Williamson, "Estimating the support of a high-dimensional distribution," *Neural computation*, vol. 13, no. 7, pp. 1443–1471, 2001.
- [20] F. Pedregosa et al., "Scikit-learn: Machine learning in Python," *Journal of Machine Learning Research*, vol. 12, pp. 2825–2830, 2011.
- [21] M. Abadi, P. Barham, J. Chen, Z. Chen, A. Davis, J. Dean, M. Devin, S. Ghemawat, G. Irving, M. Isard et al., "Tensorflow: a system for large-scale machine learning," in *Proceedings of the 12th USENIX conference on Operating Systems Design and Implementation*. USENIX Association, 2016, pp. 265–283.