

One-Class Convolutional Neural Network

Poojan Oza, *Student Member, IEEE*, and Vishal M. Patel, *Senior Member, IEEE*

Abstract—We present a novel Convolutional Neural Network (CNN) based approach for one class classification. The idea is to use a zero centered Gaussian noise in the latent space as the pseudo-negative class and train the network using the cross-entropy loss to learn a good representation as well as the decision boundary for the given class. A key feature of the proposed approach is that any pre-trained CNN can be used as the base network for one class classification. The proposed One Class CNN (OC-CNN) is evaluated on the UMDAA-02 Face, Abnormality-1001, FounderType-200 datasets. These datasets are related to a variety of one class application problems such as user authentication, abnormality detection and novelty detection. Extensive experiments demonstrate that the proposed method achieves significant improvements over the recent state-of-the-art methods. The source code is available at : github.com/otkupjnoz/oc-cnn.

Index Terms—One Class Classification, Convolutional Neural Networks, Representation Learning.

I. INTRODUCTION

Multi-class classification entails classifying an unknown object sample into one of many pre-defined object categories. In contrast, in one-class classification, the objective is to identify objects of a particular class (also known as positive class data or target class data) among all possible objects by learning a classifier from a training set consisting of only the target class data. The absence of data from the negative class(es) makes the one-class classification problem difficult.

One-class classification has many applications such as anomaly or abnormality detection [1], [2], [3], [4], [5], novelty detection [6], [7], [8], and user authentication [9], [10], [11], [12], [13], [14]. For example, in novelty detection, it is normally assumed that one does not have *a priori* knowledge of the novel class data. Hence, the learning process involves only the target class data.

Various methods have been proposed in the literature for one-class classification. In particular, many one-class classification methods are based on the Support Vector Machines (SVM) formulation [15], [8], [16]. SVMs are based on the concept of finding a boundary that maximizes the margin between two classes and are shown to work well for binary and multi-class classification. However, in one-class problems the information regarding the negative class data is unavailable. To deal with this issue, Scholkopf et al. [17] proposed one-class SVM (OC-SVM), which tackles the absence of negative class data by maximizing the boundary with respect to the origin. Another popular approach inspired by the SVM formulation is Support Vector Data Description (SVDD) introduced by Tax et al. [18], in which a hypersphere that encloses the target class data is sought. Various extensions of OC-SVM




Fig. 1: A graphical illustration of popular statistical one-class classification methods. Green circles show the target class data, red crosses show the unknown data (i.e. anomaly, novelty, outlier etc.), blue dotted lines/circles show the decision boundaries captured by the respective methods. Pink dotted line in Fig. 1(c) shows the boundary of zero error of probability. (a) OC-SVM, maximizing the margin of a hyperplane with respect to the origin. (b) SVDD, finding a hypersphere that encloses the given data. (c) MPM, finding a hyperplane that minimizes the misclassification probability.

and SVDD have been proposed in the literature over the years. We refer readers to [19] for a survey of different one-class classification methods. Another approach for one-class classification is based on the Minimax Probability Machines (MPM) formulation [20]. Single class MPM [21], [22] seeks to find a hyper-plane similar to that of OC-SVM by taking second order statistics of data into consideration. Hence, single class MPM learns a decision boundary that generalizes well to the underlying data distribution. Fig. 1 presents a high-level overview of different one-class classification methods. Though, these approaches are powerful tools in identifying the decision boundary for target data, their performance depends on the features used to represent the target class data.

Over the last five years, methods based on Deep Convolutional Neural Networks (DCNNs) have shown impressive performance improvements for object detection and recognition problems. Taking classification task as an example, the top-5 error rate of vision systems on the ImageNet dataset [23] has dropped from $\sim 25\%$ to 2.25% in the last five years. This has been made possible due to the availability of large annotated datasets, a better understanding of the non-linear mapping between input images and class labels as well as the affordability of Graphics Processing Units (GPUs). These networks learn distinct features for a particular class against another using the cross entropy loss. However, for one class problems training such networks in an end-to-end manner becomes difficult due to the absence of negative class data.

In recent years, several attempts have been made to counter the problem of training a neural network for one-class classification [5], [24], [25], [26], [27], [28], [29], [30]. These approaches can be broadly classified in to two categories, generative approaches [27], [28], [29] and discriminative ap-

P. Oza and V. M. Patel are with the Department of Electrical and Computer Engineering, Johns Hopkins University, Baltimore, MD, USA. Email: {poza2, vpatel36}@jhu.edu. This work was supported by the NSF grant 1801435.

proaches [25], [30]. Generative approaches use generative frameworks such as auto-encoders or Generative Adversarial Networks (GAN) [31] for one-class classification. For example, Ravanbakhsh et al. [27] and Sabokrou et al. [24] proposed deep auto-encoder networks for event anomaly detection in surveillance videos. However, in their approaches the focus is mainly on the image-level one-class classification. Work by Lawson et al. [26] developed a GAN-based approach for abnormality detection. Sabokrou et al. [5] extended that idea for detecting outliers from image data using an auto-encoder based generator with adversarial training. In general, these generative models such as GANs are very difficult to train as compared to the discriminative classification networks [32].

Compared to the generative approaches, discriminative approaches for one-class classification have not been well explored in the literature. One such approach by Perera and Patel [30] utilize an external reference dataset as the negative class to train a deep network for one-class classification using a novel loss function. In contrast to this method, we do not make use of any negative class data in our approach. In another approach, Chalapathy et al. [25] proposed a novel SVM inspired loss function to train a neural network for anomaly detection. With some inspirations from other statistical approaches for one-class classification (i.e. taking origin as a reference to find the decision boundary), we propose a novel method called, One-Class CNN (OC-CNN), to learn representations for one-class problems with CNNs trained end-to-end in a discriminative manner. This paper makes the following contributions:

- A new approach is proposed based on CNN for one class classification which is end-to-end trainable.
- Through experiments, we show that proposed approach outperforms other statistical and deep learning-based one class classification methods and generalizes well across a variety of one class applications.

II. PROPOSED APPROACH

Fig. 2 gives an overview of the proposed CNN-based approach for one-class classification. The overall network consists of a feature extractor network and a classifier network. The feature extractor network essentially embeds the input target class images into a feature space. The extracted features are then appended with the pseudo-negative class data, generated from a zero centered Gaussian in the feature space. The appended features are then fed into a classification network which is characterized by a fully connected neural network. The classification network assigns a confidence score for each feature representation. The output of the classification network is either 1 or 0. Here, 1 corresponds to the data sample belonging to the target class and 0 corresponds to the data sample belonging to the negative class. The entire network is trained end-to-end using binary cross-entropy loss.

A. Feature Extractor

Any pre-trained CNN can be used as the feature extractor. In this paper, we use the pre-trained AlexNet [33] and VGG16 [34] networks by removing the softmax regression layers




Fig. 2: Block diagram of the proposed approach. Here, $\bar{\mu}$ and σ are mean and standard deviation parameters of a Gaussian, respectively and \mathbf{I} is the identity matrix.

(i.e. the last layer) from their networks. During training, we freeze the convolution layers and only train the fully-connected layers. Assuming that the extracted features are D -dimensional, the features are appended with the pseudo-negative data generated from a Gaussian, $\mathcal{N}(\bar{\mu}, \sigma^2 \cdot \mathbf{I})$, where σ and $\bar{\mu}$ are the parameters of the Gaussian and \mathbf{I} is a $D \times D$ identity matrix. Here, $\mathcal{N}(\bar{\mu}, \sigma^2 \cdot \mathbf{I})$ can be seen as generating D independent one dimensional gaussian with σ standard deviation.

B. Classification Network

Due to the appending of the pseudo-negative data with the original features, the classifier network observes the input in the batch size of 2. A simple fully-connected layer followed by a softmax regression layer is used as the classifier network. The dimension of the fully-connected layer is kept the same as the feature dimension. The number of outputs from the softmax layer are set equal to two.

C. Loss Function

The following binary cross-entropy loss function is used to train the entire network

$$L_c = -\frac{1}{2K} \sum_{j=1}^{2K} (y \log(p) + (1-y) \log(1-p)), \quad (1)$$

where, $y \in \{0, 1\}$ indicates whether the classifier input corresponds to the feature extractor, (i.e. $y = 0$), or it is sampled from $\mathcal{N}(\bar{\mu}, \sigma^2 \cdot \mathbf{I})$, (i.e. $y = 1$). Here, p denotes the softmax probability of $y = 0$.

The network is optimized using the Adam optimizer [35] with learning rate of 10^{-4} . The input image batch size of 64 is used in our approach. For all experiments, the parameters $\bar{\mu}$ and σ are set equal to $\bar{0}$ and 0.01, respectively. Instance normalization [36] is used before the classifier network as it was found to be very useful in stabilizing the training procedure.

III. EXPERIMENTAL RESULTS

We evaluate the performance of the proposed approach on three different one-class classification problems - abnormality detection, face-based user authentication, and novelty detection. Abnormality-1001 [37], UMDAA-02 [38] and




Fig. 3: (a) Sample images from the Abnormality-1001 dataset [37]. Normal and abnormal images are shown on the top row and the bottom row, respectively. (b) Samples images from the UMDAA-02 Face Dataset [38]. (c) Sample images from the FounderType-200 dataset [39]. Different types of Chinese fonts.

FounderType-200 [39] datasets are used to conduct experiments for the abnormality detection, user authentication and novelty detection problems. For all methods compared here, the data is aligned such that objects are at the center with minimal background.

The proposed approach is compared with following one-class classification methods:

- **OC-SVM:** One-Class Support Vector Machine is used as formulated in [15], trained using the AlexNet and VGG16 features.
- **BSVM:** Binary SVM is used where the zero centered Gaussian noise is used as the negative data. AlexNet and VGG16 features extracted from the target class data are used as the positive class data.
- **MPM:** MiniMax Probability Machines are used as formulated in [20]. Since, the MPM algorithm involves computing covariance matrix from the data, Principal component analysis (PCA) is used to reduce the dimensionality of the features before computing the covariance matrix.
- **SVDD:** Support Vector Data Description is used as formulated in [18], trained on the AlexNet and VGG16 features.
- **OC-NN:** One-class neural network (OC-NN) is used as formulated in [25]. Here, for fair comparison, instead of using the feature extractor trained using an auto-encoder (as per [25] methodology), AlexNet and VGG16 networks, the same as the proposed method, are used. As described in [25], we evaluate OC-NN using three different activation functions - linear, Sigmoid and ReLU.
- **OC-CNN:** One-class CNN is the method proposed in this paper.
- **OC-SVM⁺:** OCSVM⁺ is another method used in this paper where OC-SVM is utilized on top of the features extracted from the network trained using OC-CNN. However, since it uses OC-SVM for classification, it is not end-to-end trainable.

A. Abnormality Detection

Abnormality detection (also referred as anomaly detection or outlier rejection) deals with identifying instances that are dissimilar to the target class instances (i.e. abnormal instances). Note that, the abnormal instances are not known a priori and only the normal instances are available during training. Such problem can be addressed by one-class classification

algorithms. The Abnormality-1001 dataset [37] is widely used for visual abnormality detection. This dataset consists of 1001 abnormal images belonging to six classes such as Chair, Car, Airplane, Boat, Sofa and Motorbike which have their respective normal classes in the PASCAL VOC dataset [40]. Sample images from the Abnormality-1001 dataset are shown in Fig. 3 (a). Normal images obtained from the PASCAL VOC dataset are split into train and test sets such that the number of abnormal and normal images in test set are equal. Reported results are averaged for all six classes.

B. User Active Authentication

Active authentication refers to the problem of identifying the enrolled user based on his/her biometric data such as face, swipe patterns, and accelerometer patterns [13]. The problem can be viewed as identifying the abnormal user behaviour to reject the unauthorized user. The active authentication problem has been viewed as one-class classification problem [12]. The UMDAA-02 dataset [38] is widely used dataset for user active authentication on mobile devices. The UMDAA-02 dataset has multiple modalities corresponding to each user such as face, accelerometer, gyroscope, touch gestures, etc. Here, we only use the face data provided in this dataset since face is one of the most commonly used modality for authentication. The face data consists of 33209 face images corresponding to 48 users. Sample images corresponding to a few subjects from this dataset are shown in Fig. 3(b). As can be seen from this figure, the images contains large variations in pose, illumination, appearance, and occlusions. For each class, train and test sets are created by maintaining 80/20 ratio. Network is trained using the train set of a target user and tested on the test set of the target user against the rest of the user test set data. This process is repeated for all the users and average results are reported.

C. Novelty Detection

The FounderType-200 dataset was introduced for the purpose of novelty detection by Liu et al. in [39]. The FounderType-200 dataset, contains 6763 images from 200 different types of fonts created by the company FounderType. Fig. 3(c) shows some sample images from this dataset. For experiments, first 100 classes are used as the target classes and remaining 100 classes are used as the novel data. The first 100 class data are split into train and test set having equal number of images. For novel data, a novel set is created having 50

Dataset	OC-SVM	BSVM	MPM	SVDD	OC-NN-lin	OC-NN-sig	OC-NN-relu	OC-CNN	OC-SVM ⁺
Abnormality-1001	0.6057	0.6126	0.5806	0.7873	0.8090	0.6391	0.7372	<i>0.8264</i>	0.8334
UMDAA-02 Face	0.5746	0.5660	0.5418	0.6448	0.6173	0.6452	0.5943	0.7017	<i>0.6736</i>
FounderType-200	0.7124	0.7067	0.7085	0.8998	0.8884	0.8696	0.8505	<i>0.9303</i>	0.9350

TABLE I: Comparison between the proposed and other methods using **AlexNet** as the base network. Results are mean of performance on all classes. Best and the second best performance are highlighted in bold fonts and italics, respectively.

Dataset	OC-SVM	BSVM	MPM	SVDD	OC-NN-lin	OC-NN-sig	OC-NN-relu	OC-CNN	OC-SVM ⁺
Abnormality-1001	0.6475	0.6418	0.5909	0.8031	0.7740	0.8373	0.5821	<i>0.8424</i>	0.8460
UMDAA-02 Face	0.5829	0.5751	0.5473	0.6424	0.6193	0.6200	0.5788	0.7350	<i>0.7230</i>
FounderType-200	0.7490	0.7067	0.7444	0.8885	0.8986	0.8677	0.8506	<i>0.9290</i>	0.9419

TABLE II: Comparison between proposed and other methods using **VGG16** as the base network. Results are mean of performance on all classes. Best and the second best performance are highlighted in bold fonts and italics, respectively.

images from each of the novel classes. For each class, train set from the known data is used for training the network and known class test set and novel set data are used for evaluation. For example, class i ($i \in \{1, 2, \dots, 100\}$) train set is used for training the network. The trained network is then evaluated with class i test set tested against the novel set (containing data of class 101-200). This is repeated for all classes i where, $i \in \{1, 2, \dots, 100\}$ and average results are reported.

IV. RESULTS AND DISCUSSION

The performance is measured using the area under the receiver operating characteristic (ROC) curve (AUROC), most commonly used metric for one-class problems. The results are tabulated in Table II and Table I corresponding to the VGG16 and AlexNet networks. AlexNet and VGG16 pretrained features are used to compute the results for OC-SVM, BSVM, SVDD and MPM. The OC-NN results are computed using the linear, sigmoid and relu activations after training on the target class data. The OC-CNN results are computed after training on the target class and for OC-SVM⁺, an one-class SVM is trained on top of the features extracted from the trained AlexNet/VGG16, and AUROC is computed from the SVM classifier scores.

From the Tables I and II, it can be observed that either OC-CNN or OC-SVM⁺ achieves the best performance on all three datasets. MPM and OC-SVM achieve similar performances, while BSVM with Gaussian data as the negative class doesn't work as well. With the BSVM baseline, we show that similar trick we used for proposed algorithm doesn't work well for statistical approaches like SVM. Among the other one-class approaches, OC-NN with linear activation performs the best. However, OC-NN results are inconsistent. For couple of experiments, SVDD was found to be working better than OC-NN. The reason behind this inconsistent performance can be due to the differences in the evaluation protocol used for OC-NN in [25] and this paper. The ratio of the number of target class images to novel/abnormal class images in our evaluation protocol is much higher than the ratio used by Chalpathy et al. [25]. When the ratio is close to one, as is the case for Abnormality-1001 dataset, the OC-NN performs better than SVDD for both AlexNet and VGG16. However, when the ratio is increased (which is more realistic scenario), as is the case for UMDAA-02 and FounderType-200, the performance of OC-NN becomes inconsistent. Whereas, using the proposed approach performs consistently well, providing $\sim 4\%$, $\sim 10\%$

and $\sim 5\%$ improvements over OC-NN for Abnormality-1001, UMDAA02- Face and FounderType-200 datasets, respectively. Since, the proposed approach is built upon the traditional discriminative learning framework for deep neural networks, it is able to learn better features than OC-NN.

Also as expected, methods based on the VGG16 network work better than the methods based on the AlexNet network. Apart from the FounderType-200 dataset where, OC-CNN with AlexNet works better than VGG16, for all methods VGG16 works better than AlexNet. However, it should be noted that better OC-SVM⁺ performance for VGG16 indicates that features learned with the proposed approach for VGG16 are better than AlexNet for FounderType-200. Overall, VGG16 gives $\sim 2\%$ improvement over AlexNet.

Another interesting comparison is between OC-SVM and OC-SVM⁺. OC-SVM uses features extracted from a pre-trained AlexNet/VGG16 network. On the other hand, OC-SVM⁺ uses features extracted from AlexNet/VGG16 network trained using the proposed approach. OC-SVM⁺ performs $\sim 18\%$ and $\sim 17\%$ better than OC-SVM on average across all datasets for AlexNet and VGG16, respectively. This result shows the ability of our approach to learn better representations. So, apart from being an end-to-end learnable standalone system, our approach can also be used to extract target class friendly features. Also, using sophisticated classifier has shown to improve the performance over OC-CNN (i.e., OC-SVM⁺) in majority of cases.

V. CONCLUSION

We proposed a new one-class classification method based on CNNs. A pseudo-negative Gaussian data was introduced in the feature space and the network was trained using a binary cross-entropy loss. Apart from being a standalone one-class classification system, the proposed method can also be viewed as good feature extractor for the target class data (i.e. OCSVM⁺) as well. Furthermore, the consistent performance improvements over all the datasets related to authentication, abnormality and novelty detection showed the ability of our method to work well on a variety of one-class classification applications. In this paper, experiments were performed over data with objects centrally aligned. In the future, we will explore the possibility of developing an end-to-end deep one class method that does joint detection and classification.

REFERENCES

- [1] V. Chandola, A. Banerjee, and V. Kumar, "Anomaly detection: A survey," *ACM computing surveys (CSUR)*, vol. 41, no. 3, p. 15, 2009.
- [2] L. Akoglu, H. Tong, and D. Koutra, "Graph based anomaly detection and description: a survey," *Data mining and knowledge discovery*, vol. 29, no. 3, pp. 626–688, 2015.
- [3] V. A. Sindagi and S. Srivastava, "Domain adaptation for automatic oled panel defect detection using adaptive support vector data description," *International Journal of Computer Vision*, vol. 122, no. 2, pp. 193–211, 2017.
- [4] C. You, D. P. Robinson, and R. Vidal, "Provable selfrepresentation based outlier detection in a union of subspaces," in *Proceedings of the 2017 IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 1–10.
- [5] M. Sabokrou, M. Khalooei, M. Fathy, and E. Adeli, "Adversarially learned one-class classifier for novelty detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2018, pp. 3379–3388.
- [6] D. Abati, A. Porrello, S. Calderara, and R. Cucchiara, "And: Autoregressive novelty detectors," *arXiv preprint arXiv:1807.01653*, 2018.
- [7] M. A. Pimentel, D. A. Clifton, L. Clifton, and L. Tarassenko, "A review of novelty detection," *Signal Processing*, vol. 99, pp. 215–249, 2014.
- [8] M. Markou and S. Singh, "Novelty detection: a reviewpart 1: statistical approaches," *Signal processing*, vol. 83, no. 12, pp. 2481–2497, 2003.
- [9] M. E. Fathy, V. M. Patel, and R. Chellappa, "Face-based active authentication on mobile devices," in *Acoustics, Speech and Signal Processing (ICASSP), 2015 IEEE International Conference on*. IEEE, 2015, pp. 1687–1691.
- [10] P. Perera and V. M. Patel, "Quickest intrusion detection in mobile active user authentication," in *Biometrics Theory, Applications and Systems (BTAS), 2016 IEEE 8th International Conference on*. IEEE, 2016, pp. 1–8.
- [11] G. Guo, L. Wen, and S. Yan, "Face authentication with make-up changes," *IEEE Transactions on Circuits and Systems for Video Technology*, vol. 24, no. 5, pp. 814–825, 2014.
- [12] M. Antal and L. Z. Szabó, "An evaluation of one-class and two-class classification algorithms for keystroke dynamics authentication on mobile devices," in *Control Systems and Computer Science (CSCS), 2015 20th International Conference on*. IEEE, 2015, pp. 343–350.
- [13] V. M. Patel, R. Chellappa, D. Chandra, and B. Barbelo, "Continuous user authentication on mobile devices: Recent progress and remaining challenges," *IEEE Signal Processing Magazine*, vol. 33, no. 4, pp. 49–61, July 2016.
- [14] P. Perera and V. M. Patel, "Efficient and low latency detection of intruders in mobile active authentication," *IEEE Transactions on Information Forensics and Security*, vol. 13, no. 6, pp. 1392–1405, 2018.
- [15] B. Schölkopf, R. C. Williamson, A. J. Smola, J. Shawe-Taylor, and J. C. Platt, "Support vector method for novelty detection," in *Advances in neural information processing systems*, 2000, pp. 582–588.
- [16] S. M. Erfani, S. Rajasegarar, S. Karunasekera, and C. Leckie, "High-dimensional and large-scale anomaly detection using a linear one-class svm with deep learning," *Pattern Recognition*, vol. 58, pp. 121–134, 2016.
- [17] 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.
- [18] D. M. Tax and R. P. Duin, "Support vector data description," *Machine learning*, vol. 54, no. 1, pp. 45–66, 2004.
- [19] S. S. Khan and M. G. Madden, "One-class classification: taxonomy of study and review of techniques," *The Knowledge Engineering Review*, vol. 29, no. 3, pp. 345–374, 2014.
- [20] G. Lanckriet, L. E. Ghaoui, C. Bhattacharyya, and M. I. Jordan, "Minimax probability machine," in *Advances in neural information processing systems*, 2002, pp. 801–807.
- [21] L. E. Ghaoui, M. I. Jordan, and G. R. Lanckriet, "Robust novelty detection with single-class mpm," in *Advances in neural information processing systems*, 2003, pp. 929–936.
- [22] P. Perera and V. M. Patel, "Dual-minimax probability machines for one-class mobile active authentication," in *Biometrics Theory, Applications and Systems (BTAS), 2016 IEEE 8th International Conference on*. IEEE, 2018.
- [23] J. Deng, W. Dong, R. Socher, L.-J. Li, K. Li, and L. Fei-Fei, "Imagenet: A large-scale hierarchical image database," in *Computer Vision and Pattern Recognition, 2009. CVPR 2009. IEEE Conference on*. Ieee, 2009, pp. 248–255.
- [24] M. Sabokrou, M. Fayyaz, M. Fathy, Z. Moayed, and R. Klette, "Deep-anomaly: Fully convolutional neural network for fast anomaly detection in crowded scenes," *Computer Vision and Image Understanding*, 2018.
- [25] R. Chalapathy, A. K. Menon, and S. Chawla, "Anomaly detection using one-class neural networks," *arXiv preprint arXiv:1802.06360*, 2018.
- [26] W. Lawson, E. Bekele, and K. Sullivan, "Finding anomalies with generative adversarial networks for a patrolbot," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition Workshops*, 2017, pp. 12–13.
- [27] M. Ravanbakhsh, M. Nabi, E. Sangineto, L. Marcenaro, C. Regazzoni, and N. Sebe, "Abnormal event detection in videos using generative adversarial nets," in *Image Processing (ICIP), 2017 IEEE International Conference on*. IEEE, 2017, pp. 1577–1581.
- [28] C. Zhou and R. C. Paffenroth, "Anomaly detection with robust deep autoencoders," in *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM, 2017, pp. 665–674.
- [29] R. Chalapathy, A. K. Menon, and S. Chawla, "Robust, deep and inductive anomaly detection," in *Joint European Conference on Machine Learning and Knowledge Discovery in Databases*. Springer, 2017, pp. 36–51.
- [30] P. Perera and V. M. Patel, "Learning deep features for one-class classification," *arXiv preprint arXiv:1801.05365*, 2018.
- [31] 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*, 2014, pp. 2672–2680.
- [32] T. Salimans, I. Goodfellow, W. Zaremba, V. Cheung, A. Radford, X. Chen, and X. Chen, "Improved techniques for training gans," in *Advances in Neural Information Processing Systems 29*, D. D. Lee, M. Sugiyama, U. V. Luxburg, I. Guyon, and R. Garnett, Eds. Curran Associates, Inc., 2016, pp. 2234–2242.
- [33] A. Krizhevsky, I. Sutskever, and G. E. Hinton, "Imagenet classification with deep convolutional neural networks," in *Advances in neural information processing systems*, 2012, pp. 1097–1105.
- [34] K. Simonyan and A. Zisserman, "Very deep convolutional networks for large-scale image recognition," 2015.
- [35] D. P. Kingma and J. Ba, "Adam: A method for stochastic optimization," *arXiv preprint arXiv:1412.6980*, 2014.
- [36] V. Dumoulin, J. Shlens, and M. Kudlur, "A learned representation for artistic style," *Proc. of ICLR*, 2017.
- [37] B. Saleh, A. Farhadi, and A. Elgammal, "Object-centric anomaly detection by attribute-based reasoning," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2013, pp. 787–794.
- [38] U. Mahbub, S. Sarkar, V. M. Patel, and R. Chellappa, "Active user authentication for smartphones: A challenge data set and benchmark results," in *Biometrics Theory, Applications and Systems (BTAS), 2016 IEEE 8th International Conference on*. IEEE, 2016, pp. 1–8.
- [39] J. Liu, Z. Lian, Y. Wang, and J. Xiao, "Incremental kernel null space discriminant analysis for novelty detection," in *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, 2017, pp. 792–800.
- [40] M. Everingham, L. Van Gool, C. K. Williams, J. Winn, and A. Zisserman, "The pascal visual object classes (voc) challenge," *International journal of computer vision*, vol. 88, no. 2, pp. 303–338, 2010.