

Learning Class-Conditional GANs with Active Sampling

Ming-Kun Xie and Sheng-Jun Huang

College of Computer Science and Technology, Nanjing University of Aeronautics and Astronautics

MIT Key Laboratory of Pattern Analysis and Machine Intelligence

Collaborative Innovation Center of Novel Software Technology and Industrialization, Nanjing, 211106

{mkxie,huangs}@nuaa.edu.cn

ABSTRACT

Class-conditional variants of Generative adversarial networks (GANs) have recently achieved a great success due to its ability of selectively generating samples for given classes, as well as improving generation quality. However, its training requires a large set of class-labeled data, which is often expensive and difficult to collect in practice. In this paper, we propose an active sampling method to reduce the labeling cost for effectively training the class-conditional GANs. On one hand, the most useful examples are selected for external human labeling to jointly reduce the difficulty of model learning and alleviate the missing of adversarial training; on the other hand, fake examples are actively sampled for internal model retraining to enhance the adversarial training between the discriminator and generator. By incorporating the two strategies into a unified framework, we provide a cost-effective approach to train class-conditional GANs, which achieves higher generation quality with less training examples. Experiments on multiple datasets, diverse GAN configurations and various metrics demonstrate the effectiveness of our approaches.

CCS CONCEPTS

• **Computing methodologies** → **Active learning settings; Neural networks.**

KEYWORDS

Active sampling, generative adversarial networks, class-conditional GANs

ACM Reference Format:

Ming-Kun Xie and Sheng-Jun Huang. 2019. Learning Class-Conditional GANs with Active Sampling. In *Proceedings of 25th ACM SIGKDD Conference on Knowledge Discovery and Data Mining (KDD'19)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/1122445.1122456>

1 INTRODUCTION

In machine learning, generative modeling has been extensively studied to generate samples indistinguishable from real data. Recently, the emergence of deep generative models offers a powerful

framework for tackling the problem. Among them, generative adversarial networks (GANs) [5] as a new way of learning generative model, have been proved to be a promising one owing to their excellent performances. Recent works have shown GANs can produce convincing results in various challenging tasks, such as realistic image generation [24, 38], conditional image generation [12, 13] and text generation [37].

Along with this success, many variants of GANs have emerged over the past few years. For example, many models are proposed to improve the quality of generated images [3, 14], and some other methods try to stabilize the training procedure [1, 2, 20]. Class-conditional GANs is another branch of research which attracted many research interests. They try to generate images of a given class by utilizing additional information, such as the class label [22, 25]. Among them, auxiliary classifier GAN(AC-GAN) [25] is a representative approach that jointly trains the real-fake discriminator and an auxiliary classifier for predicting the specific class label. It can selectively generate images for given classes, and usually leads to both higher generation quality and stability.

Despite the great successes AC-GANs have achieved, a potential limitation is that the model training requires a large set of labeled examples. The performance of AC-GANs seriously depends on the size of labeled training data. As shown in Figure 1, when a AC-GAN is applied to conditional image generation tasks, the quality of generated images is significantly improved as the number of labeled training data increases. The examples in the figure are generated by AC-GAN trained on different numbers of training data from CIFAR10 [16] (the training size varies from 1k, 5k, 20k to 50k). In real-world scenarios, it is expensive and difficult to acquire a large number of labeled data. Therefore, training AC-GANs with lower labeling cost is an important issue with great significance.

Active learning is a primary approach for reducing the labeling cost [7, 11, 31]. It progressively selects the most useful samples and queries their labels, with the target of training an effective model with less queries. The selection strategy thus plays an important role in active learning. One of the most common strategies is the uncertainty-based selection [18, 32, 36], which measures the uncertainty of unlabeled samples from the predictions of previous classifiers. There are some recent studies trying to combine informativeness and representativeness, which estimate the potential contribution of an example on improving the classifier [8, 10, 11, 35]. However, these strategies are designed for traditional classification tasks.

In this paper, we propose an Active Learning approach for Class-conditional GANs (ALCG), which try to produce high-quality images for given classes with low annotation cost under the AC-GAN [25] framework. Different from existing active learning methods, our ALCG framework proposes to automatically select samples

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

KDD'19, August 4 - 8, Anchorage, Alaska

© 2019 Copyright held by the owner/author(s). Publication rights licensed to ACM.

ACM ISBN 978-1-4503-9999-9/18/06...\$15.00

<https://doi.org/10.1145/1122445.1122456>



Figure 1: Generated image samples on CIFAR10 with various numbers of training data.

from training data as well as fake data during the two different stages, i.e., the external human labeling stage and internal model retraining stage, respectively. In the external human labeling stage, we propose a certainty sampling strategy to actively select examples from unlabeled data and query their labels for model training. Specifically, the most certain examples with high prediction confidence are selected. On one hand, these examples are usually clean images and are relatively easier for AC-GANs training; on the other hand, the most certain samples contribute less to the classification loss so as to alleviate the unstable training of AC-GANs, which is caused by the missing of adversarial training in the auxiliary classifier [39].

During the internal model retraining stage, regular GANs typically update the model by generating a batch of samples from random noise, which may lead to an inaccurate discriminator. We propose an adversarial sampling strategy to actively select the most discriminative examples for updating both discriminator and generator. The adversarial sampling (AS) strategy leads to a more intensive adversarial training between discriminator and generator, and thus can make the model converge faster and often produce better results.

Our main contributions are summarized as follows:

- A general framework of active learning for class-conditional GANs is proposed. It actively selects examples from both training data and fake data to produce high-quality images with lower annotation cost.
- Two novel sampling strategies are proposed. In the external human labeling stage, certainty sampling is proposed to lower the difficulty of the model learning. In the internal model retraining stage, the adversarial sampling is proposed to enhance the adversarial training between discriminator and generator.
- Experiments on multiple benchmark datasets, diverse GAN configurations and various performance metrics demonstrate that the proposed method can achieve higher generation quality with less training data.

The rest of the paper is organized as follows: In section 2, we briefly review related works. In section 3, the proposed approach is introduced. In section 4, experimental results are reported, followed by the conclusion in section 5.

2 RELATED WORK

The goal of generative modeling is to learn the true generative distribution of training data and then to generate or reproduce the new data point from the same distribution. Recently, the emergence

of deep generative models offers a powerful framework for this task. Among them, two of the most popular and efficient approaches are GANs [5] and variational autoencoders (VAEs) [15, 29], in which VAEs are proposed to maximize the lower bound of the data log-likelihood and GANs aim at achieving a Nash equilibrium between discriminator and generator. Each of these two frameworks has its own merits. VAEs use deterministic approximation to maximize likelihood so as to avoid intractable density functions [5], whereas GANs learn a generative distribution through adversarial training without explicit density estimation [15].

Recently, researchers are interested in conditional variants of GAN because they can selectively generate samples for given classes or improve generation quality and stability as mentioned in section 1. In addition to class labels [22, 25, 39], supervised information including object locations [27], texts [28, 38], images [13, 17] and videos [34] are used for conditional variants of GANs, and they have achieved great successes in various tasks. As a representative of class-conditional GANs, AC-GAN [25] incorporates class label information by introducing an auxiliary classifier for classes in the original GAN framework. It can selectively generate images for given classes, and always leads to both higher generation quality and stability. The proposed approaches are implemented under the AC-GAN framework.

Active learning has been actively studied for reducing the annotation cost [10, 31]. The basic approach of the active learning methods is to progressively select and annotate most useful unlabeled samples to boost the model. The key element of active learning is the selection criterion, and it is typically designed according to the classification uncertainty of samples [18, 32]. Recently, researchers have also combined active learning and deep learning to propose more effective methods. Some works apply active learning methods in deep image classification tasks to reduce the annotation cost of large scale data [9, 33]. There are a few methods using deep generative models, such as GANs, to generate or acquire more effective queries [21, 40]. While these methods are utilizing GANs to help active learning, in this paper, we propose to employ active learning techniques to improve the GANs.

3 THE PROPOSED APPROACH

In following content, we use superscript r and g to denote the real distribution and generative distribution, respectively. Let $\mathbf{x} \in \mathcal{X}$ denote the training sample, and $\mathbf{y} \in \mathcal{Y}$ denote the corresponding class label. Here, $\mathcal{X} \subseteq \mathbb{R}^d$ is feature space and $\mathcal{Y} = \{1, 2, \dots, m\}$ is label space, where m is the number of classes. We denote $L = \{\mathbf{x}_i\}_{i=1}^{n_l}$ the labeled data set with n_l examples, and $U = \{\mathbf{x}_i\}_{i=1}^{n_u}$ the unlabeled data set with n_u examples. We will first introduce the

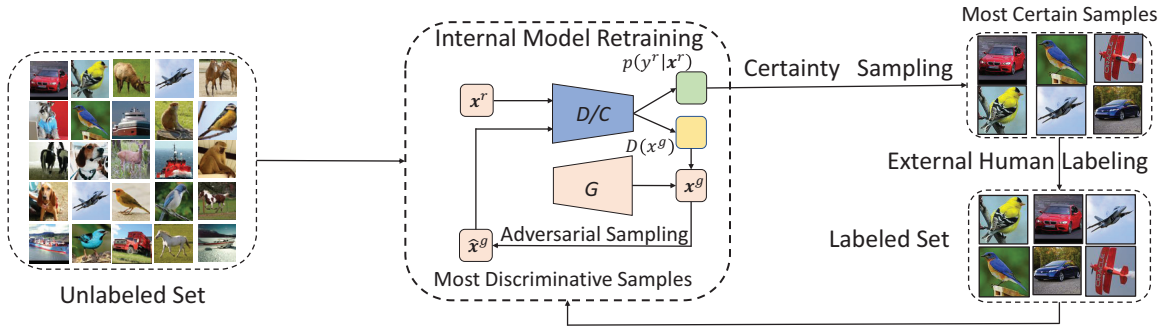


Figure 2: The framework of active learning for class-conditional GANs.

basic idea of AC-GAN, and then propose our ALCG framework. Finally, we describe the detail of certainty sampling and adversarial sampling progressively.

3.1 AC-GAN

In our proposed approach, we employ AC-GAN [25] as the base model. AC-GAN is one of the most popular conditional variants of GANs [5]. In the AC-GAN, the conditional generator is composed of two neural networks, in which one is the discriminator D and the other is the auxiliary classifier C . Discriminator D assigns probability $p = D(x)$ for examples $x \sim p^r(x)$ or assigns probability $1 - p$ for examples $x \sim p^g(x)$. The auxiliary classifier $C(y|x)$ gives a probability distribution over class labels given x . The objective function has two parts: an adversarial loss and an auxiliary classifier loss.

The adversarial loss is defined as:

$$\mathcal{L}_{AC} = \mathbb{E}_{x^r \sim p^r(x)} [\log D(x^r)] + \mathbb{E}_{z \sim p(z), y^g \sim p(y)} [\log (1 - D(G(z, y^g)))] \quad (1)$$

where D attempts to find the best decision boundary between real and generated data by maximizing this loss, and G attempts to generate data indistinguishable by D by minimizing this loss.

The auxiliary classifier loss is used to prompt the generator to generate samples for given classes. To achieve this, C is first optimized using classification loss of real data:

$$\mathcal{L}_{AC}^r = \mathbb{E}_{(x^r, y^r) \sim p^r(x, y)} [-\log C(y = y^r | x^r)], \quad (2)$$

where C learns to assign correct class labels to the real data by minimizing this loss. Furthermore, G is optimized by using a classification loss of generated data:

$$\mathcal{L}_{AC}^g = \mathbb{E}_{z \sim p(z), y^g \sim p(y)} [-\log C(y = y^g | G(z, y^g))] \quad (3)$$

where G attempts to generate data for given classes by minimizing this loss.

Overall, the full objective function is defined as:

$$\mathcal{L}_{D/C} = \mathcal{L}_{GAN} - \mathcal{L}_{AC}^r \quad (4)$$

$$\mathcal{L}_G = \mathcal{L}_{GAN} + \mathcal{L}_{AC}^g \quad (5)$$

D/C is trained by maximizing $\mathcal{L}_{D/C}$ while G is trained by minimizing \mathcal{L}_G .

Algorithm 1 The CS_MC algorithm

- 1: **Input:**
- 2: U : the unlabeled set
- 3: L : the initially labeled set
- 4: **Process:**
- 5: Initialize the parameters of the AC-GAN with labeled set L .
- 6: **For:** $t = 1 : T$
- 7: Use certainty sampling introduced in subsection 3.3 to select a batch of samples Q from U and query their labels
- 8: Add Q to L , and remove Q from U .
- 9: Feed labeled set L into AC-GAN for retraining.
- 10: **For** each training iteration
- 11: Sample a minibatch of real samples x^r from L .
- 12: Generate a set of fake data x^g by generator from random sampling noise.
- 13: Use adversarial sampling introduced in subsection 3.4 to select a minibatch of samples \hat{x}^g from fake data x^g .
- 14: Use x^r and \hat{x}^g to update D/C by maximizing Eq.(4).
- 15: Use \hat{x}^g to update G by minimizing Eq.(5).
- 16: **End For**
- 17: **End For**

3.2 The ALCG framework

As Figure 2 illustrates, our ALCG framework actively selects samples from two kinds of source, i.e., training data from unlabeled set and fake data generated by generator G during two different stages, i.e., external human labeling stage and internal model retraining stage. During the external human labeling stage, our ALCG framework progressively feeds the samples from unlabeled set U into the AC-GAN, and then certainty sampling is proposed to estimate the certainty of samples based on predictions from the auxiliary classifier C . The most certain ones are selected to add into labeled set after annotator labeling and are fed into the model for retraining. During the internal model retraining stage, adversarial sampling is used to select the most discriminative samples \hat{x}^g from fake data x^g based on the outputs of discriminator D . These samples are used to update both discriminator and generator. The entire algorithm can be summarized in Algorithm 1. In the next subsections, we will respectively introduce the detail of certainty sampling and adversarial sampling.



Figure 3: Visualization of images with low/high certainty.

3.3 Certainty Sampling

In this subsection, we focus on the external human labeling, in which the most useful samples are selected to query their labels and add them into labeled set L . In contrast to classical active learning that considers the most uncertain samples, e.g., samples of low classification confidence, we instead select the most certain samples, i.e., samples of high classification confidence for AC-GAN training. We expect the certain examples can reduce the difficulty of the model learning and alleviate the unstable training of AC-GAN. Firstly, the most certain examples tend to be clean images and are relatively easy for AC-GANs training. Figure 3 shows the samples of low confidence (in the first row) and the samples of high confidence (in the second row), in which the confidences are estimated based on the predictions of the auxiliary classifier in AC-GAN. Compared to low confidence examples, the certain ones tend to be clearer or contain cleaner backgrounds and more explicit semantic objectives. Such examples allow AC-GANs to be more likely to generate high-quality samples through adversarial training. Secondly, the most certain examples contribute to a less auxiliary classifier loss so as to alleviate the unstable training of AC-GANs which is caused by the missing of adversarial training in the auxiliary classifier [39]. This helps the model avoid suffering from mode collapse and produce high-quality image samples.

Next, we propose three certainty sampling criteria to estimate the certainty of an example, and select most certain ones for querying.

In traditional active learning, the uncertainty of an example is defined by the confidence of the classifier prediction, where a lower confidence indicates a larger uncertainty. Similarly, we can define the certainty in the same way. Specifically we can define the certainty based on $p(y_i^r | \mathbf{x}_i^r)$ predicted by the auxiliary classifier which denotes the probability of \mathbf{x}_i belonging to j th class. Based on classification probability, the certainty can be defined by entropy and margin which are commonly used for designing selection criteria. The three certainty sampling criteria are defined as following. Note that we only focus on real examples in this subsection, therefore, the superscript r is omitted for convenience.

- *Most confidence (MC)*: Rank all the unlabeled samples in a descending order according to the mc_i value, which is defined as:

$$mc_i = \max_j p(y_i = j | x_i), \quad (6)$$

If the probability of the most probable class is high then the auxiliary classifier is certain about the sample.

- *Large margin (LM)*: Rank all the unlabeled samples in a descending order according to the lm_i value, which is defined as:

$$lm_i = p(y_i = j_1 | x_i) - p(y_i = j_2 | x_i) \quad (7)$$

where j_1 and j_2 represent the first and the second most probable class labels predicted by the auxiliary classifier. The larger of the margin means the auxiliary classifier is more certain about the sample.

- *Least entropy (LE)*: Rank all the unlabeled samples in an ascending order according to the le_i value, which is defined as:

$$le_i = - \sum_{j=1}^m p(y_i = j | x_i) \log p(y_i = j | x_i) \quad (8)$$

The lower entropy value means the auxiliary classifier is more certain about the sample.

At every iteration during the external human labeling stage, we first compute certainty for each example over the unlabeled set U according to each of three criteria and then select a small batch of the most certain samples to query their ground-truth labels. We name the algorithm of Certainty Sampling with Most Confidence criterion as CS_MC for short. Similarly, we have CS_LM for the large margin criterion, and CS_LE for the least entropy criterion.

3.4 Adversarial Sampling

In this subsection, we focus on the internal model retraining, in which we retrain a AC-GAN over the labeled set L . Regular GANs often generate a batch of samples from random noise and use these samples to update the model. We argue that such samples will lead to an inaccurate discriminator, as illustrated in Figure 4. Figure 4 (a) shows the distribution of real examples and potential fake examples, while the blue line represents an initial decision boundary for discriminating real and fake data. Figure 4 (b) shows that when we use the randomly generated fake samples (in cyan) to update the discriminator, the decision boundary will move toward the real samples. However, the discriminator is still not curate enough to correctly identify all fake samples (it makes mistakes for fake samples on the side closed to real samples of the decision boundary). Instead we find that more discriminative samples often make the decision boundary more accurate. As shown in Figure 4 (c), when the discriminator is updated by using the most discriminative samples (in red), it is able to identify all the fake samples, and then it will provide the generator with more accurate gradient guidance for updating and make it produce better results. However, the most discriminative samples are unavailable because fake data are unknown before sampled. Thus we propose an intuitive method to acquire samples as discriminative as possible.

Specifically, at each iteration of model retraining, the generator first generates a set of candidate examples from random sampling noise, and then the most discriminative ones are selected from the set to update both discriminator and generator. We set the size of the candidate set as γ times the size of the final selected batch. Obviously we are likely to acquire more discriminative samples with a larger γ so as to obtain a more powerful discriminator. However, an excessively powerful discriminator sometimes makes generator stop updating. We thus set it as a default value, such as 4 in the




Figure 4: Illustration of random sampling and adversarial sampling.

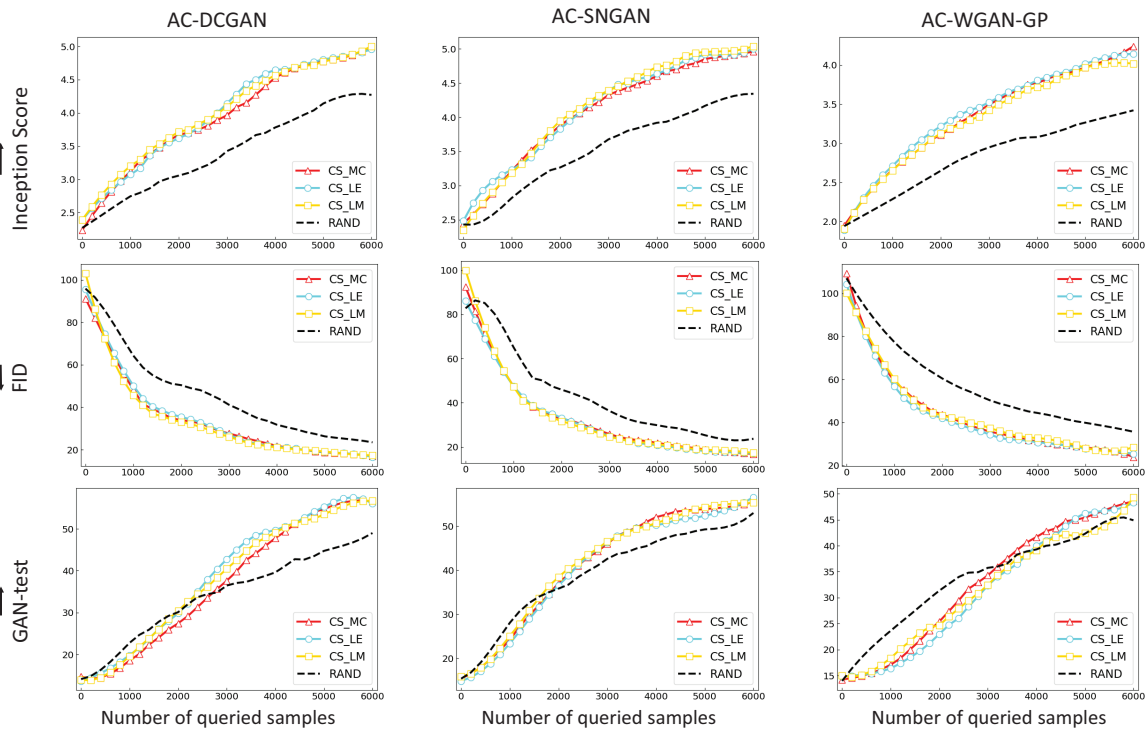


Figure 5: The comparison of certainty sampling on CIFAR10 with different GAN configurations.

experiments to avoid such cases. The selection criterion is based on the output of the discriminator $D(x_i)$, which denotes the probability of fake sample x_i belonging to real data. We first rank all fake samples in a descending order according to the value of $D(x)$, and then select a batch of the largest ones for updating discriminator and generator. This method is named by Adversarial Sampling (AS) because the most discriminative fake samples lead to a more intensive adversarial training between discriminator and generator so as to make the model converge faster and often produce better results.

4 EXPERIMENTS

4.1 Settings

We perform experiments on three benchmark dataset: CIFAR10 [16], STL10 [4] and TinyImagenet [30], which are commonly used in both image generation and active learning tasks. CIFAR10 contains 60k 32×32 natural images, which are divided into 50k training and 10k test images. STL10 contains 5k training and 8k test images. TinyImagenet contains 200 classes with 500 images for each class and we divide 500 images into 400 training and 100 test images. Both CIFAR10 and STL10 have 10 classes. For TinyImagenet, we

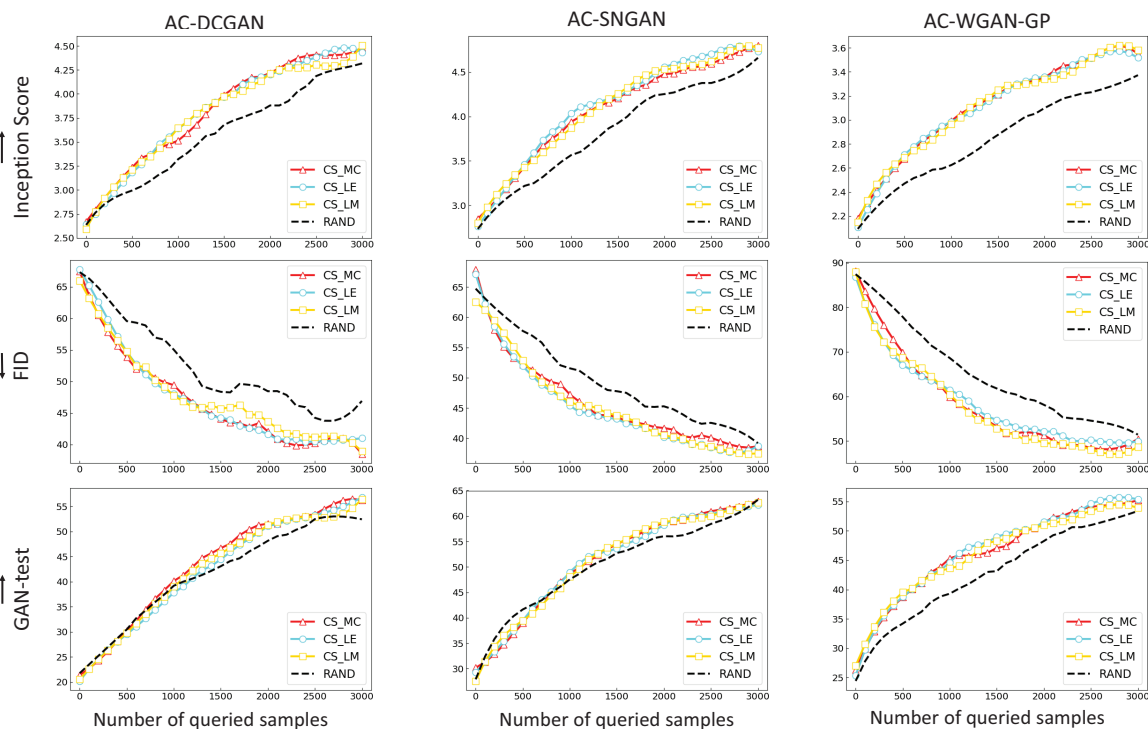


Figure 6: The comparison of certainty sampling on STL10 with different GAN configurations.

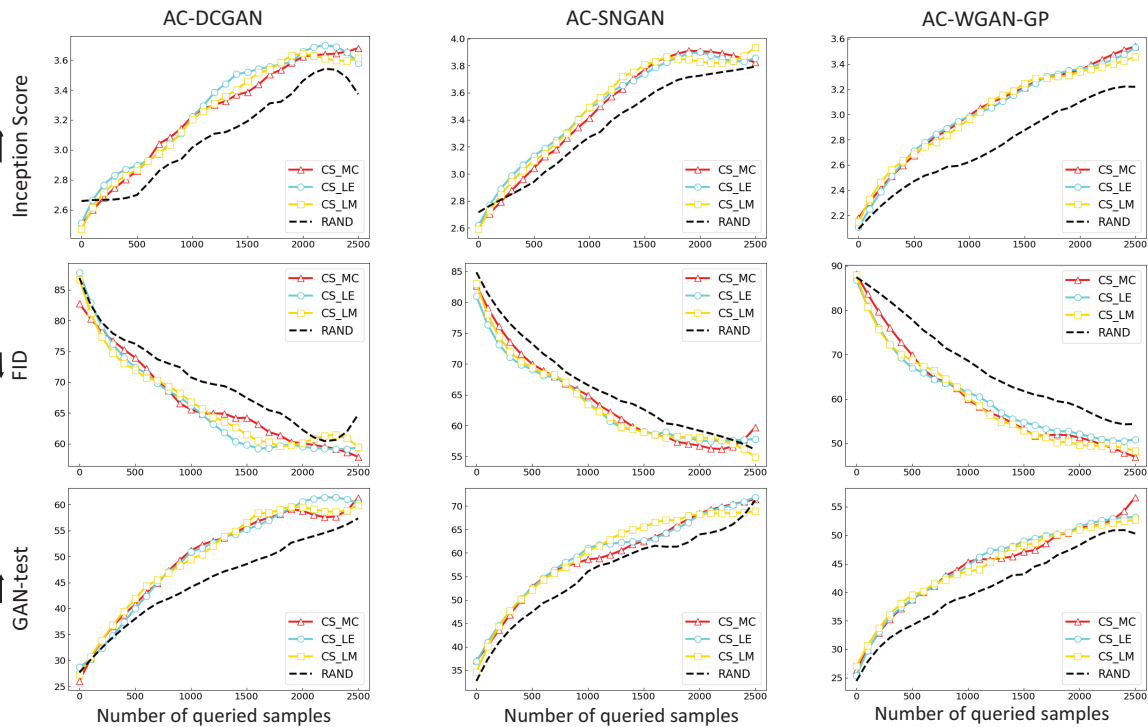


Figure 7: The comparison of certainty sampling on TinyImagenet10 with different GAN configurations.

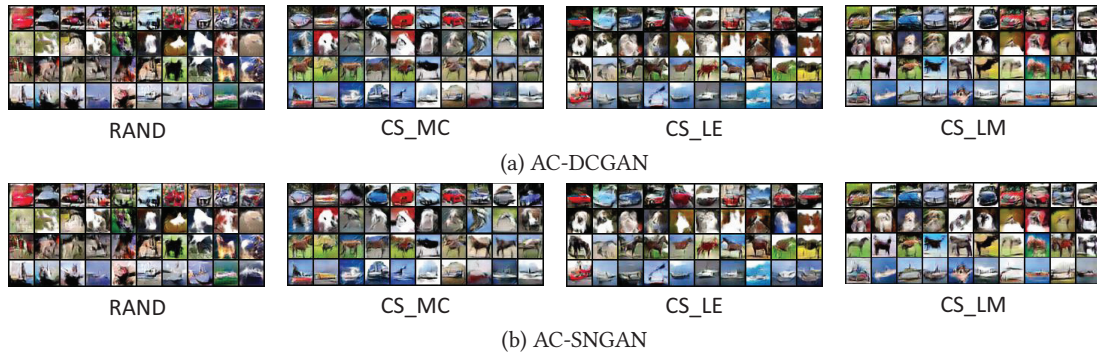


Figure 8: Generated image samples on CIFAR10 with different sampling strategy




Figure 9: The comparison of adversarial sampling on CIFAR10 with different GAN configurations.

randomly sampled a subset with 10 classes and name it TinyImagenet10. For STL10 and TinyImagenet10, we also resize the original image size to 32×32 . Among each of datasets, we randomly select 1k samples from training data to initialize the networks and the rest are left as unlabeled dataset. To avoid the influence of randomness, we repeat the experiments for 3 times and report the average results.

We examine our methods on following three popular models: DCGAN [26], SNGAN [23] and WGAN-GP [6]. We implement AC-GAN based on these three models, and we use D/C in which the layers are shared except for the last layer. We also apply the standard DCGAN architecture for each model. As shown in [19], GANs are sensitive to parameters. Actually, it is impractical to tune the parameters in active learning. Therefore, instead of searching optimal parameters for each case, we use default parameters that are typically used in the settings without active sampling for all methods.

In the experiments, we examine the performance on three metrics for a comprehensive analysis: (1) Inception Score (IS), (2) Fréchet Inception distance (FID), and (3) GAN-test. Inception score calculates the expectation of KL-divergence between the conditional class distribution and the marginal class distribution to measure both quality and diversity of generated images. FID measures the distance between the distribution of real images and generated images based on Inception embeddings. We use it to evaluate the quality of generative distribution. The GAN-test is the accuracy of a classifier trained on real images and tested on generated images.

The metric reflects the precision (i.e., image quality) of GANs and quantifies how close generated images are to a data manifold.

In the following, we firstly report the results to examine the effectiveness of certainty sampling and adversarial sampling, respectively, and then report the results of the ALCG method.

4.2 Performances of Certainty Sampling

In this subsection, we examine the performances of certainty sampling independently, in which a regular AC-GAN model is trained without adversarial sampling during the internal model retraining stage. Note that there is no existing approach applicable to our setting, we compared the proposed certainty sampling criteria most confidence, least entropy and large margin (which are denoted by CS_MC, CS_LE and CS_LM, respectively) with random sampling (which is denoted by RAND). Figure 5, 6 and 7 show the results on CIFAR10, STL10 and TinyImagenet10, respectively. The results show that our method outperforms the baseline method from the aspects of the quality of generated images, the overall generative distribution and conditional generative distribution which are reflected on Inception score, FID and GAN-test, respectively. Regarding three metrics, our certainty sampling is significantly better than random sampling, on all GAN configurations and datasets. While comparing among the three certainty sampling criteria, they are comparable to each other in most cases.

We also show some example images generated on CIFAR10 in Figure 8. For each sampling method, the model trained on all labeled data is used to generate image samples. Each row shows samples

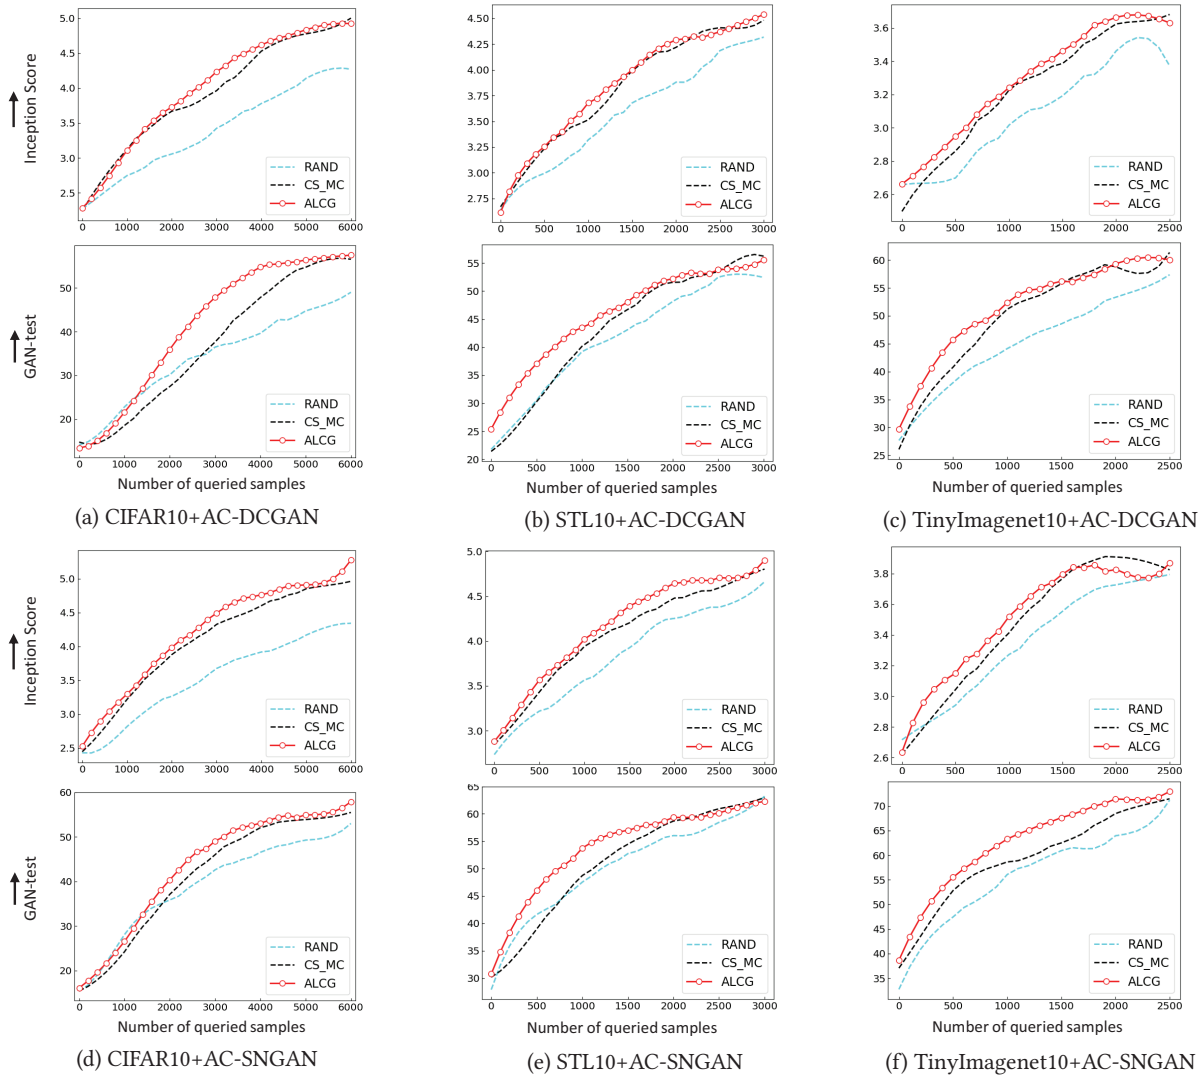


Figure 10: The comparison of ALCG on different datasets with different GAN configurations

belonging to the same class. Obviously, for both AC-DCGAN and AC-SNGAN, samples generated with certainty sampling acquire much higher fidelity as compared to random sampling.

4.3 Performances of Adversarial Sampling

In this subsection, we examine the performances of adversarial sampling. Due to page limit, we only report the results of Inception score and GAN-test in this and next subsections, where FID is omitted since it is also a Inception based metric and often yields similar result to Inception score. To validate the effectiveness of adversarial sampling independently, Figure 9 plots performance curves of AC-DCGAN and AC-SNGAN with or without adversarial sampling (which are denoted by AS and Regular, respectively) during the training process. It can be observed that the result of model with AS is significantly better than result of model without

AS, on all GAN configurations and metrics, which validates that adversarial sampling is effective for AC-GANs training.

4.4 Performances of the ALCG method

Lastly, we examine the performances of the proposed method ALCG. Results in Subsection 4.2 have shown that the performance of the three certainty sampling criteria are very similar. To save the space, we thus only implement the method with the most confidence criterion. And we compare ALCG with methods CS_MC and RAND. The performance curves are plotted in Figure 10. It can be observed that with different GAN configurations and different datasets, ALCG achieves the best result in most cases. This demonstrates that our proposed ALCG framework can improve generation quality with less training data.

5 CONCLUSION

In this paper, we perform active learning for cost-effective training of class-conditional GANs. Examples are actively selected for both external human labeling and internal model retraining. During the external human labeling stage, certainty sampling is proposed to reduce the difficulty of model learning and alleviate the missing of adversarial training in AC-GANs. During the internal model retraining stage, adversarial sampling is proposed to enhance the adversarial training between the generator and the discriminator. Experiments are performed on different datasets with various GAN configurations. The results show that the proposed approaches can achieve high-quality conditional image generation with significantly lower cost. In the future, we plan to apply our approaches on higher resolution image datasets. Also the diversity information will be considered in the active sampling to further improve the performance.

6 ACKNOWLEDGMENTS

This research was supported by NSFC(61876081, 61732006) and the Fundamental Research Funds for the Central Universities, NO. NE2019104.

REFERENCES

- [1] Martin Arjovsky, Soumith Chintala, and Léon Bottou. 2017. Wasserstein generative adversarial networks. In *International Conference on Machine Learning*. 214–223.
- [2] David Berthelot, Thomas Schumm, and Luke Metz. 2017. BEGAN: boundary equilibrium generative adversarial networks. *arXiv preprint arXiv:1703.10717* (2017).
- [3] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. 2016. Infogan: Interpretable representation learning by information maximizing generative adversarial nets. In *Advances in neural information processing systems*. 2172–2180.
- [4] Adam Coates, Andrew Y. Ng, and Honglak Lee. 2011. An Analysis of Single-Layer Networks in Unsupervised Feature Learning. In *Proceedings of the Fourteenth International Conference on Artificial Intelligence and Statistics*. 215–223.
- [5] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. 2014. Generative adversarial nets. In *Advances in neural information processing systems*. 2672–2680.
- [6] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron C. Courville. 2017. Improved Training of Wasserstein GANs. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems*. 5769–5779.
- [7] Sheng-Jun Huang and Songcan Chen. 2016. Transfer Learning with Active Queries from Source Domain. In *Proceedings of the Twenty-Fifth International Joint Conference on Artificial Intelligence, IJCAI 2016*. 1592–1598.
- [8] Sheng-Jun Huang, Nengneng Gao, and Songcan Chen. 2017. Multi-instance multi-label active learning. In *Proceedings of the Twenty-Sixth International Joint Conference on Artificial Intelligence, IJCAI 2017*. 1886–1892.
- [9] Sheng-Jun Huang, Jia-Wei Zhao, and Zhao-Yang Liu. 2018. Cost-Effective Training of Deep CNNs with Active Model Adaptation. In *Proceedings of the 24th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. 1580–1588.
- [10] Sheng-Jun Huang and Zhi-Hua Zhou. 2013. Active Query Driven by Uncertainty and Diversity for Incremental Multi-label Learning. In *2013 IEEE 13th International Conference on Data Mining*. 1079–1084.
- [11] Sheng-Jun Huang, Rong Jin, and Zhi-Hua Zhou. 2014. Active Learning by Querying Informative and Representative Examples. *IEEE Transactions on Pattern Analysis and Machine Intelligence* 10 (2014), 1936–1949.
- [12] Xun Huang, Yixuan Li, Omid Poursaeed, John Hopcroft, and Serge Belongie. 2017. Stacked generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition*, Vol. 2.
- [13] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. 2017. Image-to-Image Translation with Conditional Adversarial Networks. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*. IEEE, 5967–5976.
- [14] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen. 2018. Progressive Growing of GANs for Improved Quality, Stability, and Variation. In *ICLR*.
- [15] Diederik P. Kingma and Max Welling. [n. d.]. Auto-Encoding Variational Bayes. In *ICLR*, Vol. abs/1312.6114.
- [16] Alex Krizhevsky. 2015. Learning Multiple Layers of Features from Tiny Images. *Technique report* (2015).
- [17] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew P. Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. 2017. Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. In *2017 IEEE Conference on Computer Vision and Pattern Recognition*. 105–114.
- [18] David D. Lewis. 1995. A Sequential Algorithm for Training Text Classifiers: Corrigendum and Additional Data. *SIGIR Forum* 29, 2 (1995), 13–19.
- [19] Mario Lucic, Karol Kurach, Marcin Michalski, Sylvain Gelly, and Olivier Bousquet. 2018. Are GANs Created Equal? A Large-Scale Study. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems*. 698–707.
- [20] Xudong Mao, Qing Li, Haoran Xie, Raymond YK Lau, Zhen Wang, and Stephen Paul Smolley. 2017. Least squares generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*. 2794–2802.
- [21] Christoph Mayer and Radu Timofte. 2018. Adversarial Sampling for Active Learning. *arXiv preprint arXiv:1808.06671* abs/1808.06671 (2018).
- [22] Mehdi Mirza and Simon Osindero. 2014. Conditional generative adversarial nets. *arXiv preprint arXiv:1411.1784* (2014).
- [23] Takeru Miyato, Toshiaki Kataoka, Masanori Koyama, and Yuichi Yoshida. 2018. Spectral Normalization for Generative Adversarial Networks. *arXiv preprint arXiv:1802.05957* abs/1802.05957 (2018).
- [24] Anh Nguyen, Jeff Clune, Yoshua Bengio, Alexey Dosovitskiy, and Jason Yosinski. 2017. Plug & play generative networks: Conditional iterative generation of images in latent space. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*. 4467–4477.
- [25] Augustus Odena, Christopher Olah, and Jonathon Shlens. 2017. Conditional image synthesis with auxiliary classifier gans. In *Proceedings of the 34th International Conference on Machine Learning*. 2642–2651.
- [26] Alec Radford, Luke Metz, and Soumith Chintala. 2015. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. *arXiv preprint arXiv:1511.06434* abs/1511.06434 (2015).
- [27] Scott E. Reed, Zeynep Akata, Santosh Mohan, Samuel Tenka, Bernt Schiele, and Honglak Lee. 2016. Learning What and Where to Draw. In *Advances in Neural Information Processing Systems 29: Annual Conference on Neural Information Processing Systems*. 217–225.
- [28] Scott E. Reed, Zeynep Akata, Xinchen Yan, Lajanugen Logeswaran, Bernt Schiele, and Honglak Lee. 2016. Generative Adversarial Text to Image Synthesis. In *Proceedings of the 33rd International Conference on Machine Learning*. 1060–1069.
- [29] Danilo Jimenez Rezende, Shakir Mohamed, and Daan Wierstra. 2014. Stochastic Backpropagation and Approximate Inference in Deep Generative Models. In *Proceedings of the 31th International Conference on Machine Learning*. 1278–1286.
- [30] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael S. Bernstein, Alexander C. Berg, and Fei-Fei Li. 2015. ImageNet Large Scale Visual Recognition Challenge. *International Journal of Computer Vision* 115, 3 (2015), 211–252.
- [31] Burr Settles. 2012. Active learning. *Synthesis Lectures on Artificial Intelligence and Machine Learning* 6, 1 (2012), 1–114.
- [32] Simon Tong and Daphne Koller. 2001. Support Vector Machine Active Learning with Applications to Text Classification. *Journal of Machine Learning Research* 2 (2001), 45–66.
- [33] Keze Wang, Dongyu Zhang, Ya Li, Ruimao Zhang, and Liang Lin. 2017. Cost-Effective Active Learning for Deep Image Classification. *IEEE Trans. Circuits Syst. Video Techn.* 27, 12 (2017), 2591–2600.
- [34] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Nikolai Yakovenko, Andrew Tao, Jan Kautz, and Bryan Catanzaro. 2018. Video-to-Video Synthesis. In *Advances in Neural Information Processing Systems 31: Annual Conference on Neural Information Processing Systems*. 1152–1164.
- [35] Zheng Wang and Jieping Ye. 2015. Querying Discriminative and Representative Samples for Batch Mode Active Learning. *TKDD* 9, 3 (2015), 17:1–17:23.
- [36] Yifan Yan and Sheng-Jun Huang. 2018. Cost-Effective Active Learning for Hierarchical Multi-Label Classification. In *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018*. 2962–2968.
- [37] Lantao Yu, Weinan Zhang, Jun Wang, and Yong Yu. 2017. SeqGAN: Sequence Generative Adversarial Nets with Policy Gradient. In *AAAI*. 2852–2858.
- [38] Han Zhang, Tao Xu, Hongsheng Li, Shaoting Zhang, Xiaogang Wang, Xiao lei Huang, and Dimitris N Metaxas. 2017. Stackgan: Text to photo-realistic image synthesis with stacked generative adversarial networks. In *Proceedings of the IEEE International Conference on Computer Vision*. 5907–5915.
- [39] Zhiming Zhou, Han Cai, Shu Rong, Yuxuan Song, Kan Ren, Weinan Zhang, Yong Yu, and Jun Wang. 2017. Activation Maximization Generative Adversarial Nets. *arXiv preprint arXiv:1703.02000* (2017).
- [40] Jia-Jie Zhu and José Bento. 2017. Generative Adversarial Active Learning. *arXiv preprint arXiv:1702.07956* abs/1702.07956 (2017).