

---

# Generative Adversarial Active Learning

---

**Jia-Jie Zhu**

Department of Computer Science  
Boston College  
Chestnut Hill, Massachusetts, USA  
zhuuv@bc.edu

**José Bento**

Department of Computer Science  
Boston College  
Chestnut Hill, Massachusetts, USA  
jose.bento@bc.edu

## Abstract

We propose a new active learning approach using Generative Adversarial Networks (GAN). Different from regular active learning, the resulting algorithm adaptively synthesizes training instances for querying to increase learning speed. We generate queries according to the uncertainty principle, but our idea can work with other active learning principles. We report results from various numerical experiments to demonstrate the effectiveness of the proposed approach. In some settings, the proposed algorithm outperforms traditional pool-based approaches. To the best of our knowledge, this is the first active learning work using GAN.

## 1 Introduction

One of the most exciting machine learning breakthroughs in recent years is the generative adversarial networks (GAN) [18]. It trains a generative model by finding the Nash Equilibrium of a two-player adversarial game. Its ability to generate samples in complex domains enables new possibilities for active learners to synthesize training samples on demand, rather than relying on choosing instances to query from a given pool.

In the classification setting, given a pool of unlabeled data samples and a fixed labeling budget, active learning algorithms typically choose training samples strategically from a pool to maximize the accuracy of trained classifiers. The goal of these algorithms is to reduce label complexity. Such approaches are called pool-based active learning. This pool-based active learning approach is illustrated in Figure 1 (a).

In a nutshell, we propose to use GANs to synthesize informative training instances that are adapted to the current learner. We then ask human oracles to label these instances. The labeled data is added back to the training set to update the learner. This protocol is executed iteratively until the label budget is reached. This process is shown in Figure 1 (b).




Figure 1: (a) Pool-based active learning scenario. The learner selects samples for querying from a given unlabeled pool. (b) GAAL algorithm. The learner synthesizes samples for querying using GAN.

The main contributions of this work are as follows:

- To the best of our knowledge, this is the first active learning framework using deep generative models<sup>1</sup>. In fact, this is the first work to report numerical results in active learning synthesis for image classification. See [37, 25]. The proposed framework may inspire future GAN applications in active learning.
- While we do not claim our method is always superior to the previous active learners in terms of accuracy, in some cases, it yields classification performance not achievable even by a fully supervised learning scheme. With enough capacity from the trained generator, our method allows us to have control over the generated instances which may not be available to the previous active learners.
- We conduct experiments to compare our active learning approach with self-taught learning. The results are promising.

## 2 Related Work

Our work is related to two different subjects, active learning and deep generative models.

Active learning algorithms can be categorized into stream-based, pool-based and learning by query synthesis. Historically, stream-based and pool-based are the two popular scenarios of active learning [37].

Our method falls into the category of query synthesis. Early active learning by queries synthesis achieves good results only in simple domains such as  $X = \{0, 1\}^3$ , see [1, 2]. In [25], the authors synthesized learning queries and used human oracles to train a neural network for classifying handwritten characters. However, they reported poor results due to the images generated by the learner being sometimes unrecognizable to the human oracles. We will report results on similar tasks such as differentiating 5 versus 7, showing the advancement of our active learning scheme. Figure 2 compares image samples generated by the method in [25] and our algorithm.



Figure 2: (Left) Image queries synthesized by a neural network for handwritten digits recognition. Source: [25]. (Right) Image queries synthesized by our algorithm, GAAL.

The popular  $SVM_{active}$  algorithm from [39] is an efficient pool-based active learning scheme for SVM. Their scheme is a special instance of the uncertainty sampling principle which we also employ. [23] reduces the exhaustive scanning through database employed by  $SVM_{active}$ . Our algorithm shares the same advantage of not needing to test every sample in the database at each iteration of active learning. Although we do so by not using a pool at all instead of a clever trick. [42] proposed active transfer learning which is reminiscent to our experiments in Section 5.1. However, we do not consider collecting new labeled data in target domains of transfer learning.

There have been some applications of generative models in semi-supervised learning and active learning. Previously, [31] proposed a semi-supervised learning approach to text classification based on generative models. [22] applied Gaussian mixture models to active learning. In that work, the generative model served as a classifier. Compared with these approaches, we apply generative models to directly synthesize training data. This is a more challenging task.

One building block of our algorithm is the groundbreaking work of the GAN model in [18]. Our approach is an application of GAN in active learning.

Our approach is also related to [38] which studied GAN in a semi-supervised setting. However, our task is active learning which is different from the semi-supervised learning they discussed. Our work

<sup>1</sup>The appendix of [32] mentioned three active learning attempts but did not report numerical results. Our approach is also different from those attempts.

shares the common strength with the self-taught learning algorithm in [34] as both methods use the unlabeled data to help with the task. In Section 5.2, we compare our algorithm with a self-taught learning algorithm.

To the best of our knowledge, the only previous mentioning of using GAN for active learning is in the appendix of [32]. The authors discussed therein three attempts to reduce the number of queries. In the third attempt, they generated synthetic samples and sorted them by the information content whereas we adaptively generate new queries by solving an optimization problem. There were no reported active learning numerical results in that work.

### 3 Background

We briefly introduce some important concepts in active learning and generative adversarial network.

#### 3.1 Active Learning

In the PAC learning framework [40], label complexity describes the number of labeled instances needed to find a hypothesis with error  $\epsilon$ . The label complexity of passive supervised learning, i.e. using all the labeled samples as training data, is  $\mathcal{O}(\frac{d}{\epsilon})$  [41], where  $d$  is the VC dimension of the hypothesis class  $\mathcal{H}$ . Active learning aims to reduce the label complexity by choosing the most informative instances for querying while attaining low error rate. For example, [20] proved that the active learning algorithm from [9] has the label complexity bound  $\mathcal{O}(\theta d \log \frac{1}{\epsilon})$ , where  $\theta$  is defined therein as the disagreement coefficient, thus reducing the theoretical bound for the number of labeled instances needed from passive supervised learning. Theoretically speaking, the asymptotic accuracy of an active learning algorithm can not exceed that of a supervised learning algorithm. In practice, as we will demonstrate in the experiments, our algorithm may be able to achieve higher accuracy than the passive supervised learning in some cases.

Stream-based active learning makes decisions on whether to query the streamed-in instances or not. Typical methods include [4, 9, 13]. In this work, we will focus on comparing pool-based and query synthesis methods.

In pool-based active learning, the learner selects the unlabeled instances from an existing pool based on a certain criterion. Some pool-based algorithms make selections by using clustering techniques or maximizing a diversity measure, e.g. [6, 44, 12, 30, 45, 21]. Another commonly used pool-based active learning principle is uncertainty sampling. It amounts to querying the most uncertain instances. For example, algorithms in [39, 7] query the labels of the instances that are closest to the decision boundary of the support vector machine. Figure 3 (a) illustrates this selection process.

Mathematically, let  $P$  be the pool of unlabeled instances, and  $f = W\phi(x) + b$  be the separating hyperplane.  $\phi$  is the feature map induced by the SVM kernel. The  $\text{SVM}_{active}$  algorithm in [39] chooses a new instance to query by minimizing the distance (or its proxy) to the hyperplane

$$\min_{x \in P} \|W\phi(x) + b\|. \tag{1}$$

This formulation can be justified by the version space theory in separable cases [39] or by other analyses in non-separable cases, e.g., [7, 5]. This simple and effective method is widely applied in many studies, e.g., [16, 43].

In the query synthesis scenario, an instance  $x$  is synthesized instead of being selected from an existing pool. Previous methods tend to work in simple low-dimensional domains [2] but fail in more complicated domains such as images [25]. Our approach aims to tackle this challenge.

For an introduction to active learning, readers are referred to [37, 11].

#### 3.2 Generative Adversarial Networks

Generative adversarial networks (GAN) is a novel generative model invented by [18]. It can be viewed as the following two-player minimax game between the generator  $G$  and the discriminator  $D$ ,

$$\min_{\theta_2} \max_{\theta_1} \left\{ \mathbb{E}_{x \sim p_{data}} \log D_{\theta_1}(x) + \mathbb{E}_z \log(1 - D_{\theta_1}(G_{\theta_2}(z))) \right\}, \tag{2}$$

where  $p_{\text{data}}$  is the underlying distribution of the real data and  $z$  is uniformly distributed random variable.  $D$  and  $G$  each has its own set of parameter  $\theta_1$  and  $\theta_2$ . By solving this game, a generator  $G$  is obtained. In the ideal scenario, given random input  $z$ , we have  $G(z) \sim p_{\text{data}}$ . However, finding this Nash Equilibrium is a difficult problem in practice. There is no theoretical guarantee for finding the Nash Equilibrium due to the non-convexity of  $D$  and  $G$ . A gradient descent type algorithm is typically used for solving this optimization problem.

A few variants of GAN have been proposed since [18]. The authors of [33] use GAN with deep convolutional neural network structures for applications in computer vision(DCGAN). DCGAN yields good results and is relatively stable. Conditional GAN[15, 14, 29] is another variant of GAN in which the generator and discriminator can be conditioned on other variables, e.g., the labels of images. Such generators can be controlled to generate samples from a certain category. [8] proposed infoGAN which learns disentangled representations using unsupervised learning.

A few updated GAN models have been proposed. [35] proposed a few improved techniques for training GAN. Another potentially important improvement of GAN, Wasserstein GAN, has been proposed by [3, 19]. The authors proposed an alternative to training GAN which can avoid instabilities such as mode collapse with theoretical analysis. They also proposed a metric to evaluate the quality of the generation which may be useful for future GAN studies. Possible applications of Wasserstein GAN to our active learning framework are left for future work.

The invention of GAN triggered various novel applications. [46] performed image inpainting task using GAN. [47] proposed iGAN to turn sketches into realistic images. [28] applied GAN to single image super-resolution. [48] proposed CycleGAN for image-to-image translation using only unpaired training data. This work is the first GAN application to active learning.

For a comprehensive review of GAN, readers are referred to [17].

## 4 Generative Adversarial Active Learning

In this section, we introduce our active learning approach which we call Generative Adversarial Active Learning (GAAL). It combines query synthesis with the uncertainty sampling principle.

The intuition of our approach is to generate instances which the current learner is uncertain about, i.e. applying the uncertainty sampling principle. One particular choice for the loss function is based on uncertainty sampling principle explained in section 3.1. In the setting of a classifier with the decision function  $f(x) = W\phi(x) + b$ , the (proxy) distance to the decision boundary is  $\|W\phi(x) + b\|$ . Similar to the intuition of (1), given a trained generator function  $G$ , we formulate the active learning synthesis as the following optimization problem

$$\min_z \|W^\top \phi(G(z)) + b\|, \tag{3}$$

where  $z$  is the latent variable and  $G$  is obtained by the GAN algorithm. Intuitively, minimizing this loss will push the generated samples toward the decision boundary. Figure 3 (b) illustrates this idea. Compared with the pool-base active learning in Figure 3 (a), our hope is that it may be able to generate more informative instances than those available in the existing pool.



Figure 3: (a)  $SVM_{\text{active}}$  algorithm selects the instances that are closest to the boundary to query the oracle. (b) GAAL algorithm synthesizes instances that are informative to the current learner. Synthesized instances may be more informative to the learner than other instances in the existing pool.

The solution(s) to this optimization problem,  $G(z)$ , after being labeled, will be used as new training data for the next iteration. We outline our procedure in Algorithm 1. It is possible to use a state-of-

---

**Algorithm 1** Generative Adversarial Active Learning (GAAL)

---

- 1: Train generator  $G$  on all unlabeled data by solving (2)
- 2: Initialize labeled training dataset  $\mathcal{S}$  by randomly picking a small fraction of the data to label
- 3: **repeat**
- 4:   Solve optimization problem (3) according to the current learner by descending the gradient

$$\nabla_z \|W^\top \phi(G(z)) + b\|$$

- 5:   Use the solution  $\{z_1, z_2, \dots\}$  and  $G$  to generate instances for querying
  - 6:   Label  $\{G(z_1), G(z_2), \dots\}$  by human oracles
  - 7:   Add labeled data to the training dataset  $\mathcal{S}$  and re-train the learner, update  $W, b$
  - 8: **until** Labeling budget is reached
- 

the-art classifier, such as convolutional neural networks. To do this, we can replace the feature map  $\phi$  in Equation 3 with a feed-forward function of a convolutional neural network. In that case, the linear SVM will become the output layer of the network.

In training GAN, we follow the procedure detailed in [33]. Optimization problem (3) is non-convex with possibly many local minima. One typically aims at finding good local minima rather than the global minimum. We use a gradient descent algorithm with momentum to solve this problem. We also periodically restart the gradient descent to find other solutions. The gradient of  $D$  and  $G$  is calculated using back-propagation.

Alternatively, we can incorporate diversity into our active learning principle. Some active learning approaches rely on maximizing diversity measures, such as the Shannon Entropy. In our case, we can include in the objective function (3) a diversity measure such as proposed in [45, 21], thus increasing the diversity of samples. The evaluation of this alternative approach is left for future work.

## 5 Experiments

We perform active learning experiments using the proposed approach. We also compare our approach to self-taught learning, a type of transfer learning method. The GAN implementation used in our experiment is a modification of a publicly available TensorFlow DCGAN implementation<sup>2</sup>. The network architecture of DCGAN is described in [33].

In our experiments, we focus on binary image classification. Although this can be generalized to multiple classes using one-vs-one or one-vs-all scheme [24]. Recent advancements in GAN study show it could potentially model language as well [19]. Although those results are preliminary at the current stage. We use a linear SVM as our classifier of choice (with parameter  $\gamma = 0.001$ ). Even though classifiers with much higher accuracy (e.g., convolutional neural networks) can be used, our purpose is not to achieve absolute high accuracy but to study the relative performance between different active learning schemes.

The following schemes are implemented and compared in our experiments.

- The proposed generative adversarial active learning (GAAL) algorithm as in Algorithm 1.
- Using regular GAN to generate training data. We refer to this as simple GAN.
- SVM<sub>active</sub> algorithm from [39].
- Passive random sampling, which randomly samples instances from the unlabeled pool.
- Passive supervised learning, i.e., using all the samples in the pool to train the classifier.
- Self-taught learning from [34].

We initialize the training set with 50 randomly selected samples. The algorithms proceed with a batch of 10 queries every time.

We use two datasets for training, the MNIST and CIFAR-10. The MNIST dataset is a well-known image classification dataset with 60000 training samples. The training set and the test set follow the

---

<sup>2</sup><https://github.com/carpedm20/DCGAN-tensorflow>

same distribution. We perform the binary classification experiment distinguishing 5 and 7 which is reminiscent to [25]. The training set of CIFAR-10 dataset consists of 50000  $32 \times 32$  color images from 10 categories. One might speculate the possibility of distinguishing cats and dogs by training on cat-like dogs or dog-like cats. In practice, our human labelers failed to confidently identify most of the generated cat and dog images. Figure 4 shows generated samples. The authors of [35] reported attempts to generate high-resolution animal pictures, but with the wrong anatomy. We leave this task for future studies, possibly with improved techniques such as [3, 19]. For this reason, we perform



Figure 4: Generated samples in cat and dog categories

binary classification on the automobile and horse categories. It is relatively easy for human labelers to identify car and horse body shapes. Typical generated samples, which are presented to the human labelers, are shown in Figure 5.



Figure 5: Samples generated by GAAL (Left) MNIST dataset (Right) CIFAR-10 dataset

### 5.1 Active Learning

We use all the images of 5 and 7 from the MNIST training set as our unlabeled pool to train the generator  $G$ . Different from traditional active learning, we do not select new samples from the pool after initialization. Instead, we apply Algorithm 1 to generate a training query. For the generator  $D$  and  $G$ , we follow the same network architecture of [33]. We use linear SVM as our classifier although other classifiers can be used, e.g. [39, 36, 37].

We first test the trained classifier on a test set that follows a distribution different from the training set. One purpose is to demonstrate the adaptive capability of the GAAL algorithm. In addition, because the MNIST test set and training set follow the same distribution, pool-based active learning methods have a natural advantage over active learning by synthesis since they use real images drawn from the exact same distribution as the test set. It is thus reasonable to test on sets that follow different, albeit similar, distributions. To this end, we use the USPS dataset from [27] as the test set with standard preprocessing. In reality, such settings are very common, e.g., training autonomous drivers on simulated datasets and testing on real vehicles; training on handwriting characters and recognizing writings in different styles, etc. This test setting is related to transfer learning, where the distribution of the training domain  $P_{tr}(x, y)$  is different from that of the target domain  $P_{te}(x, y)$ . Figure 6 shows the results of our first experiment.

When using the full training set, with 11000 training images, the fully supervised accuracy is at 70.44%. The accuracy of the random sampling scheme steadily approaches that level. On the other hand, GAAL is able to achieve accuracies better than that of the fully supervised scheme. With 350 training samples, its accuracy improves over supervised learning and even  $SVM_{active}$ , an aggressive active learner [10, 39]. Obviously, the accuracy of both  $SVM_{active}$  and random sampling will eventually converge to the fully supervised learning accuracy. Note that for the  $SVM_{active}$  algorithm, an exhaustive scan through the training pool is not always practical. In such cases, the common practice is to restrict the selection pool to a small random subset of the original data.

For completeness, we also perform the experiments in the settings where the training and test set follow the same distribution. Figure 7 (Center, Right) shows these results. Somewhat surprisingly, in Figure 7 (Left), GAAL’s classification accuracy starts to drop after about 100 samples. One possible explanation is that GAAL may be generating points close to the boundary that are also close to each other. This is more likely to happen if the boundary does not change much from one active learning cycle to the next. This probably happens because the test and train sets are the identically distributed and simple, like MNIST. Therefore, after a while, the training set may be filled with many similar points, biasing the classifier and hurting accuracy. In contrast, because of the finite and discrete nature of pools in the given datasets, a pool-based approach, such as  $SVM_{active}$ , most likely explores




Figure 6: Active learning results. Train on MNIST, test on USPS. Classifying 5 and 7. The results are averaged over 10 runs. The error bars represent the empirical standard deviation of the average values. The figure is best viewed in color.

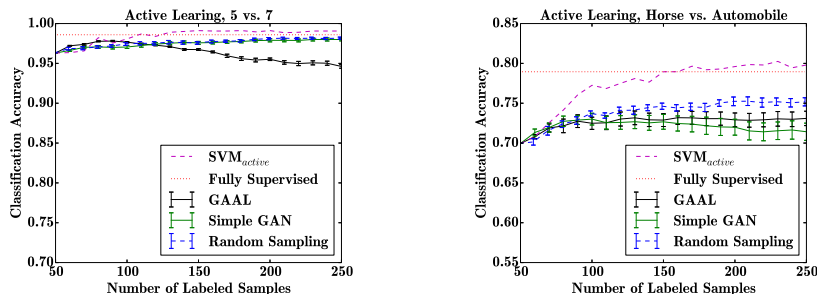


Figure 7: Active learning results. (Left) Train on MNIST, test on MNIST. Classifying 5 and 7. (Right) CIFAR-10 dataset, classifying automobile and horse. The results are averaged over 10 runs. The error bars represent the empirical standard deviation of the average values. The figures are best viewed in color.

points near the boundary that are substantially different. It is also forced to explore further points once these close-by points have already been selected. In a sense, the strength of GAAL might in fact be hurting its classification accuracy. We believe this effect is not so pronounced when the test and train sets are different because the boundary changes more significantly from one cycle to the next, which in turn induces some diversity in the generated samples.

To reach competitive accuracy when the training and test set follow the same distribution, we might incorporate a diversity term into our objective function in GAAL. We will address this in future work.

In the CIFAR-10 dataset, our human labeler noticed higher chances of bad generated samples, e.g., instances fail to represent either of the categories. This may be because of the significantly higher dimensions than the MNIST dataset. In such cases, we asked the labelers to only label the samples they can distinguish. We speculate recent improvements on GAN, e.g., [35, 3, 19], may help mitigate this issue given the cause is the instability of GANs. Addressing this limitation will be left to future studies.

## 5.2 Comparison with Self-taught Learning

One common strength of GAAL and self-taught learning [34] is that both utilize the unlabeled data to help with the classification task. As we have seen in the MNIST experiment, our GAAL algorithm seems to be able to adapt to the learner. The results in this experiment are preliminary and not meant to be taken as comprehensive evaluations.

In this case, the training domain is mostly unlabeled. Thus the method we compare with is self-taught learning [34]. Similar to the algorithm in [26], we use a Reconstruction Independent Com-

Table 1: Comparison of GAAL and self-taught learning

ALGOIRTHM	TRAINING SET SIZE	ACCURACY
<b>GAAL</b>	<b>250</b>	<b>76.42%</b>
SELF-TAUGHT	250	59.68%
SUPERVISED	250	67.87%
SELF-TAUGHT	500	65.53%
SUPERVISED	500	69.22%
SELF-TAUGHT	1000	71.96%
SUPERVISED	1000	69.58%
<b>SELF-TAUGHT</b>	<b>5000</b>	<b>78.08%</b>
SUPERVISED	5000	72.00%

ponent Analysis (RICA) model with a convolutional layer and a pooling layer. RICA is similar to a sparse autoencoder. Following standard self-taught learning procedures, We first train on the unlabeled pool dataset. Then we use trained RICA as the a feature extractor to obtain higher level features from randomly selected MNIST images. We then concatenate the features with the original image data to train the classifier. Finally, we test the trained classifier on the USPS dataset. We test the training size of 250, 500, 1000, and 5000. The reason of doing so is that deep learning type techniques are known to thrive in the abundance of training data. They may perform relatively poorly with limited amount of training data, as in the active learning scenarios. We run the experiments for 100 times and average the results. We use the same setting for the GAAL algorithm as in Section 5.1. The classifier we use is a linear SVM. Table 1 shows the classification accuracies of GAAL, self-taught learning and baseline supervised learning on raw image data. Using GAAL on the raw features achieves a higher accuracy than that of the self-taught learning with the same training size of 250. In fact, self-taught learning performs worse than the regular supervised learning when labeled data is scarce. This is possible for an autoencoder type algorithm. However, when we increase the training size, the self-taught learning starts to perform better. With 5000 training samples, self-taught learning outperforms GAAL with 250 training samples.

Based on these results, we suspect that GAAL also has the potential to be used as a self-taught algorithm<sup>3</sup>. In practice, the GAAL algorithm can also be applied on top of the features extracted by a self-taught algorithm. A comprehensive comparison with a more advanced self-taught learning method with deeper architecture is beyond the scope of this work.

## 6 Discussion and Future Work

In this work, we proposed a new active learning approach, GAAL, that employs the generative adversarial networks. One possible explanation for GAAL not outperforming the pool-based approaches in some settings is that, in traditional pool-based learning, the algorithm will eventually exhaust all the points near the decision boundary thus start exploring further points. However, this is the not the case in GAAL as it can always synthesize points near the boundary. This may in turn cause the generation of similar samples, thus reducing the effectiveness. We suspect incorporating a diversity measure into the GAAL framework as discussed at the end of Section 4 might mitigate this issue. This is left for future work.

The results of this work are enough to inspire future studies of deep generative models in active learning. However, much work remains in establishing theoretical analysis and reaching better performance. The comparison of GAAL with transfer learning is particularly interesting and worth further investigation. We also plan to investigate the possibility of using Wasserstein GAN in our framework.

## References

- [1] D Angluin. Queries and concept learning. *Mach. Learn.*, 1988.
- [2] D Angluin. Queries revisited. *Int. Conf. Algorithmic Learn.*, 2001.

<sup>3</sup>At this stage, self-taught learning has the advantage that it can utilize any unlabeled training data, i.e., not necessarily from the categories of interest. GAAL does not have this feature yet.



- [3] Martin Arjovsky, Soumith Chintala, and Léon Bottou. Wasserstein GAN. jan 2017.
- [4] Alina Beygelzimer, Sanjoy Dasgupta, and John Langford. Importance Weighted Active Learning. *Proc. 26th Annu. Int. Conf. Mach. Learn. ICML 09*, abs/0812.4(ii):1–8, 2008.
- [5] Antoine Bordes, Şeyda Ertekin, Jason Weston, and Léon Bottou. Fast Kernel Classifiers with Online and Active Learning. *J. Mach. Learn. Res.*, 6:1579–1619, 2005.
- [6] Klaus Brinker. Incorporating Diversity in Active Learning with Support Vector Machines.
- [7] Colin Campbell, Nello Cristianini, and Alex Smola. Query learning with large margin classifiers. *17th Int. Conf. Mach. Learn.*, pages 111–118, 2000.
- [8] Xi Chen, Yan Duan, Rein Houthoofd, John Schulman, Ilya Sutskever, and Pieter Abbeel. InfoGAN: Interpretable Representation Learning by Information Maximizing Generative Adversarial Nets. 2016.
- [9] David Cohn, Les Atlas, and Richard Ladner. Improving generalization with active learning. *Mach. Learn.*, 15(2):201–221, may 1994.
- [10] Sanjoy Dasgupta. Analysis of a greedy active learning strategy. In *Advances in neural information processing systems*, pages 337–344, 2005.
- [11] Sanjoy Dasgupta. Two faces of active learning. *Theor. Comput. Sci.*, 412:1767–1781, 2011.
- [12] Sanjoy Dasgupta and Daniel Hsu. Hierarchical sampling for active learning. *Proceedings of the 25th international conference on Machine learning - ICML '08*, pages 208–215, 2008.
- [13] Sanjoy Dasgupta, Daniel Hsu, and Claire Monteleoni. A general agnostic active learning algorithm. *Engineering*, 20(2):1–14, 2007.
- [14] Alexey Dosovitskiy, Jost Tobias Springenberg, Maxim Tatarchenko, and Thomas Brox. Learning to Generate Chairs, Tables and Cars with Convolutional Networks. *arXiv preprint arXiv:1411.5928*, pages 1–14, 2014.
- [15] Jon Gauthier. Conditional generative adversarial nets for convolutional face generation. *Class Project for Stanford CS231N: Convolutional Neural Networks for Visual Recognition, Winter semester 2014*, 2014.
- [16] King-Shy Goh, Edward Y. Chang, and Wei-Cheng Lai. Multimodal concept-dependent active learning for image retrieval. In *Proc. 12th Annu. ACM Int. Conf. Multimed. - Multimed. '04*, page 564, New York, New York, USA, 2004. ACM Press.
- [17] Ian Goodfellow, Yoshua Bengio, and Aaron Courville. *Deep Learning*. MIT Press, 2016.
- [18] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. In *Advances in neural information processing systems*, pages 2672–2680, 2014.
- [19] Ishaan Gulrajani, Faruk Ahmed, Martin Arjovsky, Vincent Dumoulin, and Aaron Courville. Improved training of wasserstein gans. *arXiv preprint arXiv:1704.00028*, 2017.
- [20] Steve Hanneke. A bound on the label complexity of agnostic active learning. *Proc. 24th Int. Conf. Mach. Learn. - ICML '07*, pages 353–360, 2007.
- [21] Steven C H Hoi, Rong Jin, Jianke Zhu, and Michael R Lyu. Semi-Supervised SVM Batch Mode Active Learning with Applications to Image Retrieval. *ACM Trans. Informations Syst. ACM Trans. Inf. Syst. Publ. ACM Trans. Inf. Syst.*, 27(16):24–26, 2009.
- [22] Timothy M. Hospedales, Shaogang Gong, and Tao Xiang. Finding rare classes: Active learning with generative and discriminative models. *IEEE Trans. Knowl. Data Eng.*, 25(2):374–386, 2013.
- [23] Prateek Jain, Sudheendrasvnraras Vijayanarasimhan, Kristen Grauman, Prateek Jain, and Kristen Grauman. Hashing Hyperplane Queries to Near Points with Applications to Large-Scale Active Learning. *IEEE Trans. Pattern Anal. Mach. Intell.*, 36(2):2010, 2010.
- [24] A.J. Joshi, F. Porikli, and N. Papanikolopoulos. Multi-class active learning for image classification. *IEEE Conf. Comput. Vis. Pattern Recognit.*, pages 2372–2379, 2009.
- [25] Kevin J. Lang and Eric B Baum. Query Learning Can Work Poorly when a Human Oracle is Used, 1992.

- [26] Quoc V Le, Alexandre Karpenko, Jiquan Ngiam, and Andrew Y Ng. ICA with Reconstruction Cost for Efficient Overcomplete Feature Learning.
- [27] Y. LeCun, B. Boser, J. S. Denker, D. Henderson, R. E. Howard, W. Hubbard, and L. D. Jackel. Backpropagation Applied to Handwritten Zip Code Recognition, 1989.
- [28] Christian Ledig, Lucas Theis, Ferenc Huszar, Jose Caballero, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, and Wenzhe Shi. Photo-Realistic Single Image Super-Resolution Using a Generative Adversarial Network. *arXiv*, 2016.
- [29] Mehdi Mirza and Simon Osindero. Conditional Generative Adversarial Nets. *CoRR*, pages 1–7, nov 2014.
- [30] Hieu T Nguyen and Arnold Smeulders. Active Learning Using Pre-clustering.
- [31] Kamal Nigam, Andrew Kachites McCallum, Sebastian Thrun, and Tom Mitchell. Text Classification from Labeled and Unlabeled Documents using EM. *Mach. Learn.*, 39:103–134, 2000.
- [32] Nicolas Papernot, Martín Abadi, Úlfar Erlingsson, Ian Goodfellow, and Kunal Talwar. Semi-supervised knowledge transfer for deep learning from private training data. *arXiv preprint arXiv:1610.05755*, 2016.
- [33] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. nov 2015.
- [34] Rajat Raina, Alexis Battle, Honglak Lee, Benjamin Packer, and Andrew Y Ng. Self-taught Learning : Transfer Learning from Unlabeled Data. *Proc. 24th Int. Conf. Mach. Learn.*, pages 759–766, 2007.
- [35] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. Improved Techniques for Training GANs. jun 2016.
- [36] Andrew I. Schein and Lyle H. Ungar. *Active learning for logistic regression: An evaluation*, volume 68. 2007.
- [37] Burr Settles. Active learning literature survey. *Computer sciences technical report*, 1648:University of Wisconsin–Madison, 2010.
- [38] Jost Tobias Springenberg. Unsupervised and Semi-supervised Learning with Categorical Generative Adversarial Networks. *arXiv*, (2009):1–20, 2015.
- [39] Simon Tong and Daphne Koller. Support Vector Machine Active Learning with Applications to Text Classification. *Proc. Int. Conf. Mach. Learn.*, 1(June):45–66, 2002.
- [40] L. G. Valiant and L. G. A theory of the learnable. *Commun. ACM*, 27(11):1134–1142, nov 1984.
- [41] VN Vapnik and V Vapnik. *Statistical learning theory*. 1998.
- [42] Xuezhi Wang, Tzu-Kuo Huang, and Jeff Schneider. Active transfer learning under model shift. In *International Conference on Machine Learning*, pages 1305–1313, 2014.
- [43] Manfred K Warmuth, Jun Liao, Gunnar Rätsch, Michael Mathieson, Santosh Putta, and Christian Lemmen. Active Learning with Support Vector Machines in the Drug Discovery Process. 2002.
- [44] Z Xu, R Akella, and Y Zhang. Incorporating diversity and density in active learning for relevance feedback. *European Conference on Information Retrieval*, 2007.
- [45] Yi Yang, Zhigang Ma, Feiping Nie, Xiaojun Chang, and Alexander G Hauptmann. Multi-Class Active Learning by Uncertainty Sampling with Diversity Maximization. *Int. J. Comput. Vis.*, 113(2):113–127, jun 2014.
- [46] Raymond Yeh, Chen Chen, Teck Yian Lim, Mark Hasegawa-Johnson, and Minh N. Do. Semantic Image Inpainting with Perceptual and Contextual Losses. jul 2016.
- [47] Jun-Yan Zhu, Philipp Krähenbühl, Eli Shechtman, and Alexei A. Efros. Generative Visual Manipulation on the Natural Image Manifold. pages 597–613. Springer, Cham, 2016.
- [48] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. *arXiv preprint arXiv:1703.10593*, 2017.