EC-GAN: Low-Sample Classification using Semi-Supervised Algorithms and GANs (Student Abstract)

Ayaan Haque
Saratoga High School, Saratoga, CA, USA
ayaanzhaque@gmail.com

Abstract

Semi-supervised learning has been gaining attention as it allows for performing image analysis tasks such as classification with limited labeled data. Some popular algorithms using Generative Adversarial Networks (GANs) for semi-supervised classification share a single architecture for classification and discrimination. However, this may require a model to converge to a separate data distribution for each task, which may reduce overall performance. While progress in semi-supervised learning has been made, less addressed are small-scale, fully-supervised tasks where even unlabeled data is unavailable and unattainable. We therefore, propose a novel GAN model namely *External Classifier GAN (EC-GAN)*, that utilizes GANs and semi-supervised algorithms to improve classification in fully-supervised regimes. Our method leverages a GAN to generate artificial data used to supplement supervised classification. More specifically, we attach an external classifier, hence the name EC-GAN, to the GAN's generator, as opposed to sharing an architecture with the discriminator. Our experiments demonstrate that EC-GAN's performance is comparable to the shared architecture method, far superior to the standard data augmentation and regularization-based approach, and effective on a small, realistic dataset.

Introduction

Deep generative model such as generative adversarial networks (GANs) facilitate generation of realistic images mimicking real data (Goodfellow et al. 2014). In GANs, two neural networks, a generator and a discriminator, are trained together, where the generator attempts to generate images resembling real training samples while the discriminator distinguishes the generated samples from the generated ones. Many existing methods which use GANs for semi-supervised learning employ a single network for both classification and discrimination (Salimans et al. 2016; Imran and Terzopoulos 2019). This means the network attempts to minimize two separate losses with the same parameters, which is our primary concern. Very few algorithms have attempted to separate the two networks, but those that have, such as Triple GAN (Li et al. 2017), have only been used in semi-supervised scenarios, focus on improved discrimination, and most importantly do not use artificial samples for classification. Addressing these concerns, we propose EC-GAN to aid restricted, fully-supervised learning with artificial data, while uniquely separating the tasks of discrimination and classification.

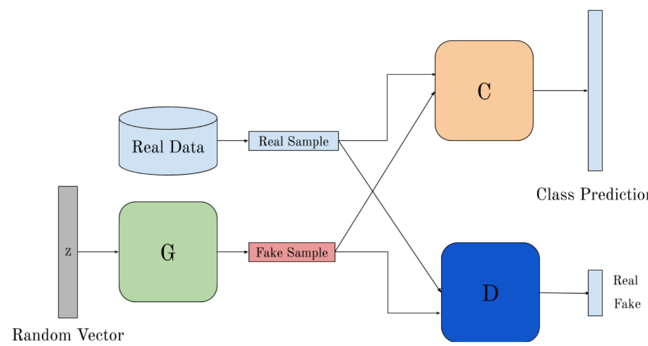

Figure 1: EC-GAN has an external classifier which is trained on generated and real images simultaneously. The discriminator and generator are trained conventionally, and the discriminator's sole output is the predicted probability of the image being real. Pseudo-labels are created for artificial images.

Methods

The EC-GAN model consists of three separate networks: a generator, a discriminator, and a classifier. For our architectures, we use the DCGAN (Radford, Metz, and Chintala 2016) and ResNet18 for the GAN and classifier respectively. At every training iteration, the generator is given random noise vectors and generates new images. The discriminator takes either the real or the generated samples and distinguishes between them.

Simultaneously, a classifier is trained in supervised fashion on the labeled data. We then use generated images as inputs for supplementing classification during training. To create labels for the artificial samples, we use a pseudo-labeling scheme which assumes a label based on the most likely class prediction according to the current state of the classifier (Lee 2013). The labels are only retained if the probability is above a specific confidence threshold.

Our training objective utilizes both supervised and unsupervised loss terms. As in standard GANs, the discriminator loss is defined as $L_D(x, z) = BCE(D(x), 1) +$

Dataset Size (%)	EC-GAN (%)		Shared DCDiscriminator	
	Classifier	GAN	Classifier	GAN
10	88.63	91.15	83.54	86.17
15	90.88	92.21	85.20	88.72
20	92.61	93.40	86.77	89.39
25	92.89	93.93	87.58	87.93
30	93.12	94.32	87.78	90.62

Table 1: EC-GAN is compared to the shared architecture method on SVHN at different dataset sizes. Left value is accuracy of a standard classifier (same architecture as GAN counterpart), followed by the accuracy of GAN classification algorithm.

$BCE(D(G(z)), 0)$, and the generator loss is defined as $L_G(z) = BCE(D(G(z)), 1)$, where BCE is binary cross-entropy, D is the discriminator, G is the generator, x is real, labeled data, and z is a random vector. For our classifier, the loss is defined as $L_C(x, y, z) = CE(C(x), y) + \lambda CE(C(G(z)), argmax(C(G(z))) > t)$, where λ is the unsupervised adversarial loss weight, CE is cross entropy loss, C is the classifier, and t is the pseudo-label threshold.

The first component of the classifier loss is the supervised cross-entropy loss using real labeled data. The next component is the unsupervised loss. We compute the cross-entropy between generated data and corresponding hypothesized labels. λ is a weight for the unsupervised loss, and similar approaches have seen success in semi-supervised learning regimes (using real instead of artificial data). We incorporate λ to regularize our model, because generated images are only meant to supplement real ones. The "labels" for the unsupervised loss are pseudo-labels, where we use the threshold t to ensure only high-quality GAN generations are used.

Experimental Evaluation

Throughout our tests, we manipulated the number of sampled used in the benchmark dataset SVHN (used for development and testing) by percent size of the dataset to show the efficacy of our algorithm in small datasets. To test our algorithm in a small, fully-supervised dataset, we used a pneumonia chest X-Ray dataset with just 5,863 images, or less than 10% of SVHN's size (Kermany, Zhang, and Goldbaum 2018).

Since the regular GAN approaches use a shared, two-headed architecture, and our EC-GAN is unique because the classifier is distinctly separate, we compared the performance of the two methods. We found that EC-GAN performs on par and occasionally better than a shared architecture in small datasets (Table 1). Through our empirical analysis, we contend that separating classification and discrimination and supplementing classification with generated images are effective approaches for improving classification performance.

Additionally, we found EC-GAN performed much better than the standard data augmentation and regularization procedures. Better generalization and improved classification performance can often be achieved with more data, which could explain why EC-GAN outperforms other methods with augmentation and regularization. Our experimental results justify the effectiveness of EC-GAN as a semi-supervised approach for restricted fully-supervised classification.

Dataset Size (%)	EC-GAN (%)	
	Classifier	GAN
25	94.37	96.48
50	95.24	97.83
75	95.64	97.40
100	96.42	97.99

Table 2: The conditional version of EC-GAN is tested on the X-ray dataset. The left value is the accuracy of a standard classifier and the right value is the accuracy of EC-GAN.

For the chest X-ray dataset, we developed a conditional-GAN. Since this dataset has a strong class imbalance, a conditional GAN would prevent imbalance in the generated samples. The experimental results reported in Table 2 reflect the strong performance by our EC-GAN. Our proposed approach improves classification, especially in realistic and small datasets. Even in extreme scenarios, with just 200 labeled images, we achieved 90.9% accuracy.

Conclusion

We have presented EC-GAN, a novel generative model that attaches an external classifier to a GAN to improve classification performance in restricted, fully-supervised datasets. Our proposed method allows classifiers to leverage GAN image generations to improve classification, while simultaneously separating the tasks of discrimination and classification. Our results showed that EC-GAN is effective and can be used to improve image classification performance in small, real-world datasets.

References

Goodfellow, I. J.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative Adversarial Nets. In *Proceedings of the 27th International Conference on Neural Information Processing Systems - Volume 2*, 2672–2680.

Imran, A.-A.-Z.; and Terzopoulos, D. 2019. Multi-adversarial Variational Autoencoder Networks. In *2019 18th IEEE International Conference On Machine Learning And Applications (ICMLA)*, 777–782.

Kermany, D.; Zhang, K.; and Goldbaum, M. 2018. Labeled optical coherence tomography (OCT) and Chest X-Ray images for classification. *Mendeley data* 2.

Lee, D.-H. 2013. Pseudo-label: The simple and efficient semi-supervised learning method for deep neural networks. In *Workshop on challenges in representation learning, ICML*, volume 3.

Li, C.; Xu, T.; Zhu, J.; and Zhang, B. 2017. Triple generative adversarial nets. *Advances in neural information processing systems* 30: 4088–4098.

Radford, A.; Metz, L.; and Chintala, S. 2016. Unsupervised Representation Learning with Deep Convolutional Generative Adversarial Networks. In *4th International Conference on Learning Representations, ICLR*.

Salimans, T.; Goodfellow, I.; Zaremba, W.; Cheung, V.; Radford, A.; and Chen, X. 2016. Improved techniques for training gans. In *Advances in neural information processing systems*, 2234–2242.