Investigating the Potential of Auxiliary-Classifier GANs for Image Classification in Low Data Regimes

Amil Dravid[1], Florian Schiffers[1], Yunan Wu[1],
Oliver Cossairt[1], Aggelos Katsaggelos[1]
[1]Northwestern University

Northwestern University

icassp 2022 Singapore

Problem Definition and Contribution

- **Convolutional neural networks (CNNs)** rely on extremely **large datasets** to perform well on new data.
- We examine the potential for **Auxiliary-Classifier GANs (AC-GANs)** as a **’one-stop-shop’** architecture for image classification and generation, particularly in low data regimes.
- We propose **modifications** to the typical AC-GAN framework: **latent space sampling scheme** and **Wasserstein loss with gradient penalty**.

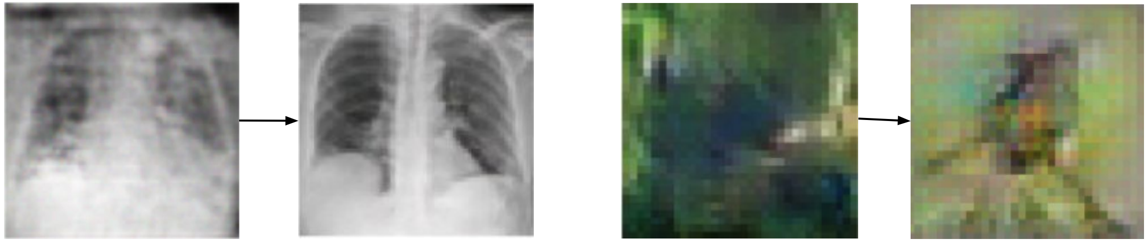

COVID-19 Chest X-ray

CIFAR10 Bird

Figure 1: Simple techniques can improve generated image quality even with limited data.

Generative Adversarial Networks (GANs)[1]

- Generator (G) tries to create samples to "fool" the discriminator (D).
- Discriminator takes turns looking at real (x) and fake images (G(*z*)).

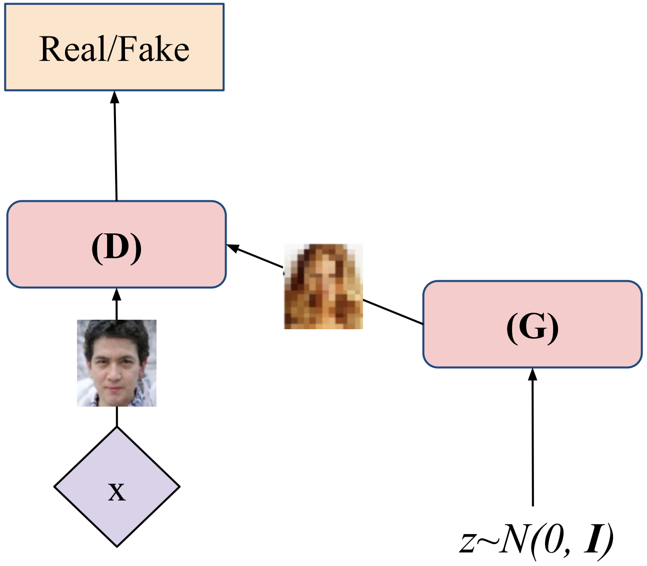

Figure 2: GAN Training Scheme.

Background: Auxiliary-Classifier GAN (AC-GAN)

- Auxiliary-Classifier GAN (AC-GAN) builds on the Conditional GAN (C-GAN) in order to improve image synthesis[2,3].

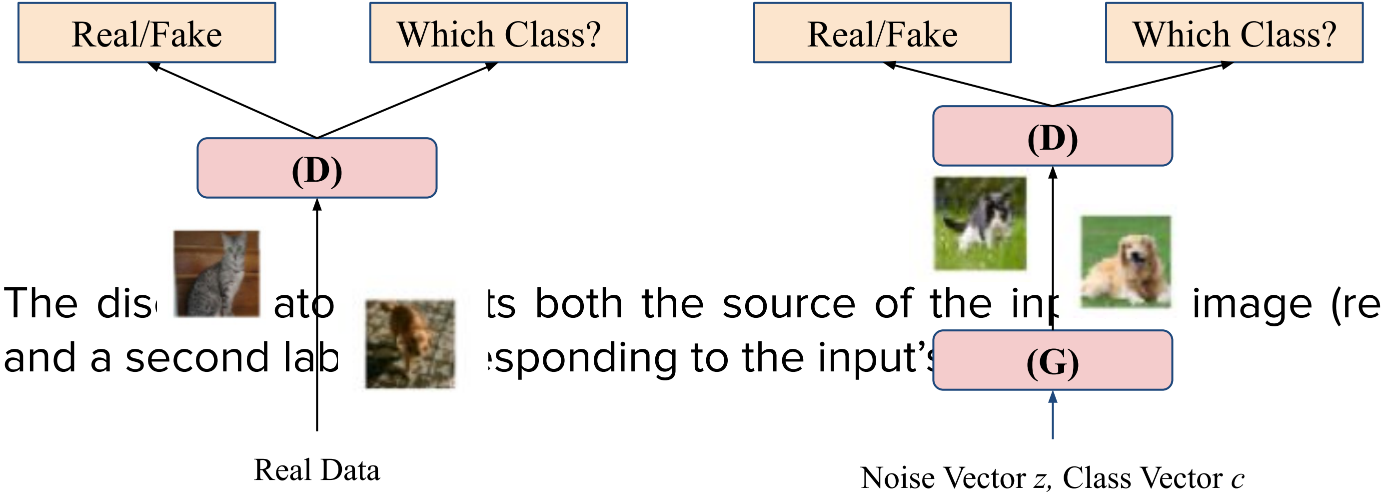

- The dis ato ts both the source of the inp image (real or fake) and a second lab esponding to the input's

Figure 3: AC-GAN Training Scheme.

Methods: Loss Function

- AC-GAN two-part objective:

$$L_S = \mathbb{E}[\log P(S = real|X_{real}))] + \mathbb{E}[\log P(S = fake|X_{fake})] \quad (1)$$

$$L_C = \mathbb{E}[\log P(C = c|X_{real}))] + \mathbb{E}[\log P(C = c|X_{fake})], \quad (2)$$

- Discriminator maximizes $L_S + L_C$. Generator maximizes $L_c - L_s$.

- L_s becomes **gradient-penalty Wasserstein loss**[4] to stabilize simultaneous image synthesis and classification.

$$L_S = \mathbb{E}[D(x)] - \mathbb{E}[D(G(z))] + \lambda\Phi \quad (3)$$

The Latent Space

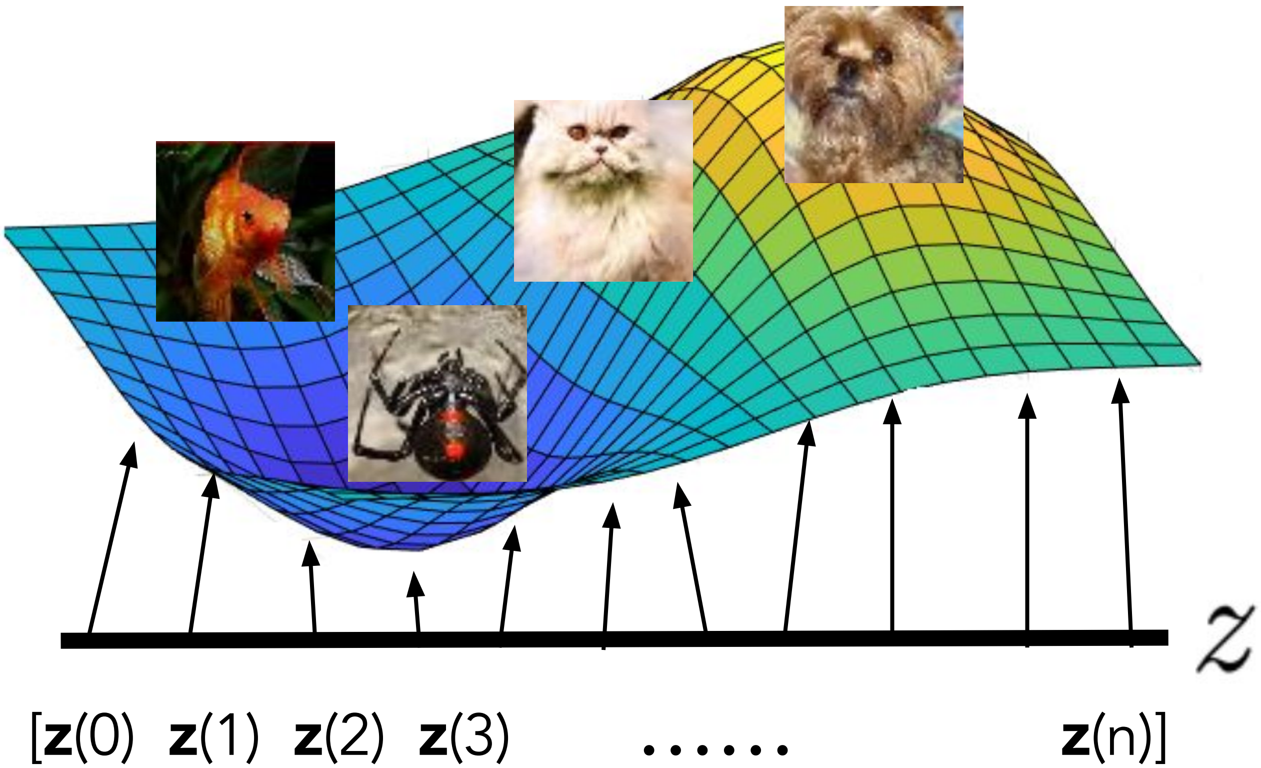

Figure 4: Illustrating the Latent Space.

Methods: Latent Sampling Scheme

- **Truncation trick**[5]: sample latent vector $\boldsymbol{z} \sim p_z$ closer to the mode of the distribution, resulting in images with greater realism, but low diversity.
- We propose **feeding truncated samples into the classifier**.

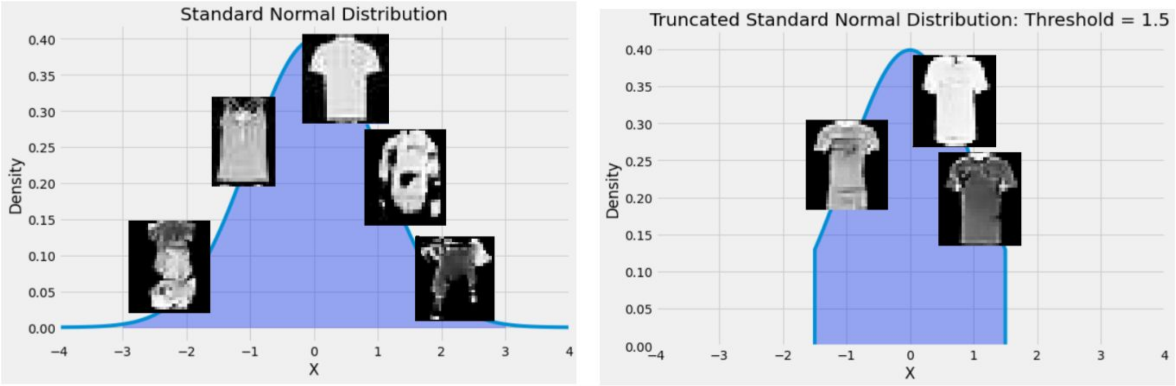

(a) Full distribution (b) Truncated distribution

Figure 5: Sampling the latent vector from a truncated distribution results in higher fidelity images, but lower diversity.

WAC-GAN-GPT

(Wasserstein AC-GAN with Gradient Penalty and Truncation).

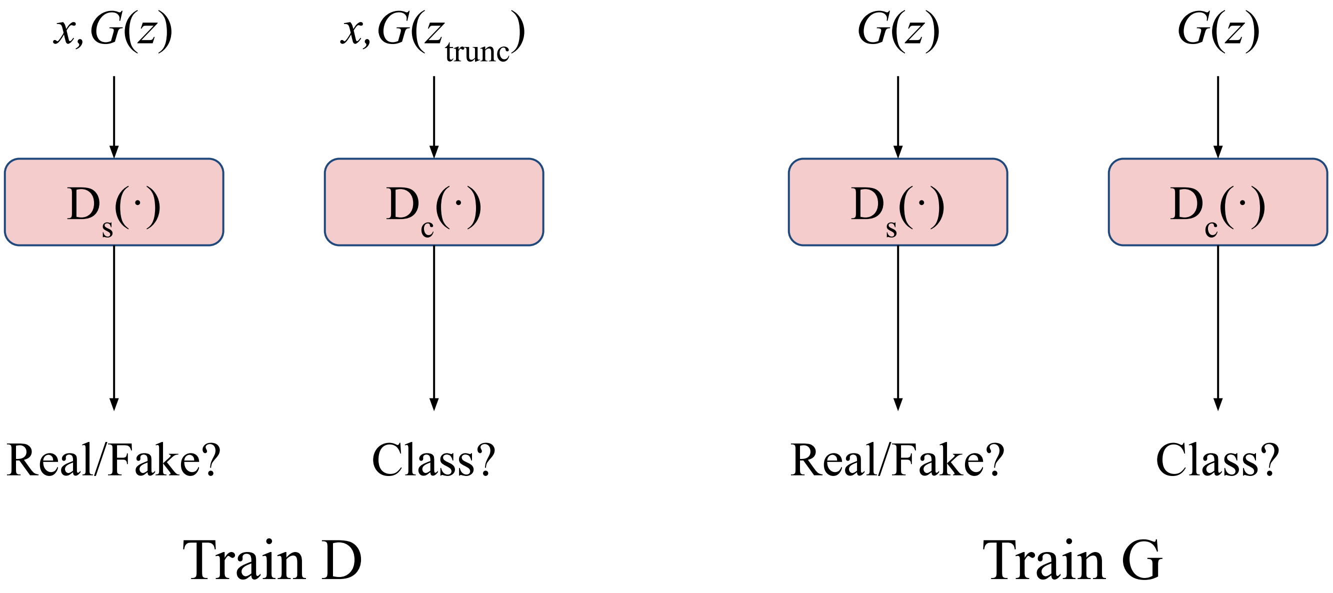

Figure 6: WAC-GAN-GPT Training Scheme

Results and Discussion

- Ablation studies on varying training set sizes on **Fashion MNIST** to compare test accuracy.

Train Size	Baseline CNN	AC-GAN	WAC-GAN-GP	AC-GAN with Truncation	WAC-GAN-GPT
500	$77.5\% \pm 1.5$	$77.6\% \pm 1.7$	$77.9\% \pm 1.5$	$78.8\% \pm 1.5$	$\mathbf{79.8}\% \pm 1.5$
2500	$83.5\% \pm 1.0$	$81.2\% \pm 2.1$	$84.4\% \pm 1.5$	$84.8\% \pm 1.1$	$\mathbf{86.0}\% \pm 1.2$
10000	$86.4\% \pm 1.5$	$87.3\% \pm 1.3$	$87.6\% \pm 0.9$	$87.8\% \pm 0.7$	$\mathbf{88.4}\% \pm 1.1$
20000	$87.5\% \pm 1.3$	$88.6\% \pm 1.6$	$88.1\% \pm 1.2$	$89.1\% \pm 0.5$	$\mathbf{89.8}\% \pm 0.9$
40000	$90.3\% \pm 0.8$	$90.9\% \pm 0.8$	$91.0\% \pm 0.4$	$90.7\% \pm 0.8$	$\mathbf{91.3}\% \pm 0.7$

Table 1: Performance on Fashion MNIST test set based on varying training set sizes.

1) Baseline CNN
2) AC-GAN
3) Wasserstein AC-GAN with Gradient Penalty (WAC-GAN-GP)
4) AC-GAN with Truncation
5) Our WAC-GAN-GPT

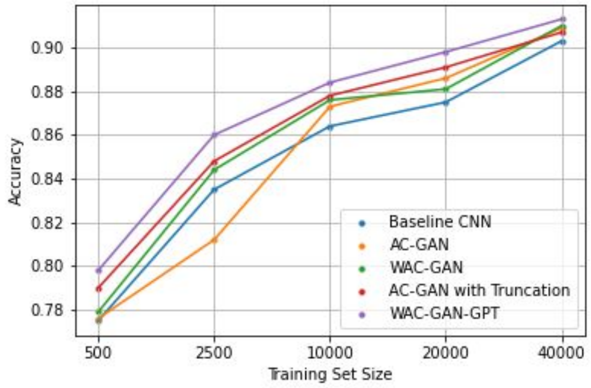

Figure 7: Graphical representation of Table 1.

Results and Discussion

- **T-SNE** on CNN embeddings for **real samples, AC-GAN samples**, and **WAC-GAN-GPT samples** based on **Fashion MNIST**.

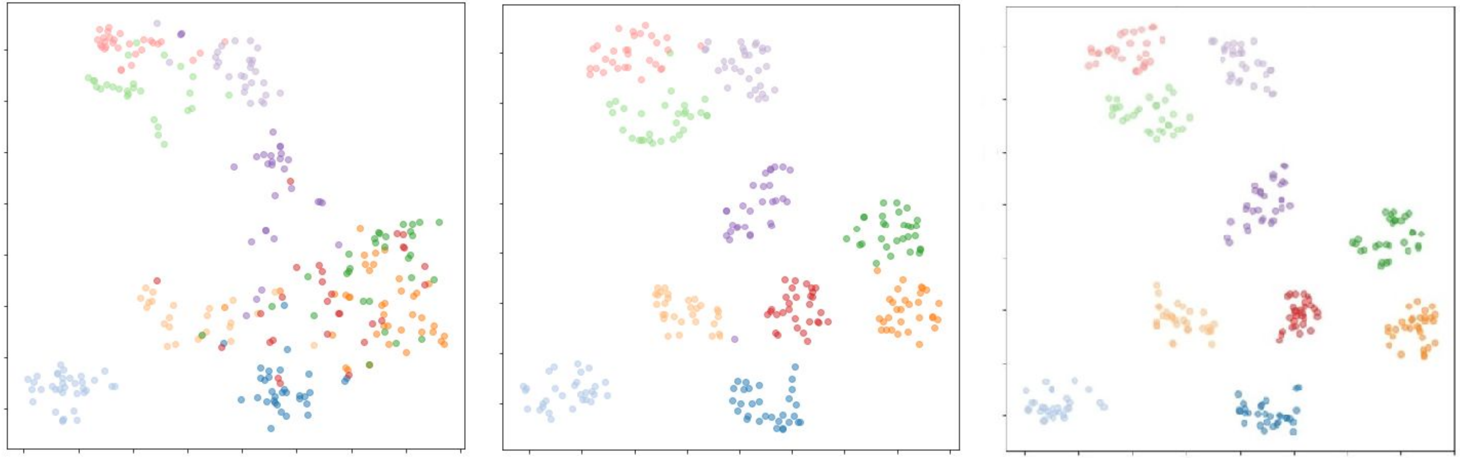

a) CNN b) AC-GAN c) WAC-GAN-GPT

Figure 8: T-SNE visualizations on Fashion MNIST samples.

- Average distance to center of class cluster: 7.83, 5.16, and 3.94 for the CNN, AC-GAN, and WAC-GAN-GPT, respectively.
- Standard deviations: 4.71, 2.17, and 1.76.

Results and Discussion

- Are low diversity but more representative images helpful?
- Find optimal truncation factor for each training set size experiment.
- Bilevel optimization: find optimal truncation factor τ while optimizing GAN:

$$\arg\min_{\tau} \mathrm{CE}(f(G(z_\tau, c)), c) \tag{4}$$

$$z \sim \mathcal{N}(0, \mathbf{I}) \tag{5}$$

$$z_\tau = sgn(z) \cdot min(|z|, \tau) \tag{6}$$

Training Size	500	2500	10000	20000	40000
τ	0.89 ± 0.14	1.05 ± 0.12	1.20 ± 0.13	1.49 ± 0.08	1.63 ± 0.05

Table 2: Optimal Truncation Factors for various training set sizes.

Results and Discussion

- Compare CIFAR10 test accuracy, and then CIFAR10.1v6[6] to compare domain generalizability.

	AlexNet	AC-GAN	WAC-GAN-GPT
CIFAR 10	$70.5\% \pm 0.5$	$70.1\% \pm 0.8$	$\mathbf{72.9}\% \pm 0.7$
CIFAR 10.1v6	$53.5\% \pm 1.0$	$56.4\% \pm 1.1$	$\mathbf{59.3}\% \pm 0.6$

Table 3: CIFAR test performance and generalizability.

- COVID-19 Detection on 128x128 chest X-rays.

	CNN	AC-GAN	WAC-GAN-GPT
COVID-19	$94.0\% \pm 1.5$	$95.5\% \pm 0.5$	$\mathbf{97.6}\% \pm 0.9$

Table 4: COVID-19 test performance.

Summary and Conclusion

- AC-GANs can achieve competitive performance with standard CNNs.
  - Particular performance gains in lower data regimes.
- Modifications: Wasserstein-GP + truncation.
- Future work: more diverse datasets, higher resolution images.
- More advanced techniques: adaptive discriminator augmentation or progressive growing.[7,8]

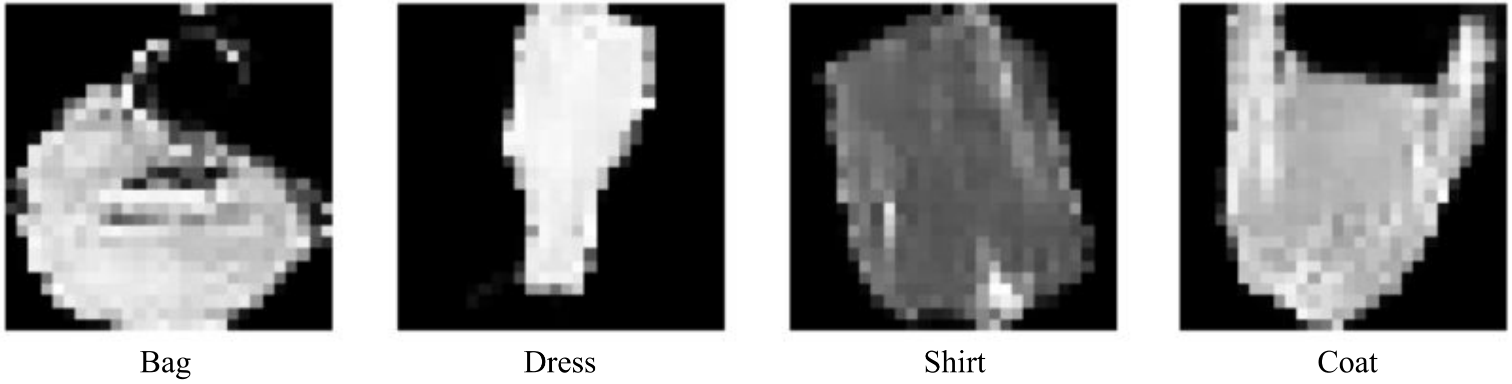

Figure 9: Naively applying data augmentation transformations leaks through to generated images.

References

[1] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio, “Generative adversarial nets,” in NeurIPS, vol. 27, pp. 2672–2680, 2014.

[2] Augustus Odena, Christopher Olah, and Jonathon Shlens, “Conditional image synthesis with auxiliary classifier gans,” in International Conference on Machine Learning. PMLR, 2017, pp. 2642–2651.

[3] Mehdi Mirza and Simon Osindero, “Conditional generative adversarial nets,” arXiv preprint arXiv:1411.1784, 2014.

[4] Ishaan Gulrajani, Faruk Ahmed, Mart́ın Arjovsky, Vincent Dumoulin, and Aaron C Courville, “Improved training of wasserstein gans,” in NeurIPS, 2017.

[5] Andrew Brock, Jeff Donahue, and Karen Simonyan, “large-scale gan training for high fidelity natural image synthesis,” in International Conference on Learning Representations, 2018.

[6] Benjamin Recht, Rebecca Roelofs, Ludwig Schmidt, and Vaishaal Shankar, “Do cifar-10 classifiers generalize to cifar10?,” arXiv preprint arXiv:1806.00451, 2018.

[7] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila, “Training generative adversarial networks with limited data,” arXiv preprint arXiv:2006.06676, 2020.

[8] Tero Karras, Timo Aila, Samuli Laine, and Jaakko Lehtinen, “Progressive growing of gans for improved quality, stability, and variation,” arXiv preprint arXiv:1710.10196, 2017.

Contact: amildravid2023@u.northwestern.edu

IEEE Xplore Link: https://ieeexplore.ieee.org/document/9747286

arXiV Link: https://arxiv.org/abs/2201.09120

Code: https://github.com/avdravid/AC-GANS-FOR-IMAGE-CLASSIFICATION