Predicting Gaussian noise injection for self-supervised Generative Adversarial Nets

Zhiyuan Wu[1], Grigorios G Chrysos[2], Volkan Cevher[2]
[1] Technical University of Munich, Germany
[2] EPFL, Switzerland

Motivation

- Self-supervised learning can encourage the discriminator to learn meaningful feature representations.
- However, the performance of pretext tasks can depend heavily on the dataset. For instance, the rotation-based pretext task [1] might fail to recognize the rotation degree of flowers.
- Motivated by [3, 4], we propose our pretext task: Insert Gaussian noise in the blocks of the discriminator and predict the position.
- Our learned representations lead to a more balanced classification accuracy across classes.

Method

- Replace the convolutional discriminator with a polynomial network from the Π-net family [2]. Suppose there are N blocks in the discriminator with the output of the t^{th} block being $z_{t+1} = z_t + Cz_t + (Cz_t) * z_t$.
- Randomly choose one block out of N uniformly and multiply the second order term by Gaussian noise.
- Suppose t is the block we insert Gaussian noise, its output will be: $z_{t+1} = z_t + Cz_t + (Cz_t) * z_t * \rho$, where $\rho \sim \mathcal{N}(\mu, \delta^2)$.
- Suppose t is not the block we apply Gaussian noise, the output of this block remains the same, i.e. $z_{t+1} = z_t + Cz_t + (Cz_t) * z_t$.
- A classifier Q, which shares all the weights except for the last layer with the discriminator D, predicts in which block the noise was injected.
- The symbol '$*$' refers to an elementwise product.
- After the training, utilize the discriminator for a classification task to evaluate the representation learning performance of the pretext task.

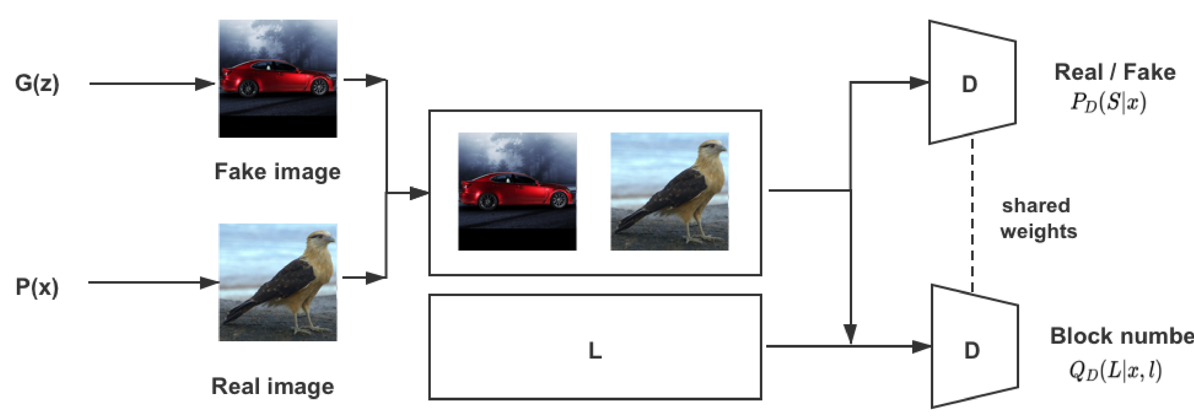

Fig. 1: Diagram of our Self-Supervised Poly-GAN (P-SS-SNGAN(N)).

Quantitative Results

- We train our model on three widely used datasets: CIFAR10, STL10, COIL20.
- For our quantitative evaluation we use the FID score.

Method	CIFAR10	STL10	COIL20
SNGAN	23.680	62.381	94.491
P-SNGAN	19.262	61.667	92.421
SNGAN(R)	21.668	61.142	92.257
P-SS-SNGAN (R)	**17.166**	**59.664**	**91.321**
P-SS-SNGAN (N)	17.886	60.163	91.782

Representation Quality

- After the training of GANs, utilize the discriminator for a classification task to assess the representation quality.
- The accuracy of the classifier is considered to indicate quality of learned representations.

Method	CIFAR10	STL10	COIL20
SNGAN	58.28%	56.84%	90.00%
P-SNGAN	59.76%	57.64%	92.67%
SNGAN(R)	60.17%	57.29%	94.00%
P-SS-SNGAN (R)	**64.17**%	**58.12**%	**96.67**%
P-SS-SNGAN (N)	63.81%	57.85%	95.67%

Performance Across Classes

- Our method results in a more balanced classification accuracy across different classes, namely larger improvements for weaker classes.

Method	CIFAR10			STL10			COIL20	
	3	6	σ^2	4	6	σ^2	2	σ^2
P-SS-SNGAN(R)	0.446	0.422	0.011	0.348	0.184	0.027	0.333	0.029
P-SS-SNGAN(N)	**0.504**	**0.504**	0.011	**0.490**	**0.282**	**0.015**	**0.920**	**0.019**

Acknowledgement

Research was sponsored by the Army Research Office and was accomplished under Grant Number W911NF-19-1-0404. This project has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (grant agreement n° 725594 - time-data).

For more details in our paper, please access through the following link or scan the QR code: https://bit.ly/3EG1GHc

References

[1] Ting Chen et al. "Self-Supervised GANs via Auxiliary Rotation Loss". In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)* (2019), pp. 12146–12155.

[2] Grigorios Chrysos et al. "Π−nets: Deep Polynomial Neural Networks". In: *IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*. 2020.

[3] C. Sønderby et al. "Amortised MAP Inference for Image Super-resolution". In: *International Conference on Learning Representations*. 2017. URL: https://arxiv.org/abs/1610.04490.

[4] Shengyu Zhao et al. "Differentiable Augmentation for Data-Efficient GAN Training". In: *Advances in Neural Information Processing Systems 33: Annual Conference on Neural Information Processing Systems 2020, NeurIPS 2020, December 6-12, 2020, virtual*. Ed. by Hugo Larochelle et al. 2020. URL: https://proceedings.neurips.cc/paper/2020/hash/55479c55ebd1efd3ff125f1337100388-Abstract.html.