Compressing Image-to-Image Translation GANs Using Local Density Structures on Their Learned Manifold

Alireza Ganjdanesh^{*1}, Shangqian Gao^{*2}, Hirad Alipanah³, Heng Huang¹

 ¹ Department of Computer Science, University of Maryland College Park, College Park, MD 20742, USA
 ² Department of Electrical and Computer Engineering, University of Pittsburgh, Pittsburgh, PA 15261, USA
 ³ Department of Mechanical Engineering and Materials Science, University of Pittsburgh, Pittsburgh, PA 15261, USA aliganj@umd.edu, shg84@pitt.edu, hia21@pitt.edu, heng@umd.edu

Abstract

Generative Adversarial Networks (GANs) have shown remarkable success in modeling complex data distributions for image-to-image translation. Still, their high computational demands prohibit their deployment in practical scenarios like edge devices. Existing GAN compression methods mainly rely on knowledge distillation or convolutional classifiers' pruning techniques. Thus, they neglect the critical characteristic of GANs: their local density structure over their learned manifold. Accordingly, we approach GAN compression from a new perspective by explicitly encouraging the pruned model to preserve the density structure of the original parameterheavy model on its learned manifold. We facilitate this objective for the pruned model by partitioning the learned manifold of the original generator into local neighborhoods around its generated samples. Then, we propose a novel pruning objective to regularize the pruned model to preserve the local density structure over each neighborhood, resembling the kernel density estimation method. Also, we develop a collaborative pruning scheme in which the discriminator and generator are pruned by two pruning agents. We design the agents to capture interactions between the generator and discriminator by exchanging their peer's feedback when determining corresponding models' architectures. Thanks to such a design, our pruning method can efficiently find performant sub-networks and can maintain the balance between the generator and discriminator more effectively compared to baselines during pruning, thereby showing more stable pruning dynamics. Our experiments on image translation GAN models, Pix2Pix and CycleGAN, with various benchmark datasets and architectures demonstrate our method's effectiveness.

Introduction

Image-to-Image translation (Isola et al. 2017; Zhu et al. 2017) Generative Adversarial Networks (I2IGANs) (Good-fellow et al. 2014) have shown excellent performance in many real-world computer vision applications: style transfer (Huang and Belongie 2017), converting a user's sketch to a real image (Park et al. 2019), super resolution (Ledig et al. 2017; Wang et al. 2018b), and pose transfer (Wang et al. 2018a; Chan et al. 2019). Yet, I2IGANs require excessive compute and memory resources. Moreover, the men-

tioned tasks require real-time user interaction, making it infeasible to deploy I2IGANs on mobile and edge devices in Artificial Intelligence of Things (AIoT) with limited resources. Thus, developing compression schemes for GANs to preserve their performance and reduce their computational burden is highly desirable. As the training dynamics of GAN models are notoriously unstable, GAN compression is much more challenging than pruning other deep models like Convolutional Neural Network (CNN) classifiers.

Despite that notable efforts have been made to compress CNNs (Han et al. 2015; Sandler et al. 2018; Rastegari et al. 2016; Li et al. 2017; Ye et al. 2020; Wan et al. 2020), GAN compression has only been explored in recent years. Early works have proposed a combination of prominent CNN pruning techniques like Neural Architecture Search (NAS) (Gong et al. 2019; Li et al. 2020; Jin et al. 2021; Gao et al. 2020; Li et al. 2022), Knowledge Distillation (Aguinaldo et al. 2019; Wang et al. 2020; Chang and Lu 2020; Chen et al. 2020a; Fu et al. 2020; Hou et al. 2021; Zhang et al. 2022b,a), and channel pruning (Li et al. 2017, 2022) to prune GANs. However, GAN Slimming (Wang et al. 2020) demonstrated that heuristically stacking several CNN pruning methods for GAN compression can degrade the final performance mainly due to the instabilities of GAN training. GCC (Li et al. 2021) empirically showed the importance of restricting the discriminator's capacity during compression. It demonstrated that the previous methods' unsatisfactory performance might be because they only pruned the generator's architecture while using the full-capacity discriminator. By doing so, the adversarial game cannot maintain the Nash Equilibrium state, and the pruning process fails to converge appropriately. Although some of these methods have shown competitive results (Li et al. 2021; Jin et al. 2021), they do not explicitly consider an essential characteristic of GANs as generative models during pruning, which is their density structure over their learned manifold.

In this paper, we propose a novel GAN Compression method by enforcing the similarity of the density structure of the original parameter-heavy model and the pruned model over the learned manifold of the original model. Our intuition is that the difference in density structures can serve as the supervision signal for pruning. Specifically, at first, we partition the learned manifold of the original model into local neighborhoods. We approximate each neighborhood

^{*}These authors contributed equally.

Copyright © 2024, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

with a generated sample and its nearest neighbors on the original model's manifold. We leverage a pretrained selfsupervised model fine-tuned on the training dataset to find the neighborhoods. Then, we introduce an adversarial pruning objective to encourage the pruned model to have a similar local density structure to the original model on each neighborhood. By doing so, we break down the task of preserving the whole density structure of the original model on its learned manifold into maintaining local density structures on neighborhoods of its manifold, which resembles kernel density estimation (Parzen 1962). In addition, we design a new adversarial GAN compression scheme in which two pruning agents (we call them gen_G and gen_D , which determine the structure of the generator G and discriminator D) collaboratively play our proposed adversarial game. Specifically, each agent takes the architecture embedding of its colleague as its input to determine the structure of its corresponding model in each iteration. By doing so, gen_G and gen_D will be able to effectively preserve the balance between the capacities of G and D and keep the adversarial game close to the Nash Equilibrium state during the pruning process. We summarize our contributions as follows:

- We propose a novel GAN compression method that encourages the pruned model to have a similar local density structure as the original model on neighborhoods of the original model's learned manifold.
- We design two pruning agents that collaboratively play our adversarial pruning game to compress both the generator and discriminator together. By doing so, our method can effectively preserve the balance between the capacities of the generator and discriminator and show more stable pruning dynamics while outperforming baselines.
- Our extensive experiments on Pix2Pix (Isola et al. 2017) and CycleGAN (Zhu et al. 2017) on various datasets demonstrate our method's effectiveness.

Related Work

GAN Compression: GANs require two orders of magnitude more computation than CNNs (Li et al. 2020). Hence, GAN compression is crucial prior to deploying them on edge devices. Search-based methods (Shu et al. 2019; Li et al. 2020; Lin et al. 2021; Wang et al. 2021) search for a lightweight architecture for the generator but are extremely costly due to their vast search space. Pruning methods (Li et al. 2020; Jin et al. 2021; Wang et al. 2020; Yu and Pool 2020) prune the redundant weights of the generator's architecture but neglect the discriminator. They result in an unbalanced generator and discriminator capacities, leading to mode collapse (Li et al. 2021). To address this problem, discriminator-free methods (Ren et al. 2021; Fu et al. 2020) distill the generator into a compressed model without using the discriminator. In another direction, GCC (Li et al. 2021) and Slimmable GAN (Hou et al. 2021) prune both the generator and discriminator together. Slimmable GAN sets the discriminator's layers' width proportional to the ones for the generator during pruning. Yet, GCC empirically showed there is no linear relation between the number of channels

of the generator and optimal discriminator, and Slimmable GAN's approach is sub-optimal. Inspired by GCC, we use two pruning agents that learn to determine the architectures of the generator and discriminator in our proposed adversarial game. Each agent gets feedback from its peer when determining its corresponding model's architecture. Thus, they can effectively preserve the balance between the generator and discriminator and stabilize the pruning process.

Manifold Learning for GANs: The manifold hypothesis indicates that high-dimensional data like natural images lie on a nonlinear manifold with much smaller intrinsic dimensionality (Tenenbaum, Silva, and Langford 2000). Accordingly, a group of methods alter the training (Casanova et al. 2021) and/or architecture of GANs by using several generators (Khayatkhoei, Singh, and Elgammal 2018), including a manifold learning step in the discriminator (Ni et al. 2022), and local coordinate coding based sampling (Cao et al. 2018). Yet, one cannot use these methods for pruning GANs that are pretrained with other methods. Another group of ideas observed that semantically meaningful paths and neighborhoods exist in the latent space of trained GANs. (Oldfield et al. 2021; Härkönen et al. 2020). Inspired by these methods, we propose to partition the learned manifold of a pretrained GAN into overlapping neighborhoods. Then, we develop an adversarial pruning scheme that encourages the pruned model to have a similar density structure to the original one over each neighborhood.

Network Pruning: Model compression (Ghimire, Kil, and Kim 2022) is a well-studied topic, and proposed methods have primarily focused on compressing CNN classifiers. They use various techniques like weight pruning (Han et al. 2015), light architecture design (Tan and Le 2019; Howard et al. 2017), weight quantization (Rastegari et al. 2016), structural pruning (Li et al. 2017; Ye et al. 2020; Ganjdanesh, Gao, and Huang 2022; He et al. 2018), knowledge distillation (Ba and Caruana 2014), and NAS (Wu et al. 2019; Ganjdanesh, Gao, and Huang 2023). We focus on pruning GANs, which is a more challenging task due to the instability of their training (Wang et al. 2020).

Proposed Method

We develop a new GAN compression method that explicitly regularizes the pruned model not to forget the density structure of the original one over its learned manifold. However, directly applying such regularization along with compression objectives can make the pruning process unstable because of the complex nature of training GANs. Thus, we simplify this objective for the pruned generator in two steps. At first, we propose to partition the learned manifold of the original model into local neighborhoods, each consisting of a sample and its nearest neighbors on the manifold. We employ a self-supervised model fine-tuned on the training dataset to approximately find such neighborhoods. Then, we propose an adversarial pruning objective to enforce the pruned model to have a density structure similar to the original model over each local neighborhood. Finally, we introduce a novel GAN compression scheme in which we use two pruning agents - called gen_G The Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI-24)



Figure 1: Our GAN pruning method. We encourage the pruned generator to preserve the density structure of the original model over its learned manifold during pruning. To do so, we partition the manifold into local neighborhoods around the samples generated by the original generator (Fig. 2) and represent each local neighborhood with a 'Center' sample (red frame) and its neighbors (blue frames). We use adversarial training to match the distributions of the original and pruned models in the neighborhoods. We implement our adversarial pruning with two pruning agents, gen_G and gen_D , that collaboratively learn to prune the original pretrained G and D by exchanging feedback.

and gen_D - that determine the architectures of the generator G and discriminator D, respectively. gen_G and gen_D play the adversarial game introduced in the previous step in a collaborative manner to find the optimal structure of Gand D. In each step, $gen_G(gen_D)$ gets feedback from its peer $gen_D(gen_G)$ about the architecture of D(G) when determining the architecture of G(D). By doing so, they can maintain the balance between the generator and discriminator during pruning and stabilize the pruning process. We show our pruning scheme in Fig. 1.

Notations

We denote an I2IGAN model's source and target domains with \mathcal{X} and \mathcal{Y} , respectively. We show the training dataset as $\mathcal{D} = \{\{(x_i)\}_{i=1}^N, \{(y_j)\}_{j=1}^M\}$ such that $(x, y) \sim \mathcal{P}(\mathbf{x}, \mathbf{y})$ where \mathcal{P} is the underlying joint distribution over the source and target domains. N and M can be equal for the paired datasets (Isola et al. 2017) or be unequal for unpaired models (Zhu et al. 2017). We represent the generator and discriminator with G and D. Also, we denote pruning agents determining the architectures of G and D during pruning with $gen_G(\cdot)$ and $gen_D(\cdot)$. The goal of the generator is to learn the distribution of corresponding $y \in \mathcal{Y}$ in the target domain conditioned on a sample $x \in \mathcal{X}$ from the source domain. We denote the learned manifold of the original generator in the domain \mathcal{Y} with \mathcal{M}_y .

Finding Local Neighborhoods on the Learned Data Manifold of the Original Generator

As mentioned above, our idea is to guide the pruning process of a GAN model by regularizing the pruned generator to have a similar density structure to the original model over its learned manifold, \mathcal{M}_y . To simplify this objective for the pruned model, we partition \mathcal{M}_y into local neighborhoods \mathcal{N}_{y_i} containing a sample y_i and its nearest neighbors on the manifold \mathcal{M}_y . The intuition is that separately modeling the density structure over each neighborhood is easier than modeling all of them simultaneously, which resembles the kernel density estimation method (Parzen 1962).

To find the local neighborhoods, on the one hand, we get inspirations from recent works that observed that latent spaces of GAN models have semantically meaningful paths and neighborhoods (Tzelepis, Tzimiropoulos, and Patras 2021; Shen and Zhou 2021; Choi et al. 2022; Voynov and Babenko 2020). On the other hand, self-supervised pretrained encoders (Chen et al. 2020b,c; Grill et al. 2020) have shown significant results in finding semantically similar samples without using labels on the manifold of their input data. Accordingly, at first, we fine-tune an encoder pretrained by SwAV (Caron et al. 2020) on the samples $\{(y_j)\}_{j=1}^M$ in the training dataset. Then, we use the training dataset and pass the training samples x_i into the original generator to obtain its predicted samples $\mathcal{D}'_y = \{y'_i\}_{i=1}^N$ on \mathcal{M}_{y} . Finally, we approximately model the neighborhood of y'_i over \mathcal{M}_y by finding its k nearest neighbor samples in \mathcal{D}'_y using our fine-tuned encoder. Formally, given a fine-tuned self-supervised encoder E, we find k nearest neighbors of y'_i in $\overline{\mathcal{D}}'_y$, denoted by $\mathcal{N}_{y_i,k}$, as follows: (Fig. 2)

$$\mathcal{E}' = \{e'_j\}_{j=1}^N, \ e'_j = E(y'_j)$$

$$Sim_i = \{Cosine(e'_i, e'_j)\}_{j=1, \ j \neq i}^N$$

$$Cosine(e'_i, e'_j) = ||e'_i{}^T e'_j||/(||e'_i|||e'_j||)$$

$$\mathcal{N}_{y'_i,k} = \{y'_j|e'_j \in \text{Top-}k(Sim_i)\}$$

(1)

i.e., we take samples in \mathcal{D}'_y that their encoded representations have the highest cosine similarity with the one for y'_i as its neighbors on \mathcal{M}_y . We use the cosine similarity metric as it has been widely used in the nearest neighbor classifiers (Vinyals et al. 2016; Touvron et al. 2021) and self-supervised learning (Caron et al. 2020).



Figure 2: Our method to find local neighborhoods on the learned manifold of the original generator. First, we obtain the original model's predictions in the target domain for training samples in the source domain (Top Left). We call the sample that we want to find its local neighborhood on the manifold 'Center' sample (red frame). We pass the predicted samples to a self-supervised encoder that is fine-tuned on the target images in the training dataset. Then, we take the samples whose representations have the highest cosine similarity with the representation of the 'Center' sample as its approximate neighbors on the manifold.

Pruning

In this section, first, we elaborate on our pruning objective. Then, we introduce our pruning agents.

Pruning Objective: We regularize the pruned generator to have a similar density structure to the original one over each local neighborhood of the learned manifold of the original model (\mathcal{M}_y). Formally, we approximately represent the neighborhood around each sample y'_i with samples in $\mathcal{N}_{y'_i,k}$ in Eq. 1. Then, we define our adversarial training objective to regularize the pruned model to preserve the local density structure of the original model in each neighborhood:

$$\min_{\theta_G} \max_{\theta_D} E_{(x,y')\sim\mathcal{P}'(\mathbf{x},\mathbf{y})} [E_{y''\sim\mathcal{N}_{y'}} [f_D(D(x,y'';v_D))]] + E_{x\sim\mathcal{P}(\mathbf{x})} E_{\xi} [f_G(D(x,G(x,\xi;v_G);v_D))] + \lambda_1 \mathcal{R}(v_G) - \lambda_2 \mathcal{L}(v_D)$$
(2)

 $\mathcal{P}'(y)$ is the marginal distribution of the original generator model on \mathcal{Y} . G and D are the pretrained generator and discriminator. θ_G and θ_D are parameters of gen_G and gen_D . v_G and v_D are architecture vectors that determine the structures of G and D. They are functions of θ_G and θ_D . f_D and f_G are GAN objectives that are least squares for CycleGAN (Zhu et al. 2017) and hinge loss for Pix2Pix (Isola et al. 2017). ξ represents the randomness in the generator implemented with Dropout (Srivastava et al. 2014) for Pix2Pix and Cycle GAN. \mathcal{R} is the regularization term to enforce the desired compression ratio on the generator, and \mathcal{L} imposes sparsity on the architecture of the discriminator. Penalizing unimportant components for discriminator is crucial to maintain the capacity balance between G and D during pruning and keeping them close to the Nash Equilibrium state, as pointed out by GCC (Li et al. 2021). Moreover, our objective is similar to the Kernel Density Estimation (KDE) (Parzen 1962) method, which aims to model the local density around each data point. Yet, in contrast with KDE, we use an implicit objective to encourage the pruned model to have a similar

density structure to the original one in our method. In addition, in Obj. 2, the parameters of G and D are inherited from the original pretrained models, and we do not train parameters of G/D along with θ_G/θ_D to prevent instability. We note that our pruning objective does not require paired images during pruning, even for GAN models such that use a paired dataset for training. The reason is that our objective employs the samples generated by the original generator to approximate its density structure over each local neighborhood. Thus, it can readily prune both paired (Isola et al. 2017) and unpaired conditional GANs (Zhu et al. 2017).

Pruning Agents: Inspired by GCC (Li et al. 2021) that demonstrated that preserving the balance between the capacity of G and D is crucial for pruning stability, we prune both G and D during the pruning process. To do so, we introduce a novel GAN compression scheme in which we use two pruning agents, gen_G and gen_D , to predict 1) architecture vectors and 2) architecture embeddings for G and D. The former is a binary vector determining the architecture of the corresponding model (G/D), and the latter is a compact representation describing its state (architecture). To preserve the balance between the capacity of G and D, we design our pruning scheme such that each pruning agent gen_G/gen_D considers the architecture embedding of D/G when determining the architecture of G/D. We implement pruning agents with a GRU (Cho et al. 2014) model and dense layers (more details in supplementary materials) and take the outputs of the GRU units of the pruning agents gen_G/gen_D as their corresponding model's (G/D) architecture embedding, summarizing the information about its architecture.

In each step of the adversarial game, gen_G and gen_D exchange their architecture embeddings. Then, each of them determines its corresponding model's architecture while knowing the other model's structure (architecture embedding), making the pruning process stable and efficient. (Fig. 1) We denote the architecture vector determining

the architecture of G/D with $v_G/v_D \in \{0, 1\}$. As these vectors have discrete values, their gradients *w.r.t* agents' parameters cannot be calculated directly. Thus, we use Straight-through Estimator (STE) (Bengio, Léonard, and Courville 2013) and Gumbel-Sigmoid reparametrization trick (Jang, Gu, and Poole 2017) to calculate the gradients. The subnetwork vector v_G is calculated as:

$$v_G = \text{round}(\text{sigmoid}(-(o_G + g)/\tau)),$$

$$o_G, h_G = gen_G(h_D; \theta_G)$$
(3)

where round(\cdot) rounds input to its nearest integer, sigmoid(\cdot) is the sigmoid function, τ is the temperature parameter to control the smoothness, and $g \in \text{Gumbel}(0,1)$ is a random vector sampled from the Gumbel distribution (Gumbel 1954). h_D is the architecture embedding for D, which is the input for gen_G . o_G and h_G are the output of gen_G and its architecture embedding for G. Similarly, v_D is calculated as:

$$v_D = \text{round}(\text{sigmoid}(-(o_D + g)/\tau)),$$

$$o_D, h_D = qen_D(h_G; \theta_D)$$
(4)

The calculation from o_* to v_* ($* \in \{G, D\}$) can be seen as using straight-through Gumbel-Sigmoid (Jang, Gu, and Poole 2017) to approximate sampling from the Bernoulli distribution. We provide our pruning algorithm and details of the calculation of o_* and h_* in the supplementary.

Predicted architecture vectors v_G and v_D determine the architectures of G and D. For the generator G, we aim to reduce MACs to reach a given budget. To do so, we use the following regularization objective:

$$\mathcal{R}(v_G) = \log(\max(T(v_G), pT_{\text{total}})/pT_{\text{total}}), \quad (5)$$

where p is the predefined threshold for pruning, $T(v_G)$ is the MACs of the current sub-network chosen by v_G , and T_{total} is the total prunable MACs.

Different from G, it is not obvious how to set a specific computation budget for D, as shown in GCC (Li et al. 2021). Instead of using a predefined threshold, we encourage gen_D to automatically identify the sub-network that can keep the Nash Equilibrium given the v_G . Inspired by the functional modularity (Csordás, van Steenkiste, and Schmidhuber 2021), we add a penalty to v_D to discover the suitable sub-module (sub-network) to keep the Nash Equilibrium:

$$\mathcal{L}(v_D) = \sum v_D / |v_D|, \tag{6}$$

where $|v_D|$ is the number of elements in v_D . The goal of Eq. 6 is to penalize unimportant elements in v_D , so the remaining elements can maintain Nash Equilibrium given v_G . **Finetuning:** After the pruning stage, we employ the trained gen_G/gen_D from the pruning process to prune G/D according to their predictions, v_G/v_D . Then, we finetune G and D together with the original objectives of their GAN methods (Isola et al. 2017; Zhu et al. 2017). Similar to GCC (Li et al. 2021), we also apply knowledge distillation (KD) for finetuning, but we show in our ablation experiments that our model can achieve high performance even without KD.



Figure 3: Qualitative results for 1) Pix2Pix: Cityscapes (top left), Edges2Shoes (bottom left), and 2) CycleGAN: Horse2Zebra (top right) and Summer2Winter (bottom right).

Experiments

We perform our experiments with Pix2Pix and Cycle-GAN models. For all experiments, we set $\lambda_1 = 3.0$, $\lambda_2 = 0.1$, and the number of neighbor samples k = 5 during pruning. We also find that our model is not very sensitive to the choice of λ_1 and λ_2 . (more details in ablation experiments) We set the number of pruning epochs to 10% of the original model's training epochs.

Results

Comparison with State of the Art Methods: We summarize the quantitative results of our method and baselines in Tab. 1. As can be seen, our method, MGGC (Manifold Guided GAN Compression), can achieve the best performance vs. computation rate trade-off compared to baselines in all experiments. For Pix2Pix on Cityscapes, MGGC reduces MACs by 83.60%, achieving the highest MACs compression rate, and improves mIoU by 1.92 compared to the original model, significantly outperforming baselines. On Edges2Shoes, MGGC prunes 0.27% more MACs than DMAD (Li et al. 2022) while outperforming it with a large margin of 4.93 FID. Although WKD (Zhang et al. 2022b) and RAKD (Zhang et al. 2022a) achieve higher compression ratio, their final model has significantly worse FID. For CycleGAN on Horse2Zebra, MGGC shows a pruning ratio very close (95.60% vs. 95.77%) to GCC (Li et al. 2021) and accomplishes 55.06 FID, significantly outperforming GCC by 4.25 FID. On Summer2Winter, MGGC attains 94.77% MACs compression rate, slightly more than DMAD with 94.40%, and shows 2.39 less FID. Remarkably, it can outperform other baselines even with an extreme MACs compression rate of 97.02%, and yet, reaching 77.33 FID. In summary, our results demonstrate the effectiveness of our method that explicitly focuses on the differences between the density structure of the pruned model and the original one over its learned manifold during pruning.

The Thirty-Eighth AAAI	Conference on Artificial	Intelligence (AAAI-24)
2 0		0 1

Model	Dataset	Method	MACs	Compression	Metric	
					$FID(\downarrow)$	mIoU (†)
	Cityscapes	Original (Isola et al. 2017)	18.60 G	-	-	42.71
		GAN Compression (Li et al. 2020)	5.66 G	69.57%	-	40.77
		CF-GAN (Wang et al. 2021)	5.62 G	69.78%	-	42.24
		CAT (Jin et al. 2021)	5.57 G	70.05%	-	42.53
		DMAD (Li et al. 2022)	3.96 G	78.71%	-	40.53
Div 2Div		WKD (Zhang et al. 2022b)	3.88 G	79.13%	-	42.93
1 1/21 1/		RAKD (Zhang et al. 2022a)	3.88 G	79.13%	-	42.81
		Norm Pruning (Li et al. 2017)	3.09 G	83.39%	-	38.12
		GCC (Li et al. 2021)	3.09 G	83.39%	-	42.88
		MGGC (Ours)	$3.05~\mathrm{G}$	83.60%	-	44.63
		Original (Isola et al. 2017)	18.60 G	-	34.31	-
	Edges2Shoes	Pix2Pix $0.5 \times$ (Isola et al. 2017)	4.65 G	75.00%	52.02	-
	Eugeszonoes	CIL (Kim, Choi, and Park 2022)	4.57 G	75.43%	44.40	-
		DMAD (Li et al. 2022)	2.99 G	83.92%	46.95	-
		WKD (Zhang et al. 2022b)	$1.56~\mathrm{G}$	91.61%	80.13	-
		RAKD (Zhang et al. 2022a)	1.56 G	91.61%	77.69	-
		MGGC (Ours)	2.94 G	84.19%	42.02	-
		Original (Zhu et al. 2017)	56.80 G	-	61.53	-
		Co-Evolution (Shu et al. 2019)	13.40 G	76.41%	96.15	-
Model Pix2Pix		GAN Slimming (Wang et al. 2020)	11.25 G	80.19%	86.09	-
	Horse2Zebra	GAN Compression (Li et al. 2020)	2.67 G	95.30%	64.95	-
		CF-GAN (Wang et al. 2021)	2.65 G	95.33%	62.31	-
		CAT (Jin et al. 2021)	2.55 G	95.51%	60.18	-
		DMAD (Li et al. 2022)	2.41 G	95.76%	62.96	-
		Norm Prune (Li et al. 2017)	2.40 G	95.77%	145.1	-
		GCC (Li et al. 2021)	2.40 G	95.77%	59.31	-
		MGGC (Ours)	$2.50~\mathrm{G}$	95.60%	55.06	-
		Original (Zhu et al. 2017)	56.80 G	-	79.12	-
	Summer2Winter	Co-Evolution (Shu et al. 2019)	11.06 G	80.53%	78.58	-
		AutoGAN-Distiller (Fu et al. 2020)	4.34 G	92.36%	78.33	-
		DMAD (Li et al. 2022)	3.18 G	94.40%	78.24	-
		MGGC (Ours)	1.69 G	97.02%	77.33	-
		MGGC (Ours)	2.97 G	94.77%	75.85	-

Table 1: Quantitative comparison of our proposed method with state-of-the-art GAN compression methods.

Qualitative Results: We visualize predictions of different models in Fig. 3. Our pruned model can preserve the fidelity of images with a much lower computational burden. Also, its superior performance compared to baselines is observable in the samples for Cityscapes (better preserving street structure) and Horse2Zebra (more faithful background color).

Stability of Our Pruning Method

We explore our method's stability by visualizing different loss values and the resource loss \mathcal{R} during pruning. Loss values for CycleGAN on Horse2Zebra are shown in Fig. 4 (a)-(c), and the ones for Pix2Pix on Cityscapes are presented in Fig. 4 (d)-(f). We found in our experiments that λ_2 has less impact on pruning dynamics than λ_1 , which is expected as the generator's capacity is more crucial to GANs' performance than the discriminator's capacity. Thus, we alter λ_1 to explore how it impacts the pruning process. We can observe that if λ_1 be large enough (Fig. 4 (a)-(e)), the loss values for G and D will get stable and remain close to each other when $\mathcal R$ converges to zero. Further, the difference between loss values for G and D in Fig. 4 (d)-(f) for our model is much smaller than the same value for GCC shown in GCC (Li et al. 2021) Fig. 5(a), which demonstrates that our method can preserve the balance between G and D more effectively than GCC during pruning. In summary, visualizations in Fig. 4 show that our method can successfully preserve the balance

between G, D after achieving the desired computation budget, leading to attain competitive final performance metrics.

Ablation Study

We present ablation results of our method with different settings in Tab. 2. We construct a Baseline by only pruning the generator G with a naive parameterization $v_G =$ round(sigmoid($-(\theta + q)/\tau$)) (Eq. 3) when using the full capacity discriminator. The results demonstrate that using pruning agents, pruning the discriminator, and establishing a feedback connection between qen_G and qen_D provide substantial performance gain compared to the baseline on Pix2Pix and CycleGAN models. This observation suggests that sophisticated designs of pruning agents are beneficial for pruning conditional GAN models. Further, considering local density structures over neighborhoods of the learned manifold (k > 0) boosts the performance of our method significantly than not leveraging them (k = 0). Remarkably, our method can almost recover the original model's performance for Pix2Pix (42.53 vs. 42.71) and even outperform it for CycleGAN (58.64 vs. 61.53) without using Knowledge Distillation (KD) in the finetuning process. These results show that the density structure over the learned manifold contains valuable information for pruning GAN models. Utilizing KD can further improve our model. Compared to GCC (Li et al. 2021), we do not use online distillation for The Thirty-Eighth AAAI Conference on Artificial Intelligence (AAAI-24)



Figure 4: Different losses given different λ_1 during the pruning process. (a)-(c) Loss values for CycleGAN on Horse2Zebra dataset. (d)-(f) Loss values for Pix2Pix on Cityscapes dataset. We normalize \mathcal{R} to the range [0, 1] for better visualization.

KD. Also, we do not learn the discriminator's architecture during finetuning, which saves computational costs.

Settings	Cityscapes		Horse2Zebra		
Settings	mIoU (†)	MACs	$FID(\downarrow)$	MACs	
Baseline	39.37		73.28		
+ D pruning	39.84		70.41		
+ Pruning Agents	40.68	2.05.0	67.60	2.50 G	
+ G↔D Feedback	40.99	3.03 G	66.72		
+ Manifold Pruning	42.53		58.64		
+ Distillation	44.63		55.06		
Original	42.71	18.60 G	61.53	56.80 G	

Table 2: Ablation results of our proposed method.

Visualization of Local Neighborhoods

We explore the difference between the learned neighborhoods of our model *vs*. the original one in Fig. 5. Samples with red frames are the predictions, and their neighbors on the right are obtained with the method in Eq. 1. In each column, green frames show the samples having the same source domain ('Edge') image, and the blue ones mean different source domain inputs. As can be seen, most of the neighbor samples of our pruned model are identical to those for the pruned model. In addition, the samples with different source images still have a semantically meaningful connection to the predicted image. These results demonstrate that our pruning objective, Eq. 2, can effectively regularize the pruned model to have a similar density structure over the neighborhoods of the original model's manifold.

Conclusions

We introduced a new compression method for image translation GANs that explicitly regularizes the pruned model to



Figure 5: Approximate neighborhoods on the learned manifold of our pruned model *vs*. the original model.

have a similar density structure to the original one on its learned data manifold. We simplify this objective for the pruned model by leveraging a self-supervised trained encoder to find approximate local neighborhoods on the manifold. Then, we proposed our adversarial pruning objective that encourages the pruned generator to have a similar density structure to the original model over each local neighborhood. In addition, we proposed a novel pruning scheme that uses two pruning agents that collaboratively play our adversarial pruning game to prune both the generator and discriminator. By doing so, our agents can effectively maintain the balance between the generator and discriminator, thereby stabilizing the pruning process. Our experimental results clearly illustrate the added value of using the learned density structure of a GAN model for pruning it compared to the baselines that combine CNN compression techniques.

Acknowledgements

This work was partially supported by NSF IIS 2347592, 2348169, 2348159, 2347604, CNS 2347617, CCF 2348306, DBI 2405416.

References

Aguinaldo, A.; Chiang, P.-Y.; Gain, A.; Patil, A.; Pearson, K.; and Feizi, S. 2019. Compressing gans using knowledge distillation. *arXiv preprint arXiv:1902.00159*.

Ba, J.; and Caruana, R. 2014. Do Deep Nets Really Need to be Deep? In *NeurIPS*.

Bengio, Y.; Léonard, N.; and Courville, A. 2013. Estimating or propagating gradients through stochastic neurons for conditional computation. *arXiv preprint arXiv:1308.3432*.

Cao, J.; Guo, Y.; Wu, Q.; Shen, C.; Huang, J.; and Tan, M. 2018. Adversarial learning with local coordinate coding. In *International Conference on Machine Learning*, 707–715.

Caron, M.; Misra, I.; Mairal, J.; Goyal, P.; Bojanowski, P.; and Joulin, A. 2020. Unsupervised learning of visual features by contrasting cluster assignments. *Advances in Neural Information Processing Systems*, 33: 9912–9924.

Casanova, A.; Careil, M.; Verbeek, J.; Drozdzal, M.; and Romero-Soriano, A. 2021. Instance-Conditioned GAN. In *Advances in Neural Information Processing Systems*.

Chan, C.; Ginosar, S.; Zhou, T.; and Efros, A. A. 2019. Everybody dance now. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, 5933–5942.

Chang, T.-Y.; and Lu, C.-J. 2020. TinyGAN: Distilling Big-GAN for Conditional Image Generation. In *Proceedings of the Asian Conference on Computer Vision*.

Chen, H.; Wang, Y.; Shu, H.; Wen, C.; Xu, C.; Shi, B.; Xu, C.; and Xu, C. 2020a. Distilling portable generative adversarial networks for image translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*.

Chen, T.; Kornblith, S.; Norouzi, M.; and Hinton, G. 2020b. A simple framework for contrastive learning of visual representations. In *International conference on machine learning*.

Chen, X.; Fan, H.; Girshick, R.; and He, K. 2020c. Improved baselines with momentum contrastive learning. *arXiv* preprint arXiv:2003.04297.

Cho, K.; van Merrienboer, B.; Gülçehre, Ç.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning Phrase Representations using RNN Encoder-Decoder for Statistical Machine Translation. In *EMNLP 2014*.

Choi, J.; Lee, J.; Yoon, C.; Park, J. H.; Hwang, G.; and Kang, M. 2022. Do Not Escape From the Manifold: Discovering the Local Coordinates on the Latent Space of GANs. In *International Conference on Learning Representations*.

Csordás, R.; van Steenkiste, S.; and Schmidhuber, J. 2021. Are Neural Nets Modular? Inspecting Functional Modularity Through Differentiable Weight Masks. In *International Conference on Learning Representations*.

Fu, Y.; Chen, W.; Wang, H.; Li, H.; Lin, Y.; and Wang, Z. 2020. AutoGAN-Distiller: Searching to Compress Generative Adversarial Networks. In *ICML*.

Ganjdanesh, A.; Gao, S.; and Huang, H. 2022. Interpretations steered network pruning via amortized inferred saliency maps. In *European Conference on Computer Vision*, 278–296. Springer.

Ganjdanesh, A.; Gao, S.; and Huang, H. 2023. EffConv: Efficient Learning of Kernel Sizes for Convolution Layers of CNNs. In *Thirty-Seventh AAAI Conference on Artificial Intelligence*. AAAI Press.

Gao, C.; Chen, Y.; Liu, S.; Tan, Z.; and Yan, S. 2020. Adversarialnas: Adversarial neural architecture search for gans. In *CVPR*, 5680–5689.

Ghimire, D.; Kil, D.; and Kim, S.-h. 2022. A Survey on Efficient Convolutional Neural Networks and Hardware Acceleration. *Electronics*, 11(6): 945.

Gong, X.; Chang, S.; Jiang, Y.; and Wang, Z. 2019. Autogan: Neural architecture search for generative adversarial networks. In *ICCV*, 3224–3234.

Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. *NeurIPS*, 27.

Grill, J.-B.; Strub, F.; Altché, F.; Tallec, C.; Richemond, P.; Buchatskaya, E.; Doersch, C.; Avila Pires, B.; Guo, Z.; Gheshlaghi Azar, M.; et al. 2020. Bootstrap your own latenta new approach to self-supervised learning. *NeurIPS*.

Gumbel, E. J. 1954. *Statistical theory of extreme values and some practical applications: a series of lectures*, volume 33. US Government Printing Office.

Han, S.; Pool, J.; Tran, J.; and Dally, W. 2015. Learning both weights and connections for efficient neural network. In *Advances in neural information processing systems*.

Härkönen, E.; Hertzmann, A.; Lehtinen, J.; and Paris, S. 2020. Ganspace: Discovering interpretable gan controls. *Advances in Neural Information Processing Systems*.

He, Y.; Lin, J.; Liu, Z.; Wang, H.; Li, L.-J.; and Han, S. 2018. Amc: Automl for model compression and acceleration on mobile devices. In *ECCV*.

Hou, L.; Yuan, Z.; Huang, L.; Shen, H.; Cheng, X.; and Wang, C. 2021. Slimmable Generative Adversarial Networks. In *AAAI*. AAAI Press.

Howard, A. G.; Zhu, M.; Chen, B.; Kalenichenko, D.; Wang, W.; Weyand, T.; Andreetto, M.; and Adam, H. 2017. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*.

Huang, X.; and Belongie, S. 2017. Arbitrary style transfer in real-time with adaptive instance normalization. In *ICCV*.

Isola, P.; Zhu, J.-Y.; Zhou, T.; and Efros, A. A. 2017. Imageto-image translation with conditional adversarial networks. In *CVPR*, 1125–1134.

Jang, E.; Gu, S.; and Poole, B. 2017. Categorical Reparameterization with Gumbel-Softmax. In *ICLR 2017*.

Jin, Q.; Ren, J.; Woodford, O. J.; Wang, J.; Yuan, G.; Wang, Y.; and Tulyakov, S. 2021. Teachers Do More Than Teach: Compressing Image-to-Image Models. In *CVPR*.

Khayatkhoei, M.; Singh, M. K.; and Elgammal, A. 2018. Disconnected manifold learning for generative adversarial networks. *NeurIPS*, 31.

Kim, B.-K.; Choi, S.; and Park, H. 2022. Cut Inner Layers: A Structured Pruning Strategy for Efficient U-Net GANs. *arXiv preprint arXiv:2206.14658*.

Ledig, C.; Theis, L.; Huszár, F.; Caballero, J.; Cunningham, A.; Acosta, A.; Aitken, A.; Tejani, A.; Totz, J.; Wang, Z.; et al. 2017. Photo-realistic single image super-resolution using a generative adversarial network. In *CVPR*.

Li, H.; Kadav, A.; Durdanovic, I.; Samet, H.; and Graf, H. P. 2017. Pruning filters for efficient convnets. *ICLR*.

Li, M.; Lin, J.; Ding, Y.; Liu, Z.; Zhu, J.-Y.; and Han, S. 2020. Gan compression: Efficient architectures for interactive conditional gans. In *CVPR*.

Li, S.; Lin, M.; Wang, Y.; Fei, C.; Shao, L.; and Ji, R. 2022. Learning efficient gans for image translation via differentiable masks and co-attention distillation. *IEEE Transactions on Multimedia*.

Li, S.; Wu, J.; Xiao, X.; Chao, F.; Mao, X.; and Ji, R. 2021. Revisiting Discriminator in GAN Compression: A Generator-discriminator Cooperative Compression Scheme. *Advances in Neural Information Processing Systems*, 34.

Lin, J.; Zhang, R.; Ganz, F.; Han, S.; and Zhu, J.-Y. 2021. Anycost gans for interactive image synthesis and editing. In *CVPR*.

Ni, Y.; Koniusz, P.; Hartley, R.; and Nock, R. 2022. Manifold Learning Benefits GANs. In *CVPR*.

Oldfield, J.; Markos, G.; Panagakis, Y.; Nicolaou, M. A.; and Ioannis, P. 2021. Tensor Component Analysis for Interpreting the Latent Space of GANs. In *BMVC*.

Park, T.; Liu, M.-Y.; Wang, T.-C.; and Zhu, J.-Y. 2019. Semantic image synthesis with spatially-adaptive normalization. In *CVPR*.

Parzen, E. 1962. On estimation of a probability density function and mode. *The annals of mathematical statistics*, 33(3).

Rastegari, M.; Ordonez, V.; Redmon, J.; and Farhadi, A. 2016. Xnor-net: Imagenet classification using binary convolutional neural networks. In *ECCV*. Springer.

Ren, Y.; Wu, J.; Xiao, X.; and Yang, J. 2021. Online multigranularity distillation for gan compression. In *ICCV*.

Sandler, M.; Howard, A.; Zhu, M.; Zhmoginov, A.; and Chen, L.-C. 2018. Mobilenetv2: Inverted residuals and linear bottlenecks. In *CVPR*.

Shen, Y.; and Zhou, B. 2021. Closed-form factorization of latent semantics in gans. In *CVPR*.

Shu, H.; Wang, Y.; Jia, X.; Han, K.; Chen, H.; Xu, C.; Tian, Q.; and Xu, C. 2019. Co-evolutionary compression for unpaired image translation. In *ICCV*.

Srivastava, N.; Hinton, G.; Krizhevsky, A.; Sutskever, I.; and Salakhutdinov, R. 2014. Dropout: a simple way to prevent neural networks from overfitting. *The journal of machine learning research*, 15(1): 1929–1958.

Tan, M.; and Le, Q. 2019. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International conference on machine learning*, 6105–6114. PMLR.

Tenenbaum, J. B.; Silva, V. d.; and Langford, J. C. 2000. A global geometric framework for nonlinear dimensionality reduction. *science*, 290(5500): 2319–2323.

Touvron, H.; Sablayrolles, A.; Douze, M.; Cord, M.; and Jégou, H. 2021. Grafit: Learning fine-grained image representations with coarse labels. In *ICCV*.

Tzelepis, C.; Tzimiropoulos, G.; and Patras, I. 2021. WarpedGANSpace: Finding non-linear RBF paths in GAN latent space. In *ICCV*.

Vinyals, O.; Blundell, C.; Lillicrap, T.; Wierstra, D.; et al. 2016. Matching networks for one shot learning. *Advances in neural information processing systems*, 29.

Voynov, A.; and Babenko, A. 2020. Unsupervised discovery of interpretable directions in the gan latent space. In *International conference on machine learning*.

Wan, A.; Dai, X.; Zhang, P.; He, Z.; Tian, Y.; Xie, S.; Wu, B.; Yu, M.; Xu, T.; Chen, K.; et al. 2020. Fbnetv2: Differentiable neural architecture search for spatial and channel dimensions. In *CVPR*.

Wang, H.; Gui, S.; Yang, H.; Liu, J.; and Wang, Z. 2020. GAN Slimming: All-in-One GAN Compression by A Unified Optimization Framework. In *ECCV*.

Wang, J.; Shu, H.; Xia, W.; Yang, Y.; and Wang, Y. 2021. Coarse-to-Fine Searching for Efficient Generative Adversarial Networks. *arXiv preprint arXiv:2104.09223*.

Wang, T.-C.; Liu, M.-Y.; Zhu, J.-Y.; Liu, G.; Tao, A.; Kautz, J.; and Catanzaro, B. 2018a. Video-to-video synthesis. *arXiv* preprint arXiv:1808.06601.

Wang, X.; Yu, K.; Wu, S.; Gu, J.; Liu, Y.; Dong, C.; Qiao, Y.; and Change Loy, C. 2018b. Esrgan: Enhanced superresolution generative adversarial networks. In *ECCV work-shops*.

Wu, B.; Dai, X.; Zhang, P.; Wang, Y.; Sun, F.; Wu, Y.; Tian, Y.; Vajda, P.; Jia, Y.; and Keutzer, K. 2019. Fbnet: Hardwareaware efficient convnet design via differentiable neural architecture search. In *CVPR*.

Ye, M.; Gong, C.; Nie, L.; Zhou, D.; Klivans, A.; and Liu, Q. 2020. Good subnetworks provably exist: Pruning via greedy forward selection. In *ICML*.

Yu, C.; and Pool, J. 2020. Self-supervised generative adversarial compression. *NeurIPS*, 33.

Zhang, L.; Chen, X.; Dong, R.; and Ma, K. 2022a. Regionaware Knowledge Distillation for Efficient Image-to-Image Translation. *arXiv preprint arXiv:2205.12451*.

Zhang, L.; Chen, X.; Tu, X.; Wan, P.; Xu, N.; and Ma, K. 2022b. Wavelet knowledge distillation: Towards efficient image-to-image translation. In *CVPR*.

Zhu, J.-Y.; Park, T.; Isola, P.; and Efros, A. A. 2017. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *ICCV*.