

Towards Device Efficient Conditional Image Generation

Nisarg A. Shah
snisarg812@gmail.com

AI Foundation
California, USA

Gaurav Bharaj
first.last@gmail.com

Abstract

We present a novel algorithm to reduce tensor compute required by a conditional image generation autoencoder without sacrificing quality of photo-realistic image generation. Our method is device agnostic, and can optimize an autoencoder for a given CPU-only, GPU compute device(s) in about normal time it takes to train an autoencoder on a generic workstation. We achieve this via a two-stage novel strategy where, first, we condense the channel weights, such that, as few as possible channels are used. Then, we prune the nearly zeroed out weight activations, and fine-tune the autoencoder. To maintain image quality, fine-tuning is done via student-teacher training, where we reuse the condensed autoencoder as the teacher. We show performance gains for various conditional image generation tasks: segmentation mask to face images, face images to cartoonization, and finally CycleGAN-based model over multiple compute devices. We perform various ablation studies to justify the claims and design choices, and achieve real-time versions of various autoencoders on CPU-only devices while maintaining image quality, thus enabling at-scale deployment of such autoencoders.

1 Introduction

High demand for consumer avatars, filters, and scene generation applications has led to an increased at-scale need of photo-realistic image generation. Such applications rely on Generative Adversarial Networks (GANs) [9] and supervised image-to-image style transfer [8, 25, 50] via autoencoders such as U-nets [40]. Technical advancements, and availability of deep learning APIs [10, 67] has helped achieve image generation. Backends of such APIs rely on fast tensor operations, parallelized via GPUs. However, real-time image generation via GAN-like methods has high deployment cost due to GPU compute costs and high break-even profitability point. Although certain edge devices are native GPUs [62] capable, they can suffer from slow inference, quality and resolutions deterioration of generated images. Thus, we desire a solution that can quickly optimize neural-nets for a given compute device, without sacrificing image quality and inference times. State-of-the-art literature suggests several approaches – neural architecture design (NAD) [19, 23], network architecture search (NAS) [68], neural-net compression (quantization [11], distillation [68], and pruning [11, 16]). We propose a novel net pruning algorithm that be employed along with NAD and NAS. Typical neural-net model compression techniques focus on image classification and detec-

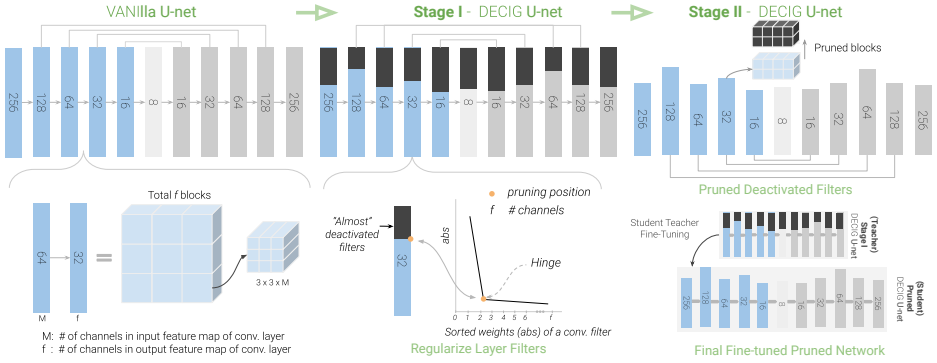


Figure 1: Our method DECIG, dynamically condenses channel filter and prunes GAN-based autoencoders for image generation. During Stage I, vanilla U-net autoencoder is trained with penalization, where weight distribution (centre-bottom) has several resulting near-zero value channels that can be pruned. During Stage II, the pruned network is fine-tuned in student-teacher manner using condensed model (from Stage-I) as the teacher.

tion [45]. Applying such methods for conditional image generation is relatively less explored and a naïve application may lead to image artifacts. Shu *et al.* [43] propose a channel pruning GAN compression evolutionary search algorithm, however, their method is designed for cyclic-consistency image generation [66], and nontrivial to extend for non-cyclic consistency GANs [25, 40]. Shu *et al.* [43] show that generators compressed by classifier compression methods [63] suffer performance decay compared to original generators. Chen *et al.* [5] propose GAN compression by training efficient generators by model distillation and remove dependency on cyclic consistency. Their student network is handcrafted and requires significant architectural engineering for good performance.

We propose a novel strategy to create compute efficient autoencoders for a given device. We do this by condensing neural-net channel filter weight distribution that condense filter usage, and later, prunes least activated filters while fine-tune using student-teacher model, where, the condensed autoencoder acts as the teacher. Our method is device agnostic, and optimizes neural-nets for given device. This also allows for a trade-off between computation complexity, and synthesized image quality, Fig 3. We summarize our novel contributions below:

1. A novel strategy to reduce compute costs via dynamic channel filter condensing and pruning GAN autoencoders for image generation.
2. A filter penalization loss for better filter weight distribution for easy pruning across layers, and detection of a “hinge” to get a minimum threshold for a particular filter structure, obtaining compute efficient autoencoders.
3. Optimized autoencoders perform real-time inference on CPU-only, CPU-GPU compute, with equivalent FID vs. vanilla autoencoders for conditional image generation.

2 Related Works

Conditional Image Generation Goodfellow *et al.* [9]’s tremendous success with image generation led to several vision and graphics applications [27, 66, 48]. For image-to-image

tasks, such as, segmentation masks to image generation [25, 86, 50] aim to achieve better image quality, these methods require high-end (Nvidia’s V100) GPUs for fast inference. At-scale use of such methods requires lower compute, and inference costs. Reducing neural-net compute is an active area of research. While most tackle non-image generation tasks, we explicitly aim to solve for conditional image generation. Such methods can be categorized as (1) Neural Architecture Design, (2) Neural Architecture Search, and (3) Neural-net Compression. Our work falls under the last category; we reduce compute by designing an algorithm that given an input autoencoder efficiently prunes filters, while maintaining generated image quality.

Neural Architecture Design of CNNs [13, 26, 42] led to massive gains for vision tasks. Such deep architectures are heavy on compute, even more so as autoencoders [40]. While several lightweight architectures have been proposed [10, 19, 21, 23, 42, 47, 55], these works exploit costly tensor blocks within architectures, and replace them with lightweight ones, or perform tensor compute in an efficient manner to improve the performance, generally for non-image generation. Wang *et al.* [62] use depth-wise separable convolutions to reduce tensor compute for denoising. On the other hand, our method is a novel strategy to reduce compute given an autoencoder, and can leverage these architectures as input.

Neural Architecture Search algorithmically search for efficient architectures. This search is highly nonlinear, with high compute and time complexity. Several works search architectures via reinforcement learning, and genetic algorithms [3, 35, 45]. Zoph *et al.* [58] search for transferable network blocks, surpasses manually designed architectures [24, 46]. While Cai *et al.* [9] speed up exploration for better architectures via network transformation. Fu *et al.* [2]’s distiller framework does adaptive search for operators types and channel widths. AutoML [30] framework searches for channel widths for existing generators which can be computationally expensive. In comparison, our method optimizes accuracy and weight distribution using a penalization loss and trains in a similar time as vanilla autoencoder. Also, our approach is complementary and can be combined with NAS.

Classification based pruning methods for conditional Image generation Though, use of classification(*Encoder only*) pruning methods seem to be a logical first choice for pruning in conditional image generation, several SOTA methods note otherwise. Yu *et al.* [64] conduct extensive experiments using *standard* pruning methods – manual pruning [12] and gradual pruning (AGP) [57] for GAN-based tasks (face generation). They note that image generated from manual pruning is of very poor quality, whereas, the models didn’t converge using AGP [57]. [49] and [31] highlight the instability in GAN minimax training and not being able to achieve Nash’s equilibrium for retaining performance metrics (FID), on applying classification based pruning methods for conditional image generation. Similarly, Shu *et al.* [43] show that applying classification neural-net compression [20, 34, 44] for generators compression suffers performance decay compared with the original generator. Based on these arguments we conclude that a naïve application of classification neural-net compression for image generation is less suitable, and thus the need for our approach.

Neural-Net Compression (Distillation, Quantization, Pruning) Wang *et al.* [49]’s unified GAN compression framework uses model distillation, channel pruning, and quantization to address this issue. Their channel pruning method is based on [33]’s batch-norm [22]

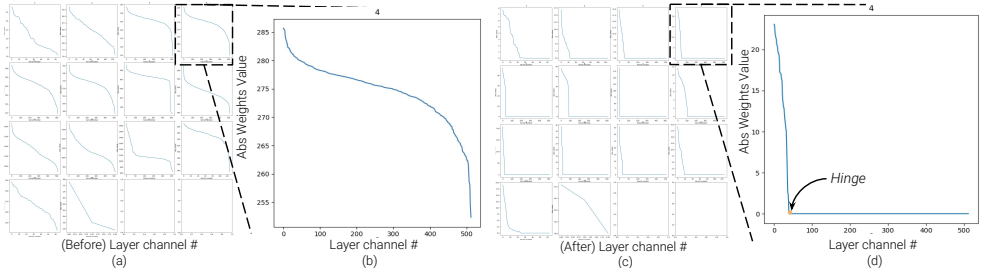


Figure 2: (a) & (c) shows weight distribution of Unet-64’s layers after training, without and with DECIG penalization; (b) and (d) shows zoomed-in version of corresponding *fourth* layer. (d) show the “hinge” for pruning *fourth* layer achieved by our method. X-axis indicates channel # in the layer, Y-axis indicates absolute value of weights of corresponding channel.

scale parameters, and measures importance among kernels of a convolutional layer. This dependence on batch-norm makes it inherently infeasible for architectures without batch-norm, while we seek a more generic approach. Aguinaldo *et al.* [4] compress DC-GAN [69] via knowledge distillation. Hou *et al.* [18] explore a method to generate nets with different channel sizes, while it leads to identity preserving image generation for face images, their FIDs may deteriorate. Lin *et al.* [52] propose a student-teacher learning method for interactive photo-realistic image generation and Lin *et al.* [51] propose a online collaborative distillation scheme to learn intermediate features of teacher generator and discriminator to optimize the student generator. Compared, we don’t have to pre-design student network and can condense a model using magnitude distribution plots and “hinge” modeling, Sec 3. Further, our algorithm can be used along with student-teacher training methods for joint training and device specific optimizations.

3 Method

A vanilla autoencoder generator G learns to synthesize an image I from an input segmentation map, $S \in \{H \times W \times 3\}$. Our pix2pix-like [25] setup uses a U-net [40] backbone generator, G , while the optimized generator G^* aims to be compute efficient s.t. the quality of generated images from both generators (G, G^*) is nearly equivalent, while G^* can be deployed across diverse hardware – CPUs, (e)GPUs optimizing for latency and image quality trade-off. The optimization condense (regularizes) filters used convolution layers of the autoencoder (Stage I), and later prunes least used filters (Stage II) and fine-tunes the pruned generator, Fig. 1.

As per, Li *et al.* [24] and Wen *et al.* [63] of the three levels for sparsity regularization, *coarse* channel-level sparsity provides a better trade-off between flexibility and ease of deployment. It can be applied on any neural-net with convolutional layers, and generates a *sparser* version of the original model. Channel-level sparsity requires pruning all adjacent connections associated with a particular channel, and makes it challenging to apply it directly on a pre-trained model due to generally non-existent zero weight channels in the neural-net, see Fig. 2 (a). To alleviate the problem of non-existent zero weight channels for sparsity regularization, we enforce a penalization loss in the training objective. Specifically, we introduce a loss function that operates on absolute value of filter weights and systematically

pushes the filter weights towards zero during training.

Unlike Liu *et al.* [63] that regularize added scaling factors after convolution or on adjacent scaling factor of normalization layer, our method operates on layer weights. We observe that using extra scaling factors not only add computational burden, but also such normalization specific methods increase the complexity of the approach when dealing with new methods with pre-activation structures and cross-connecting layers like ResNets [64], and DenseNets [62]. Further, such methods designed with batch-norm (or normalization layers in general) become unusable when working with newer, normalization free architectures. Our loss function operates on magnitude of channel weights, and works with such newer architectures.¹

Stage I - Channel Weight Regularization Channel pruning methods [62], utilize kernel magnitude as the criterion for relative importance across filters. On the other hand, when we train a network, a per channel importance factor γ , equivalent to magnitude of the weights of the corresponding channels is introduced. We train the network weights, and optimize the importance factor to condense the weights to as few channels as possible. This training objective for the i^{th} layer is given by:

$$L_i = \sum_{j=1}^n f(j) * ||W_{i,j}||_1 \quad (1)$$

Where, n is layer # in the network, j the channel # of the convolutional filter, and $W_{i,j}$ the filter weight of the i^{th} layer and j^{th} sorted channel. We explore three different channel regularization strategies with $j \in (1, n)$:

- *Uniform* feature channel regularization, $f(j) = 1.0$
- *Linear* feature channel regularization, $f(j) = j$
- *Exponential* feature channel regularization, $f(j) = e^{0.01j}$

We have selected three different channel regularisation methods, i.e. linear, uniform and exponential, to explore and verify the effects of relative penalisation (faster increase in penalisation as number of channels increase for a layer) and its effects on compression and quality.

Stage I - Layer Device Performance Regularization. GPU devices exploit tensor compute parallelism in convolution layers and simultaneously process large number of weight channels. CPU devices, however, carry out these operations sequentially and don't benefit from similar parallelism. Depending on the device and their memory allocation, relative convolutional operation speed across different spatial resolutions and feature map sizes differs considerably. For example, a convolution(kernel=3, stride=2) at 8×8 resolution with 512 input and output channels require 7.179 milliseconds (ms) on CPU and 1.132 ms on GPU. However, same convolution at 16×16 resolution takes $3 \times$ time and $1.6 \times$ time on CPU and GPU, respectively. Similarly, for a convolution(kernel=3, stride=2) at 128×128 resolution with input channel 1, if the number of output channels are increased from 32 to 128, the run-time for CPU gets quadrupled ($4 \times$ cost) while for GPU it remains nearly same. Based on this insight, we make the neural-net optimization device specific.

For model deployment, the compute devices are usually fixed, thus we propose run-time layer level (device-dependent) channel regularization strategy. We calculate the run-time

¹A layer indicates convolutional layer in the deep neural network, where a collection of channels make a single layer. For instance, if there is a 3×3 convolutional layer, with the input feature map of shape $H \times W \times M$ and output feature map of shape $H \times W \times N$, there would be total N channels each with shape $3 \times 3 \times M$.

Segmentation Mask	Ground Truth	Vanilla Unet-16	Vanilla Unet-32	Vanilla Unet-64	DECIG Unet-64 High reg.	DECIG Unet-64 Low reg.
CPU (FPS) ↑		70	25.9	7.3	25.2	16.4
GPU (FPS) ↑		200	168	25.9	156	131
FID ↓		74.5	58.7	47.3	48.6	37.9
Params (M) ↓		2.62	10.46	41.83	3.74	9.5




Figure 3: (Left to Right) Baseline variants: Unet- $\{16, 32, 64\}$, and DECIG versions of Unet-64 with high and low regularization. *Note: While better image generation methods exist, our emphasis is to maintain image quality given baseline autoencoders.*

for each layer across a particular device, and use it as a multiplicative factor $l(i)$ for that layer, to calculate total penalization. Our method allows device agnostic, multiply-accumulate (MAC) operations based, layer-level channel regularization. To this end, we calculate the multiplicative factor of each layer based on corresponding MAC operations of that particular layer. The penalization L_{PENAL} , minimax optimization [9] for a GAN is $\min_G \max_D L_{GAN}$ objective function, and final objective L_{GAN}^{ALAP} are given as:

$$\begin{aligned}
 L_{PENAL} &= \sum_{i=0}^n l(i) * L_i \\
 L_{GAN} &= \mathbb{E}_{y \in \mathcal{Y}} [\log(D(y))] + \mathbb{E}_{x \in \mathcal{X}} [\log(1 - D(G(x)))] \\
 L_{GAN}^{ALAP} &= L_{GAN} + L_{l1} + \alpha * L_{PENAL}
 \end{aligned} \tag{2}$$

Where, \mathcal{X} corresponds to random noise distribution, while \mathcal{Y} corresponds to real image distribution. $l1$ loss (ground-truth and generated images) is used during training as well. α is decided such that, for the first training iteration, ratio of $\alpha * L_{PENAL}$ and L_{l1} is 0.1

Stage II - Pruning and Distillation After Stage-I training, Eq 2, we obtain a model with a considerable amount of inactivated (near zero weight) channels. Due to the penalization loss, distinction between near zero and important channels is easily identifiable, Fig. 2(d). The inclination point that shows the threshold between these two types of channels is identified as the “hinge”. Fig 2 shows an example resulting weight distribution plot. We sort channels of a layer in Unet-64 model by magnitude (importance factor), and the hinge is identified at 50th channel for this particular layer. This way, we do not require to take an arbitrary guess or a global threshold [43] on the number of channels to be pruned, making the trained model *device efficient*. After identifying the “hinge”, channels lower in magnitude than hinge channel are pruned. Along with pruning these channels, we also remove corresponding incoming and outgoing connections and weights across all layers and obtain a compact network with fewer – parameters, run-time memory, and compute operations.

This hinge-based pruning has a minimal effect on the perceptual quality of generated images because the pruned channels magnitude is far less compared to that of non-pruned ones, sharp slope in Fig. 2(d). We also observe that, *almost* deactivated filters (Fig. 1) act as

Segmentation Mask	Ground Truth	Aguinaldo et al. 2019	Wang et al. 2020	DECIG - Unet-64 (high reg.)	DECIG - Unet-64 (low reg.)
CPU (FPS) ↑		25.9	20.3	25.2	16.4
GPU (FPS) ↑		168	142	156	131
FID ↓		55.9	52.6	48.6	37.9
Parameters (Million) ↓		10.46	6.72	3.74	9.5

Figure 4: Comparison with state-of-art Distillation and Pruning methods: (Left to Right) Segment Mask, G.T., Aguinaldo *et al.* [24], Wang *et al.* [49], DECIG-Unet-64 (high reg.), DECIG-Unet-64 (low reg.)

noisy channels. Thus when pruned, the autoencoder generates images with higher perceptual quality without further training. The minimal negative effect on perceptual quality, if any, can be further nullified by fine-tuning the pruned network via student-teacher training, where Stage-I trained model acts as the teacher model. Additionally, in several cases, like the over-parameterized or low weight penalization models discussed in the subsequent section, we find that the fine-tuned pruned network can reach higher perceptual scores than the vanilla network. After this stage, we finally obtain the optimized generator G^* .

4 Results and Ablation

We present our results, comparison and ablation results for paired and unpaired conditional image generation tasks. Specifically, results with generator Unet-16, 32, 64, 192 variants² for pix2pix [25] and Resnet for CycleGAN for respective vanilla versions, and DECIG (ours) optimized models. Since, the discriminator does not affect inference time, the student and teacher discriminator structure was kept the same.

To verify the generalizability across tasks and comparisons, we test DECIG on three applications: (1) segmentation mask to face generation, (2) face to cartoonized images for pix2pix model and (3) horse-to-zebra for CycleGAN-like model. We use CelebA-HQ dataset [28], that contains 30K high-quality face images and corresponding pixel-level segmentation masks face generation, and for cartoon images generation using animegan2³. We also evaluate the optimized generator via G^* using perceptual quality FID [17] score, memory consumption, run-times, and qualitative comparisons. We use MAC layer regularization and linear $f(j)$, Eq. 1, feature channel regularization for experiments, unless otherwise specified and measure the latency speed-up on Intel Xeon CPU E5-2686, and low-cost NVidia K80 GPU.

²Unet- x [25] indicates Unet autoencoder, where number of channels is x , that doubles after every strided convolution with an upper limit of 512.

³<https://github.com/bryandlee/animegan2-pytorch>

Source Image	UNet-64	ResNet	Shu et al.	Wang et al.	DECIG-ResNet
FID ↓	138	71.87	96	88	62.72
Flops (Giga) ↓	6.03 G	52.90 G	12.51 G	11.36 G	10.9 G
Memory (MB) ↓	160 MB	43.51 MB	10.19 MB	8.7 MB	7.42 MB

Figure 5: CycleGAN compression: (Left to Right) Source Image, Style transfer UNet-64, ResNet CycleGAN, CEC [44], GAN-Sliming [49] and DECIG-ResNet, respectively.

Results & Comparisons: Face Generation Application. Fig. 3 shows results of variants on U-net [40] architecture. Although, compared with Unet-64, *miniature* Unet variants (Unet-16 and Unet-32) have fewer parameters, their generated images have several artifacts – austere and blurry repeated patches. Images generated by our condensed generators look sharper and more realistic with lower inference times. Our method (Fig. 3) distributes weights in each layer, s.t., not only was it able to achieve better FID compared with corresponding vanilla models, but also hinge-based pruning had a minimal effect on FID. After student-teacher training, the model recovers from pruning artifacts and the resulting model has improved FIDs by 17%, and run-time (FPS), parameters and memory improve significantly.

Inference time for DECIG-Unet generators improves $3.5\times$ on CPU and $1.63\times$ on GPU. Primary objective of *high-reg* version of DECIG is to enable higher model compression with equivalent perceptual scores, compared to it’s vanilla variant, whereas, in *low-reg* versions enables higher perceptual quality over compression metrics. *high* and *low* indicates the penalization in the overall loss function. *high-reg* and *low-reg* is a relative term, we have scaled high-reg with 0.1 to get low-reg before training. The DECIG-Unet-64 high-reg variant achieves a real-time inference on CPU-only device and makes it feasible for cost-effective deployment.

Results & Comparisons: CycleGAN application. We go beyond Unets, and apply DECIG on unpaired image-to-to image translation, and use ResNet [56] as our generator, which is more complex than a Unet due to its cross-connections, and compare with recent SOTA methods like GAN-sliming [60]. Following [43, 49], we use GFLOps, model size (MB), and FID (between source style-transfer and target-style images) as a quantitatively measure for the ResNet generator.

On horse-to-zebra task, Fig. 5, DECIG-ResNet achieves $5.9\times$ reduction in model size compared to baseline and a 30% improvement in FID score. Our methods achieves significant improvement of 28% in FID scores compared to Wang *et al.* [49] along with improved results over compression matrices. Using *high-reg* version of DECIG-ResNet we achieve 33% reduction in the number of flops and $11\times$ reduction in memory usage with improvements in perceptual quality metrics, Table 7.

Face Images	Ground Truth	UNet-64	DECIG-UNet-64 Linear (low reg.)	DECIG-UNet-64 Uniform(Stage-1)	DECIG-UNet-64 Uniform(Stage-2)
FID ↓		29.20	27.52	25.65	24.31
Parameters (Million) ↓		41 M	18.5 M	41 M	7.8 M
Memory (MB) ↓		160 MB	96 MB	160 MB	30 MB
CPU (FPS) ↑		7.4	8.2	7.4	11.6
GPU (FPS) ↑		96	107	96	119

Figure 6: Face Cartoonization Compression: (Left to Right) Input, Cartoonized GT, UNet-64, DECIG-UNet-64 Linear (low-reg), DECIG-UNet-64 Uniform, Stage I and II.

Results: Cartoonization Application. Experiments are conducted on two different channel weight regularization approaches, i.e., Uniform and Linear, and the weighing of penalization is controlled to get balanced improvements over both FID scores and compression metrics. The qualitative and quantitative results of the cartoon-style transfer, together with model statistics (FLOPs and model sizes), are shown in Fig. 6. We also show results after only stage-I training to quantify student-teacher’s effect on our approach. Exploring improved student-teacher methods for training is orthogonal to our work and could be leveraged for additional improvements. For the cartoonization task, as fine-scale features are less critical, uniform channel weight regularization performs better compared with the linear version. DECIG-UNet performs well compared to UNet, with a 20% improvement in FID scores and 6× reduction in number of parameters. These results validate DECIG’s efficacy compared to models trained with the conventional approach in terms of perceptual quality and model parameter reduction.

Ablation: Overparametrization and Regularization. We evaluate our method on *overly-parametrized* Unet-192 and observe its advantages on over-fit reduction. DECIG-Unet-192 (Fig. 7(b)) typically achieves a 6× reduction in the number of parameters vs Unet-192. The pruned DECIG-Unet-192 achieves a lower FID score (47.3, 30% improvement) compared to the original model’s FID (65.3). We attribute this to our penalization strategy that addresses the over-fitting over the training set; see supplementary for qualitative results.

Ablation: Channel Weight and Layer Device Regularization Our channel weight regularization method supports several multiplicative functions like Uniform, Linear, and Exponential, and the corresponding quantitative predictions are shown in Tab. 1(b).

In our experiments, we observe linear $f(x)$ as the standard strategy based on trade-off perceptual image quality scores and run-time improvements, and uniform $f(x)$ performs slightly better when fine-scale features are less critical in the generation process like tasks such as cartoonization. We also test our approach on device agnostic layer level regularization (Tab. 1(a)). Here, FPS for model’s MAC optimized for CPU specifically (DECIG-CPU) has better FPS on CPU device compared to model optimized for general device, see supplementary

	Unet-64	DECIG CPU	DECIG GPU	DECIG MAC	Linear Pen. (High reg.)	Uniform Penalization	Exponential Penalization	Linear Pen. (Low reg.)
CPU (FPS)	7.3	28.6	26.1	25.2	25.2	22.5	21.9	16.4
GPU (FPS)	96	161	169	156	156	149	144	131
FID	47.3	51.48	51.61	48.6	48.6	45.61	44.46	37.9

Table 1: Comparison of Unet-64’s DECIG variants with layer weight {CPU, GPU, MAC} and channel weight {Linear (high-reg), Uniform, Exponential, Linear (low-reg)} regularization.

	ResNet	Shu et al. [43]	Wang et al. [49]	DECIG- ResNet		Vanilla ResNet- 128	DECIG - ResNet- 128	Vanilla Unet- 192	DECIG - Unet- 192
FID	148.81	139.88	120	133	<i>CPU, fps</i>	0.15	0.63	0.45	2.5
GFlops	52.90	12.16	12.05	8.21	<i>GPU, fps</i>	3.43	6.1	7.8	25
Mem. (mb)	43.51	10.02	9.04	4.83	FID	52.2	49.7	65.26	42.8
					<i>param</i>	54	4.9	376	61.5

(a)

(b)

Figure 7: (a) Comparison with CycleGAN Zebra-to-Horse Compression [43, 49] (b) Baseline conditional-GAN {Unet-192, ResNet}, and their DECIG versions.

for qualitative results.

Comparisons Summary with SOTA Pruning Methods. Aguinaldo *et al.* [4] uses model distillation to train a slim student generator from the corresponding bigger teacher network. We used Unet-64 as the teacher network and Unet-32, i.e., half slimmed channels for every layer as the student network and train using the distillation method. We observe slight FID improvement from vanilla Unet-32 model, however, generated results are blurry, Fig. 4. We also implement Wang *et al.* [49] for vanilla Unet-64 and obtain improved results in terms of perceptual quality (FID), parameter reduction, and memory usage compared to [4]. However, images generated from [49] have several artifacts such as, random color patterns, illuminations and thus, quality of generated images have more artifacts that our methods, while both networks have similar sizes and runtime, Figure 4.

5 Limitations and Conclusion

We present DECIG, a tensor compute reduction method that optimizes autoencoders for conditional image generation and achieves near real-time inference capabilities on CPU-only, and GPU devices. Our method achieves significant improvements over state-of-the-art methods wrt run-time and perceptual quality for various conditional image generation tasks – pix2pix (segmentation mask to images, images to cartoonization) and CycleGAN using autoencoders UNet and ResNet, respectively.

Several limitations exist: Although, task specific, our method does not completely preserve face identity attribute after compression and requires manual penalization strategy. We also require picking the hinge during Stage-II pruning manually. In the future, we want to explore improvements for these limitations via better perceptual losses, and introduce identity preserving loss [9], and propose automated hinge selection via clustering, and sharp curvature change modeling during compression aware training.

References

- [1] Martín Abadi, Paul Barham, Jianmin Chen, Zhifeng Chen, Andy Davis, Jeffrey Dean, Matthieu Devin, Sanjay Ghemawat, Geoffrey Irving, Michael Isard, et al. Tensorflow: A system for large-scale machine learning. In *12th {USENIX} symposium on operating systems design and implementation ({OSDI} 16)*, pages 265–283, 2016.
- [2] Angeline Aguineldo, Ping-Yeh Chiang, Alex Gain, Ameya Patil, Kolten Pearson, and Soheil Feizi. Compressing gans using knowledge distillation. *arXiv preprint arXiv:1902.00159*, 2019.
- [3] Andrew Brock, Theodore Lim, James M Ritchie, and Nick Weston. Smash: one-shot model architecture search through hypernetworks. *arXiv preprint arXiv:1708.05344*, 2017.
- [4] Han Cai, Tianyao Chen, Weinan Zhang, Yong Yu, and Jun Wang. Reinforcement learning for architecture search by network transformation. *arXiv preprint arXiv:1707.04873*, 4, 2017.
- [5] Hanting Chen, Yunhe Wang, Han Shu, Changyuan Wen, Chunjing Xu, Boxin Shi, Chao Xu, and Chang Xu. Distilling portable generative adversarial networks for image translation. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 2020.
- [6] Jiankang Deng, Jia Guo, Niannan Xue, and Stefanos Zafeiriou. Arcface: Additive angular margin loss for deep face recognition. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4690–4699, 2019.
- [7] Yonggan Fu, Wuyang Chen, Haotao Wang, Haoran Li, Yingyan Lin, and Zhangyang Wang. Autogan-distiller: Searching to compress generative adversarial networks. *arXiv preprint arXiv:2006.08198*, 2020.
- [8] Leon A Gatys, Alexander S Ecker, and Matthias Bethge. Image style transfer using convolutional neural networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 2414–2423, 2016.
- [9] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial nets. *Advances in neural information processing systems*, 27, 2014.
- [10] Kai Han, Yunhe Wang, Qi Tian, Jianyuan Guo, Chunjing Xu, and Chang Xu. Ghostnet: More features from cheap operations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 1580–1589, 2020.
- [11] Song Han, Huizi Mao, and William J Dally. Deep compression: Compressing deep neural networks with pruning, trained quantization and huffman coding. *arXiv preprint arXiv:1510.00149*, 2015.
- [12] Song Han, Jeff Pool, John Tran, and William Dally. Learning both weights and connections for efficient neural network. *Advances in neural information processing systems*, 28, 2015.

- [13] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 770–778, 2016.
- [14] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Identity mappings in deep residual networks. In *European conference on computer vision*, pages 630–645. Springer, 2016.
- [15] Yang He, Ping Liu, Ziwei Wang, Zhilan Hu, and Yi Yang. Filter pruning via geometric median for deep convolutional neural networks acceleration. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4340–4349, 2019.
- [16] Yihui He, Xiangyu Zhang, and Jian Sun. Channel pruning for accelerating very deep neural networks. In *Proceedings of the IEEE international conference on computer vision*, pages 1389–1397, 2017.
- [17] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. *Advances in neural information processing systems*, 30, 2017.
- [18] Liang Hou, Zehuan Yuan, Lei Huang, Huawei Shen, Xueqi Cheng, and Changhu Wang. Slimmable generative adversarial networks. *arXiv preprint arXiv:2012.05660*, 2020.
- [19] Andrew G Howard, Menglong Zhu, Bo Chen, Dmitry Kalenichenko, Weijun Wang, Tobias Weyand, Marco Andreetto, and Hartwig Adam. Mobilenets: Efficient convolutional neural networks for mobile vision applications. *arXiv preprint arXiv:1704.04861*, 2017.
- [20] Hengyuan Hu, Rui Peng, Yu-Wing Tai, and Chi-Keung Tang. Network trimming: A data-driven neuron pruning approach towards efficient deep architectures. *arXiv preprint arXiv:1607.03250*, 2016.
- [21] Jie Hu, Li Shen, and Gang Sun. Squeeze-and-excitation networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 7132–7141, 2018.
- [22] Gao Huang, Zhuang Liu, Laurens Van Der Maaten, and Kilian Q Weinberger. Densely connected convolutional networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4700–4708, 2017.
- [23] Forrest N Iandola, Song Han, Matthew W Moskewicz, Khalid Ashraf, William J Dally, and Kurt Keutzer. Squeezenet: Alexnet-level accuracy with 50x fewer parameters and < 0.5 mb model size. *arXiv preprint arXiv:1602.07360*, 2016.
- [24] Sergey Ioffe and Christian Szegedy. Batch normalization: Accelerating deep network training by reducing internal covariate shift. In *International conference on machine learning*, pages 448–456. PMLR, 2015.
- [25] Phillip Isola, Jun-Yan Zhu, Tinghui Zhou, and Alexei A Efros. Image-to-image translation with conditional adversarial networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1125–1134, 2017.
- [26] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Communications of the ACM*, 60(6):84–90, 2017.

- [27] Christian Ledig, Lucas Theis, Ferenc Huszár, Jose Caballero, Andrew Cunningham, Alejandro Acosta, Andrew Aitken, Alykhan Tejani, Johannes Totz, Zehan Wang, et al. Photo-realistic single image super-resolution using a generative adversarial network. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4681–4690, 2017.
- [28] Cheng-Han Lee, Ziwei Liu, Lingyun Wu, and Ping Luo. Maskgan: Towards diverse and interactive facial image manipulation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [29] Hao Li, Asim Kadav, Igor Durdanovic, Hanan Samet, and Hans Peter Graf. Pruning filters for efficient convnets. *arXiv preprint arXiv:1608.08710*, 2016.
- [30] Muyang Li, Ji Lin, Yaoyao Ding, Zhijian Liu, Jun-Yan Zhu, and Song Han. Gan compression: Efficient architectures for interactive conditional gans. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pages 5284–5294, 2020.
- [31] Shaojie Li, Jie Wu, Xuefeng Xiao, Fei Chao, Xudong Mao, and Rongrong Ji. Revisiting discriminator in gan compression: A generator-discriminator cooperative compression scheme. *Advances in Neural Information Processing Systems*, 34, 2021.
- [32] Ji Lin, Richard Zhang, Frieder Ganz, Song Han, and Jun-Yan Zhu. Anycost gans for interactive image synthesis and editing. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 14986–14996, 2021.
- [33] Zhuang Liu, Jianguo Li, Zhiqiang Shen, Gao Huang, Shoumeng Yan, and Changshui Zhang. Learning efficient convolutional networks through network slimming. In *Proceedings of the IEEE international conference on computer vision*, pages 2736–2744, 2017.
- [34] Jian-Hao Luo, Jianxin Wu, and Weiyao Lin. Thinet: A filter level pruning method for deep neural network compression. In *Proceedings of the IEEE international conference on computer vision*, pages 5058–5066, 2017.
- [35] Risto Miikkulainen, Jason Liang, Elliot Meyerson, Aditya Rawal, Daniel Fink, Olivier Francon, Bala Raju, Hormoz Shahrzad, Arshak Navruzyan, Nigel Duffy, et al. Evolving deep neural networks. In *Artificial intelligence in the age of neural networks and brain computing*, pages 293–312. Elsevier, 2019.
- [36] Taesung Park, Ming-Yu Liu, Ting-Chun Wang, and Jun-Yan Zhu. Semantic image synthesis with spatially-adaptive normalization. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 2337–2346, 2019.
- [37] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style, high-performance deep learning library. *Advances in neural information processing systems*, 32:8026–8037, 2019.
- [38] Antonio Polino, Razvan Pascanu, and Dan Alistarh. Model compression via distillation and quantization. *arXiv preprint arXiv:1802.05668*, 2018.

- [39] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. *arXiv preprint arXiv:1511.06434*, 2015.
- [40] Olaf Ronneberger, Philipp Fischer, and Thomas Brox. U-net: Convolutional networks for biomedical image segmentation. In *International Conference on Medical image computing and computer-assisted intervention*, pages 234–241. Springer, 2015.
- [41] Artsiom Sanakoyeu, Dmytro Kotovenko, Sabine Lang, and Bjorn Ommer. A style-aware content loss for real-time hd style transfer. In *proceedings of the European conference on computer vision (ECCV)*, pages 698–714, 2018.
- [42] Mark Sandler, Andrew Howard, Menglong Zhu, Andrey Zhmoginov, and Liang-Chieh Chen. Mobilenetv2: Inverted residuals and linear bottlenecks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 4510–4520, 2018.
- [43] Han Shu, Yunhe Wang, Xu Jia, Kai Han, Hanting Chen, Chunjing Xu, Qi Tian, and Chang Xu. Co-evolutionary compression for unpaired image translation. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 3235–3244, 2019.
- [44] Karen Simonyan and Andrew Zisserman. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556*, 2014.
- [45] Kenneth O Stanley and Risto Miikkulainen. Evolving neural networks through augmenting topologies. *Evolutionary computation*, 10(2):99–127, 2002.
- [46] Christian Szegedy, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich. Going deeper with convolutions. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 1–9, 2015.
- [47] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning*, pages 6105–6114. PMLR, 2019.
- [48] Ayush Tewari, Mohamed Elgharib, Gaurav Bharaj, Florian Bernard, Hans-Peter Seidel, Patrick Pérez, Michael Zöllhofer, and Christian Theobalt. Stylerig: Rigging stylegan for 3d control over portrait images, cvpr 2020. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*. IEEE, june 2020.
- [49] Haotao Wang, Shupeng Gui, Haichuan Yang, Ji Liu, and Zhangyang Wang. Gan slimming: All-in-one gan compression by a unified optimization framework. In *European Conference on Computer Vision*, pages 54–73. Springer, 2020.
- [50] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Guilin Liu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. Video-to-video synthesis. *arXiv preprint arXiv:1808.06601*, 2018.
- [51] Ting-Chun Wang, Ming-Yu Liu, Jun-Yan Zhu, Andrew Tao, Jan Kautz, and Bryan Catanzaro. High-resolution image synthesis and semantic manipulation with conditional gans. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8798–8807, 2018.

- [52] Yuzhi Wang, Haibin Huang, Qin Xu, Jiaming Liu, Yiqun Liu, and Jue Wang. Practical deep raw image denoising on mobile devices. In *European Conference on Computer Vision*, pages 1–16. Springer, 2020.
- [53] Wei Wen, Chunpeng Wu, Yandan Wang, Yiran Chen, and Hai Li. Learning structured sparsity in deep neural networks. *Advances in neural information processing systems*, 29:2074–2082, 2016.
- [54] Chong Yu and Jeff Pool. Self-supervised generative adversarial compression. *Advances in Neural Information Processing Systems*, 33:8235–8246, 2020.
- [55] Xiangyu Zhang, Xinyu Zhou, Mengxiao Lin, and Jian Sun. Shufflenet: An extremely efficient convolutional neural network for mobile devices. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 6848–6856, 2018.
- [56] Jun-Yan Zhu, Taesung Park, Phillip Isola, and Alexei A Efros. Unpaired image-to-image translation using cycle-consistent adversarial networks. In *Proceedings of the IEEE international conference on computer vision*, pages 2223–2232, 2017.
- [57] Michael Zhu and Suyog Gupta. To prune, or not to prune: exploring the efficacy of pruning for model compression. *arXiv preprint arXiv:1710.01878*, 2017.
- [58] Barret Zoph, Vijay Vasudevan, Jonathon Shlens, and Quoc V Le. Learning transferable architectures for scalable image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 8697–8710, 2018.