

GANORCON: Are Generative Models Useful for Few-shot Segmentation?

Oindrila Saha Zezhou Cheng Subhansu Maji
University of Massachusetts Amherst

{osaha, zezhoucheng, smaji}@cs.umass.edu

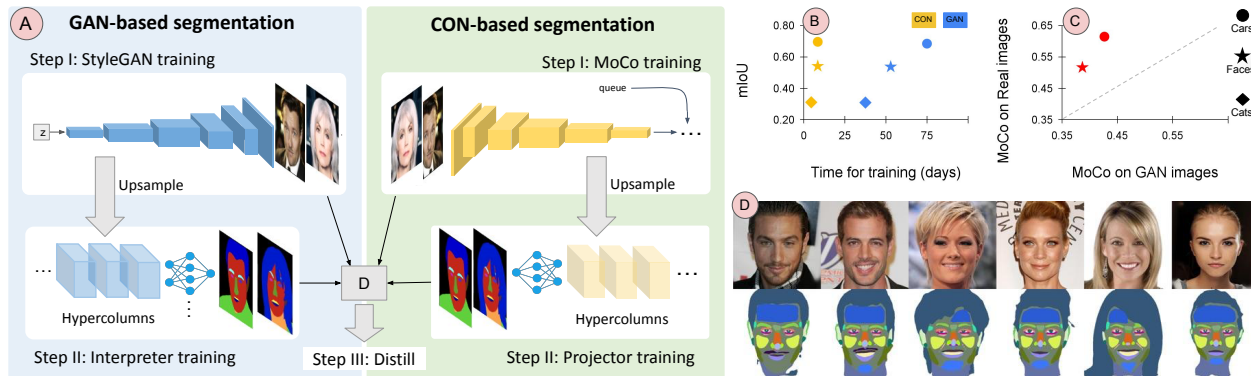


Figure 1. **GANORCON** — A comparison of GAN and contrastive learning (CON) approaches for few-shot part segmentation. **(A)** Both approaches consist of three steps. **Step I:** Train a GAN decoder for generating images or a contrastive learning based image encoder. **Step II:** Train a projector given a few labeled examples using hypercolumn representations from the decoder or encoder. **Step III:** For efficient inference, distill the GAN sampled images and their labels to an off-the-shelf feed-forward segmentation model. This step is optional for CON as the model is feed-forward. **(B)** CON representations achieve better performance while being nearly an order-of-magnitude faster to train than GAN representations on several datasets. **(C)** CON representations are significantly worse when trained on GAN generated images. **(D)** CON representations are effective at fine-grained part segmentation tasks — the figure shows the output of our method trained with 16 labeled faces.

Abstract

Advances in generative modeling based on GANs has motivated the community to find their use beyond image generation and editing tasks. In particular, several recent works have shown that GAN representations can be re-purposed for discriminative tasks such as part segmentation, especially when training data is limited. But how do these improvements stack-up against recent advances in self-supervised learning? Motivated by this we present an alternative approach based on contrastive learning and compare their performance on standard few-shot part segmentation benchmarks. Our experiments reveal that not only do the GAN-based approach offer no significant performance advantage, their multi-step training is complex, nearly an order-of-magnitude slower, and can introduce additional bias. These experiments suggest that the inductive biases of generative models, such as their ability to disentangle shape and texture, are well captured by standard feed-forward networks trained using contrastive learning.

1. Introduction

“When solving a problem of interest, do not solve a more general problem as an intermediate step.”

—Vladimir Vapnik

Generative modeling has long been used for unsupervised representation learning for recognition tasks. But recently, several works (e.g., [41, 46]) have shown that representations from off-the-shelf GANs (e.g., [3, 20]) can be re-purposed for part segmentation tasks, in part due to their ability to generate highly realistic images. These models can be trained with a few labeled samples, often outperforming transfer learning baselines for a range of object categories. Yet, generating images is arguably more challenging than learning representations for part labeling. This begs the question if the intermediate step of training a GAN is useful. What information, if any, is lost in the process of training a GAN, and if alternatives, in particular those based on unsupervised contrastive learning (e.g., [5, 8, 14]), are just as effective. Despite recent advances in generative and contrastive learning, a careful comparison of their ef-

effectiveness for few-shot part segmentation is lacking in the literature. We thus begin by presenting a strong contrastive learning baseline for this task. We then identify the key steps in these approaches and presents a series of experiments to evaluate their importance.

The two schemes are illustrated in Fig. 1 and described in §3. Generative modeling approaches [41, 46] first train a generator $\mathbf{x} = f(\mathbf{z})$ on the unlabeled samples $\mathbf{x} \sim \{\mathbf{x}_i\}_{i=1}^n$ and $\mathbf{z} \sim N(0, \sigma^2\mathbb{I})$. Once trained a few generated samples are labeled by an annotator to train a part label predictor h . This is used to generate a large labeled dataset $(\mathbf{x}, \mathbf{y}) \sim \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^m$ and “distilled” to a feed-forward segmentation network g using the dataset as a source of supervision. Contrastive learning (CON) techniques train a network $f(\mathbf{x})$ to encode an image using an objective that maximizes the similarity of images under transformations (e.g., [5, 10, 14]). Once trained a classifier h is trained to predict part labels given labeled examples on top of f . The combined model $h \circ f$ can be optionally distilled to another network g . While the details vary, we identify three factors that affect the overall performance on the task:

1. Complexity – how long does it take to train the model and how many steps are involved?
2. Effectiveness – how effective are the learned representations on the few-shot segmentation task. In particular how does the choice of the labeling network h and distillation step affect the overall performance.
3. Robustness – How do differences in distribution of GAN generated data from the real data affect labeling difficulty and performance.

Our experiments suggest that GAN representations offer no significant advantage over CON representations in each of these dimensions. GAN training is nearly an order-of-magnitude slower (Fig. 1 and §3.1). Across four datasets (Face-34, Face-8, Cat16, Car20) and two evaluation metrics (IoU and Weighted IoU), CON representations consistently outperform GAN representations, as well as strong baselines based on transfer learning and semi-supervised learning (Tab. 1). GANs performance is especially poor when measured directly using latent code optimization to embed the test images, but distillation to a feed-forward segmentation network improves their generalization significantly. Yet, their performance remains below CON representations (Tab. 2 & 3). CON representations see a smaller benefit from distillation unlike the GAN representations, suggesting that a significant portion of the generalization ability of GANs can be attributed to distillation to the final segmentation network. CON representations also generalize better when trained on domain specific datasets than ImageNet (Tab. 5). Remarkably, we find that while the performance of the final segmentation network trained on real and synthetically generated data are comparable, but CON representations trained on GAN generated images perform

significantly worse (Tab. 6). This might indicate that GAN generated imagery contain artifacts that makes instance discrimination artificially easy (See §5 for details), and puts caution to their use for fine-grained classification tasks.

In summary our contributions are: 1) a strong contrastive learning approach for part segmentation. We evaluate the effectiveness learning frameworks (MoCo [14] vs SimSiam [5] in Tab. 4) and architectures for predicting part segmentations on top of the learned representations. The approach outperforms recent StyleGAN-based approaches on standard few-shot part segmentation benchmarks. 2) An analysis of the GAN and CON representations in terms of the three factors listed above. We find that GAN-based approaches take longer to train, are less accurate, and suffer from labeling difficulty due to their inability to generate fine-grained parts. These experiments suggest that the inductive biases of GANs, such as their ability to disentangle shape and texture, are well captured by standard feed-forward networks trained using contrastive learning. The source code and data associated with the paper are publicly available at <https://people.cs.umass.edu/~osaha/ganorcon>.

2. Related Work

Generative models learn a probability distribution using unlabeled samples and extract representations from the learned distribution $p(\mathbf{x}|\theta)$. For example, Jaakkola and Haussler [18] map data points \mathbf{x} based on Fisher score $U_{\mathbf{x}} = \nabla_{\theta} \log p(\mathbf{x}|\theta)$, enabling the use of Hidden Markov Models to represent DNA sequences. Fisher vector representations [35, 36] have been used to encode the distribution of local descriptors in images (e.g., SIFT [26], LBP [30]) using Gaussian mixture models [27]. RBMs [37], VAEs [22], GANs [12], among others employ deep networks to learn distributions over complex, high-dimensional signals representing images, 3D shapes, speech, and text. Once trained, the internal representations of these networks can be used for discriminative tasks.

Besides their use in image generation and editing tasks, the ability to generate highly realistic data has motivated the community to explore GANs as a source of training data for discriminative tasks. Besnier et al. [2] train image classifiers based on class-conditional samples from BigGAN [3] and identify heuristics to improve generalization. Jahanian et al. [19] use GANs to sample views of data to learn representations in a contrastive learning framework, while Tanaka and Aranha [38] use GANs for data augmentation. While much of prior work has focused on image classification, two recent works, DatasetGAN [46] and ReGAN [41], show that hypercolumn representations [13] from StyleGANs [20] are effective for fine-grained image segmentation. These representations generalize well when a small number of labeled examples are provided, and the model can be distilled to an off-the-shelf image segmen-

tation model by training it on a large number of GAN-generated images and their labels.

Yet, generating data can be more challenging than learning the desired invariances directly. Self-supervised learning provides an alternative for learning these invariances from unlabeled data. These techniques are based on training deep networks to solve a proxy task derived from the data alone. While early works proposed tasks such as image colorization [45], orientation prediction [11], jigsaw puzzle tasks [29], and inpainting [32], **contrastive learning** [1, 5–7, 10, 14, 17, 31, 40, 43] has emerged as the leading technique. These techniques are based on applying synthetic geometric and photometric image transformations and maximizing the mutual information between the transformations of a single image based on variations of the InfoNCE loss [31]. Contrastive learning has been shown to be an effective alternative to ImageNet pre-training for various downstream tasks. While the goal is often to learn a global image representation, several works have proposed variants for dense image prediction tasks such as part segmentation or landmark detection [9, 33, 39, 42]. The dense variants are based on extracting hypercolumns across layers of a network to obtain spatially varying representations, or modifying the contrastive learning objective to account for different locations within an image, e.g., by adding a within-image contrastive loss.

The goal of this work is to compare and contrast GAN-based and contrastive learning (CON) based approaches for few-shot segmentation based on their complexity, effectiveness and robustness and highlight some of the challenges in each. GANs act as a compact representation of the dataset and their internals can be exploited for few-shot discriminative tasks. Yet, there is a loss of information in the process of learning the data distribution which we quantify based on the ability to project samples into the latent space of GANs and the loss in performance of contrastive models trained on the GAN generated images. The next few sections describe these tradeoffs in detail.

3. GANORCON: Generative or Contrastive

We present a contrastive learning framework for few-shot fine-grained part segmentation. We construct our approach in phases similar to GAN-based segmentation methods as presented in DatasetGAN [46] and ReGAN [41], so as to directly compare various aspects such as computational complexity, effectiveness, labeling difficulty and robustness. In the following paragraphs, we compare the construction of the GAN and CON methods for each phase. The steps include I) a representation learning framework (generative or contrastive), II.0) extracting hypercolumns from the trained model, II) training a few-shot projector over feature activation maps learnt in the previous step, and III) a supervised distillation training step using examples

generated from the combined model of steps I and II.

Setup: Let $\mathbf{x} \in \mathbb{R}^{H \times W \times 3}$ denote an image of an object and $\mathbf{y} \in \mathbb{M}^{H \times W} | \mathbb{M} = \{1, \dots, C\}$ denote the pixel-wise classification labels, with C denoting the number of classes. Given a large number of unlabeled examples $\{\mathbf{x}\}$ and a limited number of labeled examples $\{(\mathbf{x}, \mathbf{y})\}$, the goal is to learn a mapping function $\mathbf{f}(\mathbf{x}) : \mathbf{x} \rightarrow \mathbf{y}$ that generates a pixel-wise classification label.

Step I – Representation Learning: Given a large set of unlabeled images $\{\mathbf{x} | \mathbf{x} \sim p_{data}(\mathbf{x})\}$, a generative model (f_{GEN}) is trained to map from a noise distribution $p_z(\mathbf{z})$ to the real data space via a min-max objective:

$$\min_{f_{GEN}} \max_{f_D} \mathbb{E}_{\mathbf{x} \sim p_{data}(\mathbf{x})} [f_D(\mathbf{x})] - \mathbb{E}_{\mathbf{z} \sim p_z(\mathbf{z})} [f_D(f_{GEN}(\mathbf{z}))] \quad (1)$$

where f_{GEN} denotes the generator network and f_D the discriminator network. The generator tries to mimic a real image while the discriminator tries to distinguish between a real image and an image from the generator network.

On the other hand, in our approach we train a contrastive learning model f_{CON} which learns underlying representations by differentiating between images in the same unlabeled dataset. We experiment with methods that work both with (MoCo) and without (SimSiam) negative examples.

The Momentum Contrast (MoCo [15]) method minimizes the Noise-Contrastive Estimation (InfoNCE [31]) loss. Given a sample \mathbf{x} and its transformation \mathbf{q}_+ as well as other samples $\mathbf{q}_i, i \in \{1, 2, \dots, N\}$, such that $\mathbf{q}_+ \in \mathbf{q}_i$, InfoNCE is defined as:

$$\mathcal{L}_{InfoNCE} = -\log \frac{\exp(\langle f_{CON}(\mathbf{x}), f_{CON}(\mathbf{q}_+) \rangle)}{\sum_{i=1}^N \exp(\langle f_{CON}(\mathbf{x}), f_{CON}(\mathbf{q}_i) \rangle)} \quad (2)$$

Here \mathbf{q}_+ is generated by applying photometric and geometric transformations on the anchor sample \mathbf{x} . We use MoCo V2 [7] in our experiments which uses the above loss.

Simple Siamese (SimSiam [8]) network trains without the use of negative examples, wherein given an additional predictor network (h), it minimizes the cosine feature distance between two views (\mathbf{x}, \mathbf{q}_+) of the same image using:

$$\mathcal{L}_{cosine} = -\frac{h(f_{CON}(\mathbf{x})) \cdot f_{CON}(\mathbf{q}_+)}{\|h(f_{CON}(\mathbf{x}))\|_2 \cdot \|f_{CON}(\mathbf{q}_+)\|_2} \quad (3)$$

Networks trained using the above methods complete the first part of the method. Next we extract hypercolumns [13] to learn segmentation maps from.

Step II.0 – Hypercolumn Extraction: A deep network of n layers (or blocks¹) can be written as $\Phi(\mathbf{x}) = \Phi^{(n)} \circ \Phi^{(n-1)} \circ \dots \circ \Phi^{(1)}(\mathbf{x})$. A representation $\Phi(\mathbf{x})$ of size $H' \times W' \times K$ can be spatially interpolated to input size $H \times W \times K$

¹Due to skip-connections, we cannot decompose the encoding over layers, but can across blocks.

Method	Face34	Face34 weighted	Face8	Car20	Cat16
Transfer Learning	0.4577 ± 0.0151	-	0.6283	0.3391 ± 0.0094	0.2252 ± 0.0061
Semi-Supervised [28]	0.4817 ± 0.0066	-	0.6336	0.4451 ± 0.0057	0.3015 ± 0.0035
DatasetGAN [46]	0.5346 ± 0.0121	-	0.7001	0.6233 ± 0.0055	0.3126 ± 0.0071
DatasetGAN*	0.5365 ± 0.0043	0.7959 ± 0.0237	0.6971 ± 0.0095	0.6840 ± 0.0322	0.3083 ± 0.0068
Ours - CONV + distill	0.5406 ± 0.0015	0.8241 ± 0.0133	0.7048 ± 0.0028	0.6956 ± 0.0178	0.3101 ± 0.0042

Table 1. **Comparison of our best method with prior art:** We report mean IOU and standard deviation of each approach for various datasets. Face8 is an out of domain experiment where the training set is same as that of Face34 but the testing set is from a different dataset. Our approach with f_{PROJ} as CONV with distillation outperforms previous methods in all of the test settings. DatasetGAN* denotes results after re-implementation using their updated code, where they change their upsampling technique and uncertainty calculation.

Method		Face34		Face34 weighted		Face8		Car20		Cat16		ReGAN/Face10	
Pretrain	Projector	Direct	Distill	Direct	Distill	Direct	Distill	Direct	Distill	Direct	Distill	Direct	Distill
StyleGAN	MLP	0.4301	0.5365	0.7572	0.7959	0.6331	0.6971	0.5153	0.6840	0.1585	0.3083	0.7624	0.7505
MoCo	MLP	0.5162	0.5341	0.7784	0.7953	0.6874	0.7052	0.6143	0.6788	0.2902	0.3048	0.7509	0.7737
MoCo	CONV	0.5414	0.5406	0.8038	0.8241	0.6994	0.6980	0.6447	0.6956	0.3021	0.3101	0.7683	0.7811

Table 2. **Detailed comparison of the effect of each step on mIoU :** The first row shows scores for generative models where "Direct" is obtained by performing latent optimization on test images to predict masks using obtained latent and "Distill" is the score after distillation. The other two rows are scores of contrastive models with different architectures of f_{PROJ} . The blue shaded column is the comparison with ReGAN method, where we use different architectures for the GAN method and different train/test settings following their paper.

to produce a pixel representation $\Phi_i(\mathbf{x}) \in \mathbb{R}^{H \times W \times K}$. The hypercolumn representation of layers l_1, l_2, \dots, l_n is obtained by concatenating interpolated features from corresponding layers i.e. $\Phi_i(\mathbf{x}) = \Phi_i^{(l_1)}(\mathbf{x}) \oplus \Phi_i^{(l_2)}(\mathbf{x}) \oplus \dots \oplus \Phi_i^{(l_n)}(\mathbf{x})$. Both generative and contrastive methods extract hypercolumns from their respective trained models for few-shot segmentation. For StyleGAN this is comprised of the outputs of the AdaIN layers.

Step II – Few-shot Segmentation: Using the extracted $\Phi_i(\mathbf{x})$, both approaches train a projector network (f_{PROJ}) to obtain pixel-wise confidence maps of shape $H \times W \times C$ where C denotes the number of classes.

DatasetGAN utilizes an ensemble of MLP classifiers as f_{PROJ} and use majority voting at test time². Our method is more simpler in comparison. We use two variations for the architecture of f_{PROJ} , one being a MLP and the other a modified UNet [34] (see Appendix B for architectures) which we will call CONV hereon.

At this point, we can evaluate on test images directly to obtain our performance measure, however to compare with the generative method we also perform distillation.

Step III – Supervised Distillation: To infer on unseen images, generative models can use two approaches 1) latent code optimization to obtain latents which can nearly estimate the test image, 2) perform distillation using a standard supervised segmentation network. Since the first method affects performance (see Tab. 2) and it is time consuming to obtain latents for test images, the second method is preferred by DatasetGAN. For generative method, the optimal

solution of the distillation process can be formulated as:

$$\theta_S^* = \operatorname{argmin}_{\theta_S} \mathbb{E}_{\mathbf{z}}[\mathbf{CE}(T(\mathbf{z}; \theta_T), S(f_{GEN}(\mathbf{z}; \theta_S)))] \quad (4)$$

where S is the Student network with parameters θ_S , T is the teacher network with parameters θ_T and \mathbf{CE} denotes the Cross Entropy loss. For the generative method, $T(\mathbf{z}; \theta_T)$ involves generating images from random latents \mathbf{z} using f_{GEN} and then obtaining corresponding masks using the Interpreter network.

In our case, $T(\mathbf{x}; \theta_T)$ is inferencing over our trained contrastive + projector networks with images \mathbf{x} from a pre-existing dataset to obtain masks. The student parameters can be optimized as below:

$$\theta_S^* = \operatorname{argmin}_{\theta_S} \mathbb{E}_{\mathbf{x}}[\mathbf{CE}(T(\mathbf{x}; \theta_T), S(\mathbf{x}; \theta_S))] \quad (5)$$

It is important to note that the distillation step is optional for our method since it can directly infer on any given test image. However this step offers a little improvement in performance in some cases.

3.1. Design considerations

3.1.1 Computational Complexity

The computational bottleneck for the GAN based methods is the StyleGAN [20] training, whereas for our case its the contrastive learning network. According to the official repository of StyleGAN, it takes 11 days and 8 hours to train on 70000 images of resolution 1024×1024 on 4 Tesla V100 GPUs. Our network trained on the same dataset takes 7 days and 10 hours on 4 1080TI GPUs. By crude linear scaling using the ratios of RAM and FLOPs of the GPUs, we can infer that our method is around $7 \times$ faster.

²DatasetGAN also uses this ensemble to predict uncertainty for filtering generated data before the distillation step.

We also compare with a StyleGAN2 [21] based method. In this case, our method is even more computationally efficient. StyleGAN2 takes 18 days and 14 hours on the same dataset and setting as discussed above. Thus our unsupervised training is $11\times$ faster in comparison.

The second step of both the methods i.e. the Style Interpreter training of DatasetGAN and Projector network training of our method have similar complexity as the input feature dimensions and network architectures are similar. Our method takes lesser time compared to DatasetGAN as they use an ensemble of classifiers. Note that this step takes the least amount of time as number of examples is very limited.

Furthermore, StyleGAN based methods need the distillation step to infer on unseen images. The StyleGAN + interpreter model first needs to generate a large number of images. The data generation and training of the distillation network takes ~ 13 hours on 4 1080ti GPUs with DatasetGAN’s settings for 10k training images. Our method can avoid this step and still produce competitive results.

3.1.2 Effectiveness

We evaluate our method on the same train and test settings as DatasetGAN for various datasets. We also compare with ReGAN for face data. Our method using MLP is able to achieve similar results w.r.t. the GAN based approaches after distillation, while our CONV based method outperforms the baselines both with and without (except Cars data) distillation. We discuss our results in greater detail in §5.

3.1.3 Labeling Difficulty

GAN based approaches suffer from the following two major drawbacks when it comes to training and inference of few-shot segmentation, which our approach is able to avoid.

1. To create training examples for the few-shot setting, generative model based methods need manual annotations on images generated by the GAN. This is due to the fact that latent optimization is not very accurate and is especially taxing for a pixel-wise classification task. On the other hand, our approach can use any existing dataset to train the segmentation projector network.
2. To evaluate GAN based segmentation network directly, we again need to perform latent code optimization. We show in Tab. 2 that this creates a gap in the final performance. Thus, to get a fair evaluation of the GAN based methods an extra step of distillation is required. This requires the GAN to generate a large number of images and labels ($\sim 10k$) to train a supervised segmentation network.

3.1.4 Robustness

Representations learned by generative modeling techniques are very sensitive to affine transforms. GANs will not be able to produce images which have a small perturbation w.r.t. their learnt representation space such as translation.




Figure 2. **Robustness of GAN vs CON** : Generative method is not able to model test images when some perturbation is introduced. The third column shows the resultant images after 10k iterations of latent code optimization. The contrastive model is able to deal with the shift in the images, wherein the resultant masks are shifted similarly to the input as visualized in the last column.

Iterations	1000	5000	10000	20000
IoU	0.4178	0.4270	0.4302	0.4283

Table 3. **Direct evaluation of StyleGAN + interpreter** : mIoU vs iterations of latent optimization on images of Face34 test data.

Our approach is robust to such affine transforms and outputs a segmentation map with affine corresponding to the input image. We show this by shifting the test images and inferring our model on them. Fig. 2 shows our results on the transformed data. To get segmentation maps from the GAN based network for the translated images, we perform latent optimization on them. Fig. 2 shows that the generative model is not able to map these shifted images. Since the test images are shifted to the left, the StyleGAN generates images where the head is turned to the right, which is a bias of the underlying distribution that it has learned from the training data. The GAN based methods thus rely on the extra distillation step to overcome this disadvantage.

4. Experiments

In this section we present the experimental details of our approach. We outline the datasets used, implementation details for each step, evaluation metrics and the settings for comparison to baseline methods.

4.1. Datasets

4.1.1 Human Faces

Training: For training the contrastive learning models, we use the Flickr-Faces-HQ [20] and the CelebAMask-HQ [25] datasets. They have 70k and 30k images respectively, of 1024×1024 resolution. For all training, we resize images to 512×512 unlike StyleGAN which uses the original larger size. For the few shot segmentation training part, we use all 16 images from the Face34 training dataset released by DatasetGAN [46]. The number in the names of the datasets correspond to the number of classes. We then choose a subset comprising of 10k images from CelebAMask-HQ to perform inference of our trained

model. These images paired with the corresponding predicted masks form the training set for distillation. For all experiments of DatasetGAN we generate 10k images as well.

We also compare our method with ReGAN, where we train our models using the same datasets as their method. We use our unsupervised learning model that is trained on the Flickr-Faces-HQ dataset and train the few shot segmentation part using the same subset of 10 images from CelebAMask-HQ as ReGAN. Correspondingly, we convert the labels to 10 classes from 19 (see Appendix A). We refer to this 10 class train/test data as Face10. For distillation we use the same subset of 10k images as earlier.

Testing: We use the testing set with 20 images of Face34 to compare to DatasetGAN’s corresponding scores. DatasetGAN also reports scores on Face8 which is a modified version of CelebAMask-HQ with 8 classes. We use the same label transformation (see Appendix A) and sample 50 images randomly from CelebAMask-HQ’s testing set of 2k images to do the evaluation for our method and baselines. For comparing with ReGAN we use a subset of 190 images from the training set of CelebAMask-HQ following their method. For all evaluation experiments, we resize images to 512×512 resolution like both the GAN methods.

4.1.2 Cats

Training: We use LSUN Cats dataset [44] for training our unsupervised network. Unlike StyleGAN which uses the whole dataset (70M images) to train, we use only a subset of 100k images and still achieve better results. We resize images to 256×256 resolution for all training runs. For the few shot segmentation part we use Cat16 dataset released by DatasetGAN. We filter out few examples which have visible annotation errors and finally use 28 images for training.

Testing: We use the Cat16 testing set which has 20 images and resize them to 256×256 to calculate IoU.

4.1.3 Cars

Training: Similar to the setting of Cats dataset, we use a subset of 100k images from 46M total images in LSUN Cars dataset [44] to train our contrastive model. We use DatasetGAN’s Car20 training set comprising of 33 images for training the few shot model. In this case, we resize to 512×512 for both unsupervised and few shot training runs.

Testing: We use 10 images from the Car20 test set and resize to 512×512 similar to DatasetGAN for evaluation.

4.2. Implementation Details

Contrastive Learning. For unsupervised contrastive learning, we use two different frameworks namely MoCoV2 [7] and SimSiam [8]. We use ResNet50 [16] as the backbone architecture for both of these approaches. We train MoCoV2 for 800 epochs and SimSiam for 400 epochs. We use a batch size of 32 for the training where input images are resized to 512×512 (faces and cars) and 128 for the 256×256 setting (cats).

Method	MoCo	SimSiam
MLP	0.5162	0.5038
CONV	0.5414	0.5297
MLP + distill	0.5341	0.5188
CONV + distill	0.5406	0.5173

Table 4. **Effect of contrastive model on downstream IoU:** We report mIoU for the Face34 dataset for each method varied over the architecture of f_{PROJ} and before and after distillation.

Projector	Pretraining Strategy			
	ImageNet classification	Imagenet MoCo	FFHQ MoCo	CelebA MoCo
MLP	0.1893	0.4650	0.5023	0.5162
CONV	0.2551	0.4828	0.5132	0.5414

Table 5. **Effect of pretraining of f_{CON} on mIoU :** The first column reports scores using ResNet50 trained on ImageNet1K classification task. The other columns compare MoCo trained with different datasets with the same ResNet50 backbone. All scores are on the Face34 test set without distillation.

Few Shot Segmentation. After obtaining the trained ResNet50 we freeze its weights and extract hypercolumns by stacking activations of the layers from `conv_2x` to `conv_5x`. We upsample all the hypercolumns to the same resolution as the input image and concatenate them. We use the resultant 3D feature tensor to train our projector network. We use two different approaches : MLP and CONV as our f_{PROJ} . We use random resized crop, flip and color jitter as augmentations. We train the MLP for 800 epochs with an initial lr of 0.001 and cosine lr decay. We train the UNet based architecture - CONV, for 200 epochs with an initial lr of 0.0005. During this training, in the ResNet50 architecture we change the stride of the first convolution layer from 2 to 1 while extracting hypercolumns to train the projector model in order to preserve more spatial information.

Supervised Distillation. Once the few shot segmentor is trained, we use it to create a dataset to perform distillation. We use Deeplabv3 [4] with ResNet101 as the student model for all experiments. We use an initial lr of 0.001 with a batch size of 8 and train for 2 epochs for every setting.

4.3. Evaluation

We follow the same evaluation strategy as DatasetGAN for all scores that we report. The testing set is split into five folds and each set is chosen as validation set to choose the best checkpoint for testing. We finally report the mean Intersection over Union (mIoU) over all the folds.

We also carry out a comparison of weighted mIoU for human face datasets, consistent with ReGAN’s evaluation method. We compute the weight of each class as the ratio of the number of ground-truth pixels belonging to the class to the total number of pixels. We then perform the same cross validation as described above.

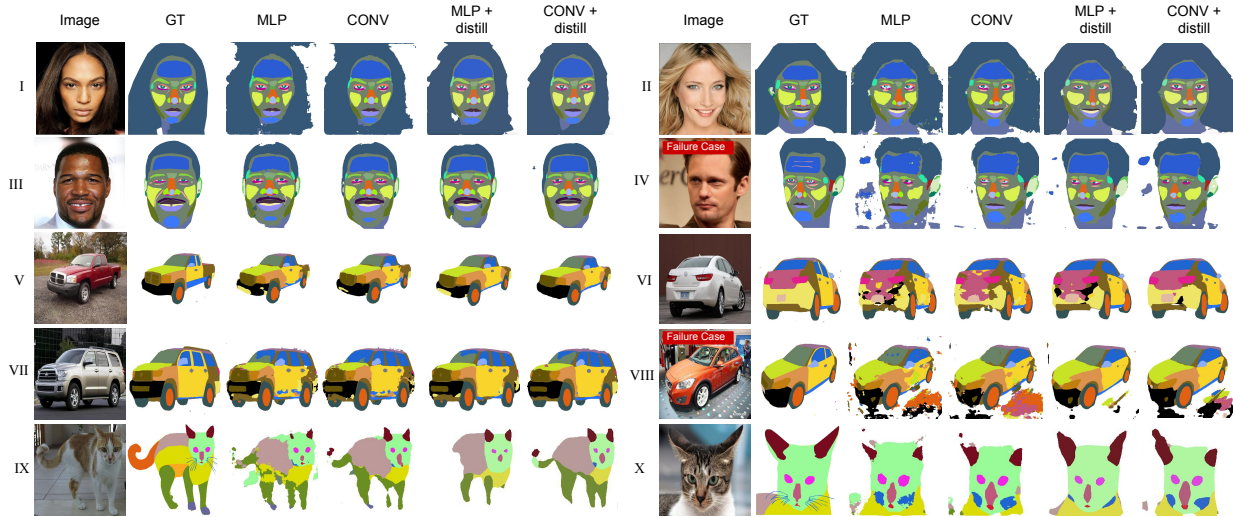


Figure 3. **Qualitative Results on various datasets:** We compare the visual results for our MLP and CONV methods both with and without distillation. We present results on the test sets of Face34, Car20 and Cat16 datasets. In image IV and VIII, we show two failure cases.

4.4. Comparison to Baselines

Direct Evaluation of StyleGAN segmentor. DatasetGAN’s method has three training steps : StyleGAN, Style interpreter and parts segmentation network (DeepLabV3). We want to evaluate their method at the second step i.e. infer using the StyleGAN + interpreter network with our test set to produce segmentation masks and then compute mIoU. To do so we first need to find the latents which most nearly produce the test images. We perform latent code optimization with different number of iterations to project test images into the latent space. We then pass these resultant latents to the combined network and obtain predictions to compute the mean IoU. Tab. 3 shows the variation of mIoU with number of iterations of latent code optimization. It can be seen that there is a considerable gap in these scores and the one after the supervised distillation step (Tab. 1 & 2).

Face8 and Face10. We compare with DatasetGAN on Face8 which is a modified version of CelebAMask-HQ dataset with 8 classes instead of 19. For these experiments we do not retrain our model on CelebA. We use the existing model trained on Face34 and merge labels after inference to form the 8 classes from 34, for both GAN and CON.

The Face10 values are in comparison with the ReGAN [41] paper. In this case, we need to retrain our model on CelebA examples modified to 10 classes since one of the classes “cloth” is not a part of the classes in Face34. We use the same 10 train images as used by ReGAN to compare to their 10-shot setting. We keep our network architectures same, whereas to replicate ReGAN’s numbers we use the architectures and hyperparameters as specified in their paper. We use weighted IoU for evaluation following their paper. Tab. 2 shows that ReGAN performs better in the direct step as compared to after distillation. This may be attributed

Training set	Cars	Faces
GAN generated images	0.4266	0.3868
Real images	0.6143	0.5162

Table 6. mIoU scores on Face34 and Car20 for MoCo trained with images generated from StyleGAN vs from real dataset.

to the fact that StyleGAN2 works better with latent code optimization. Both our methods are able to outperform ReGAN after distillation. Our CONV variant performs better in the direct phase too.

5. Results and Discussion

In this section we present comparison of our quantitative performance with prior art, discuss on latent code optimization, perform ablation over various training settings and analyze the qualitative results of our various methods.

5.1. Performance comparison with other methods

Tab. 1 compares the performance of our best method after distillation with transfer learning, semi-supervised learning [28] and DatasetGAN baselines. The transfer network was pre-trained on the semantic segmentation task of MSCOCO [24]. Our method performs better than the baseline methods on all the datasets.

Tab. 2 presents a detailed comparison of our method with the GAN methods. For Face34 we also compare with weighted mIoU, the metric used in ReGAN. Distillation improves performance in all cases, the highest jump can be seen in the Car20 dataset. Except ReGAN, in all cases the direct evaluation of GAN segmentor yields much worse results than after distillation.

We compare with ReGAN in the shaded column with weighted mIoU. In this case the direct method performs bet-

ter than the distilled. This is due to the fact that StyleGAN2 is better able to work with latent code optimization, however the issues of robustness and inferring on unseen images will still remain. StyleGAN2 is also even more computationally expensive as compared to StyleGAN (§3.1.1).

Our variant with CONV as the chosen f_{PROJ} architecture outperforms the MLP method, except for the Face8 setting. This method achieves better mIoU scores than the GAN baseline both with and without distillation.

5.2. Difficulty of latent code optimization

We analyze the performance of direct evaluation of DatasetGAN’s StyleGAN + Style Interpreter network on unseen images from the Face34 test set. We vary the number of iterations of latent code optimization to find corresponding latents for the StyleGAN network using the input test images. In Tab. 3 we show the variation of mIoU with iterations. The score does not increase after 0.4302 at 10k iterations. There is a considerable gap between this number and the score of 0.5365 mIoU obtained by DatasetGAN after distillation. This is due to the fact that latent code optimization is not able to produce images as close to the input image as is required for reasonable performance in fine-grained segmentation.

5.3. Training of Backbone network

In this section we ablate on various training strategies for our backbone network f_{CON} . For all the following experiments, we keep the architecture fixed to that of ResNet50.

MoCo vs SimSiam. We compare contrastive learning frameworks which train with (MoCoV2) and without (SimSiam) negative examples. Tab. 4 compares the scores of each with the MLP model as f_{PROJ} both before and after distillation. MoCoV2 is able to achieve better performance on the few-shot segmentation task than SimSiam. However, SimSiam needs half as many epochs (400 epochs) to reach the reported score compared to MoCoV2 (800 epochs).

Comparison with Imagenet1K pretrained. The first column in Tab. 5 shows the mIoU on Face34 dataset for backbone network f_{CON} initialized using Imagenet1K [23] classification weights. The performance boost that MoCo pre-training offers is significant.

MoCo with different real datasets. The rest of the columns of Tab. 5 compares the effect of the unlabeled dataset. We obtain the best performance using CelebA since the test images belong to CelebA’s test set. FFHQ allows better performance than ImageNet because it is comprised solely of faces. However it is interesting to note that even though ImageNet dataset has no face images, the segmentor using MoCo trained on it still obtains good scores, comparable to FFHQ/CelebA.

MoCo training on GAN generated images. To test the bias of generative models, we train MoCo on images gener-

ated by StyleGAN. We use the same resolution and number of unlabeled images for both training runs. Tab. 6 lists the scores on Face34 and Car20. There is a considerable drop in mIoU for the MoCo trained on images generated by StyleGAN. The MoCo training on generated images also reaches high accuracy much earlier ($\sim 90\%$ at 200 epochs). This indicates that GANs are not able to model a large domain i.e they are not able to learn a general distribution. This might also mean that GAN generated images have some typical artefacts that CON models are able to identify.

5.4. Qualitative Results

Fig. 3 presents a qualitative comparison of the strengths and weaknesses of our methods. In image I, because of the black background both the direct MLP and CONV models struggle with segmenting the hair, though the CONV does a better job. In this case distillation helps improve the predicted mask. For images II and III, the CONV + distill variants are able to capture the overall structure better, whereas the MLP + distill methods are able to differentiate smaller regions such as wrinkles and chin. In image IV, the MLP model confuses a large percentage of the background as neck/forehead. The CONV method is able to mitigate this problem as it can take more global cues owing to its larger receptive field (see Appendix B). In images V and VII, the CONV + distill approach is able to eliminate confusion between back and front bumpers which is present in the predictions of the other methods. In image VI, the MLP model makes some wrong predictions. The CONV model is able to clear up most of the confusion, while the CONV + distill model boosts this performance and minimizes faulty predictions. In image VIII, we show a failure case. Predictions on the cats data has suffered due to faulty annotations in the dataset and the large variation of pose and occlusion making it challenging to learn representations.

6. Conclusion

We present a contrastive learning approach for few-shot fine-grained part segmentation. We compare to previous generative approaches and show using various experiments that our method achieves better performance, is simpler, an order-of-magnitude faster and is more robust. GAN generated datasets while effective on the part-segmentation task, might not be as effective on tasks such as fine-grained image classification due to biases they introduce. Further investigation of the nature of these biases could better guide the community adoption of GANs as a replacement of datasets.

Acknowledgements The project is supported in part by Grant #1749833 from the National Science Foundation of United States. Our experiments were performed on the University of Massachusetts Amherst GPU cluster obtained under the Collaborative Fund managed by the Mass. Technology Collaborative.

References

- [1] Philip Bachman, R Devon Hjelm, and William Buchwalter. Learning representations by maximizing mutual information across views. In *NeurIPS*, 2019. 3
- [2] Victor Besnier, Himalaya Jain, Andrei Bursuc, Matthieu Cord, and Patrick Pérez. This dataset does not exist: training models from generated images. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5. IEEE, 2020. 2
- [3] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018. 1, 2
- [4] Liang-Chieh Chen, George Papandreou, Florian Schroff, and Hartwig Adam. Rethinking atrous convolution for semantic image segmentation. *arXiv preprint arXiv:1706.05587*, 2017. 6
- [5] Ting Chen, Simon Kornblith, Mohammad Norouzi, and Geoffrey Hinton. A simple framework for contrastive learning of visual representations. 2020. 1, 2, 3
- [6] Ting Chen, Simon Kornblith, Kevin Swersky, Mohammad Norouzi, and Geoffrey Hinton. Big self-supervised models are strong semi-supervised learners. *NeurIPS*, 2020. 3
- [7] Xinlei Chen, Haoqi Fan, Ross Girshick, and Kaiming He. Improved baselines with momentum contrastive learning. *arXiv preprint arXiv:2003.04297*, 2020. 3, 6
- [8] Xinlei Chen and Kaiming He. Exploring simple siamese representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 15750–15758, 2021. 1, 3, 6
- [9] Zezhou Cheng, Jong-Chyi Su, and Subhransu Maji. On equivariant and invariant learning of object landmark representations. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 9897–9906, 2021. 3
- [10] Alexey Dosovitskiy, Jost Tobias Springenberg, Martin Riedmiller, and Thomas Brox. Discriminative unsupervised feature learning with convolutional neural networks. In *NeurIPS*, 2014. 2, 3
- [11] Spyros Gidaris, Praveer Singh, and Nikos Komodakis. Unsupervised representation learning by predicting image rotations. In *ICLR*, 2018. 3
- [12] 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. 2
- [13] Bharath Hariharan, Pablo Arbeláez, Ross Girshick, and Jitendra Malik. Hypercolumns for object segmentation and fine-grained localization. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pages 447–456, 2015. 2, 3
- [14] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *CVPR*, 2020. 1, 2, 3
- [15] Kaiming He, Haoqi Fan, Yuxin Wu, Saining Xie, and Ross Girshick. Momentum contrast for unsupervised visual representation learning. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 9729–9738, 2020. 3
- [16] 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. 6
- [17] R Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Phil Bachman, Adam Trischler, and Yoshua Bengio. Learning deep representations by mutual information estimation and maximization. In *ICLR*, 2019. 3
- [18] Tommi S Jaakkola, David Haussler, et al. Exploiting generative models in discriminative classifiers. *Advances in neural information processing systems*, pages 487–493, 1999. 2
- [19] Ali Jahanian, Xavier Puig, Yonglong Tian, and Phillip Isola. Generative models as a data source for multiview representation learning. *arXiv preprint arXiv:2106.05258*, 2021. 2
- [20] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4401–4410, 2019. 1, 2, 4, 5
- [21] Tero Karras, Samuli Laine, Miika Aittala, Janne Hellsten, Jaakko Lehtinen, and Timo Aila. Analyzing and improving the image quality of stylegan. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 8110–8119, 2020. 5
- [22] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013. 2
- [23] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. Imagenet classification with deep convolutional neural networks. *Advances in neural information processing systems*, 25:1097–1105, 2012. 8
- [24] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European conference on computer vision*, pages 740–755. Springer, 2014. 7
- [25] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *Proceedings of the IEEE international conference on computer vision*, pages 3730–3738, 2015. 5
- [26] David G Lowe. Object recognition from local scale-invariant features. In *Proceedings of the seventh IEEE international conference on computer vision*, volume 2, pages 1150–1157. Ieee, 1999. 2
- [27] Geoffrey J McLachlan and Kaye E Basford. *Mixture models: Inference and applications to clustering*, volume 38. M. Dekker New York, 1988. 2
- [28] Sudhanshu Mittal, Maxim Tatarchenko, and Thomas Brox. Semi-supervised semantic segmentation with high-and low-level consistency. *IEEE transactions on pattern analysis and machine intelligence*, 2019. 4, 7
- [29] Mehdi Noroozi and Paolo Favaro. Unsupervised learning of visual representations by solving jigsaw puzzles. In *ECCV*, 2016. 3
- [30] Timo Ojala, Matti Pietikainen, and David Harwood. Performance evaluation of texture measures with classification based on kullback discrimination of distributions. In *Proceedings of 12th international conference on pattern recognition*, volume 1, pages 582–585. IEEE, 1994. 2

- [31] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. Representation learning with contrastive predictive coding. *arXiv preprint arXiv:1807.03748*, 2018. 3
- [32] Deepak Pathak, Philipp Krahenbuhl, Jeff Donahue, Trevor Darrell, and Alexei A Efros. Context encoders: Feature learning by inpainting. In *CVPR*, 2016. 3
- [33] Pedro O Pinheiro, Amjad Almahairi, Ryan Y Benmalek, Florian Golemo, and Aaron Courville. Unsupervised learning of dense visual representations. *arXiv preprint arXiv:2011.05499*, 2020. 3
- [34] 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. 4
- [35] Jorge Sánchez, Florent Perronnin, Thomas Mensink, and Jakob Verbeek. Image classification with the fisher vector: Theory and practice. *International journal of computer vision*, 105(3):222–245, 2013. 2
- [36] Karen Simonyan, Omkar M Parkhi, Andrea Vedaldi, and Andrew Zisserman. Fisher vector faces in the wild. In *BMVC*, volume 2, page 4, 2013. 2
- [37] Paul Smolensky. Information processing in dynamical systems: Foundations of harmony theory. Technical report, Colorado Univ at Boulder Dept of Computer Science, 1986. 2
- [38] Fabio Henrique Kiyoyiti dos Santos Tanaka and Claus Aranha. Data augmentation using gans. *arXiv preprint arXiv:1904.09135*, 2019. 2
- [39] James Thewlis, Samuel Albanie, Hakan Bilen, and Andrea Vedaldi. Unsupervised learning of landmarks by descriptor vector exchange. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*, pages 6361–6371, 2019. 3
- [40] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. *ECCV*, 2020. 3
- [41] Nontawat Tritrong, Pitchaporn Rewatbowornwong, and Supasorn Suwajanakorn. Repurposing gans for one-shot semantic part segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 4475–4485, 2021. 1, 2, 3, 7
- [42] Xinlong Wang, Rufeng Zhang, Chunhua Shen, Tao Kong, and Lei Li. Dense contrastive learning for self-supervised visual pre-training. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 3024–3033, 2021. 3
- [43] Zhirong Wu, Yuanjun Xiong, Stella Yu, and Dahua Lin. Unsupervised feature learning via non-parametric instance discrimination. In *CVPR*, 2018. 3
- [44] Fisher Yu, Ari Seff, Yinda Zhang, Shuran Song, Thomas Funkhouser, and Jianxiong Xiao. Lsun: Construction of a large-scale image dataset using deep learning with humans in the loop. *arXiv preprint arXiv:1506.03365*, 2015. 6
- [45] Richard Zhang, Phillip Isola, and Alexei A Efros. Colorful image colorization. In *ECCV*, 2016. 3
- [46] Yuxuan Zhang, Huan Ling, Jun Gao, Kangxue Yin, Jean-Francois Lafleche, Adela Barriuso, Antonio Torralba, and Sanja Fidler. Datasetgan: Efficient labeled data factory with minimal human effort. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pages 10145–10155, 2021. 1, 2, 3, 4, 5