

Improving the Effectiveness of Deep Generative Data

Ruyu Wang^{1,2} Sabrina Schmedding¹ Marco F. Huber^{2,3}

¹Bosch Center for Artificial Intelligence, Renningen, Germany

²Institute of Industrial Manufacturing and Management IFF, University of Stuttgart, Stuttgart, Germany

³Fraunhofer Institute for Manufacturing Engineering and Automation (IPA), Stuttgart, Germany

{ruyu.wang, sabrina.schmedding}@de.bosch.com marco.huber@ieee.org

Abstract

Recent deep generative models (DGMs) such as generative adversarial networks (GANs) and diffusion probabilistic models (DPMs) have shown their impressive ability in generating high-fidelity photorealistic images. Although looking appealing to human eyes, training a model on purely synthetic images for downstream image processing tasks like image classification often results in an undesired performance drop compared to training on real data. Previous works have demonstrated that enhancing a real dataset with synthetic images from DGMs can be beneficial. However, the improvements were subjected to certain circumstances and yet were not comparable to adding the same number of real images. In this work, we propose a new taxonomy to describe factors contributing to this commonly observed phenomenon and investigate it on the popular CIFAR-10 dataset. We hypothesize that the Content Gap accounts for a large portion of the performance drop when using synthetic images from DGM and propose strategies to better utilize them in downstream tasks. Extensive experiments on multiple datasets showcase that our method outperforms baselines on downstream classification tasks both in case of training on synthetic only (Synthetic-to-Real) and training on a mix of real and synthetic data (Data Augmentation), particularly in the data-scarce scenario.

1. Introduction

Over the past decade, DGMs powered by large-scale datasets have demonstrated revolutionary results in the field of image synthesis. Models based on different working theories such as the variational autoencoder (VAE) [28] and GANs [16] have shown astonishing performance in generating realistic-looking images. More recently, diffusion-based models [11, 21] have emerged as the new state-of-the-art family in this challenging task and have broken the long-time dominance of the GAN family. The rapidly improved image quality and mode coverage of generative models

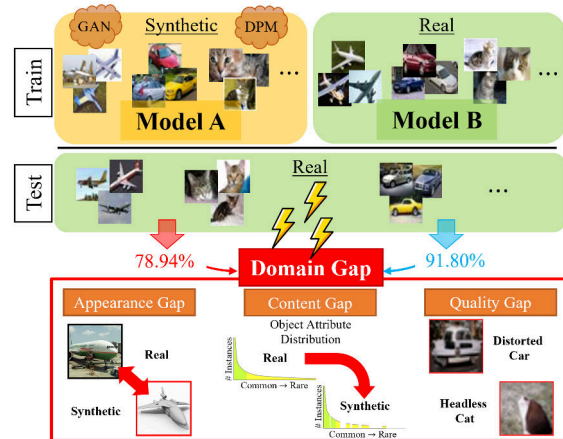


Figure 1. An illustration of the commonly observed performance degradation when training a model on synthetic data and testing it on real data. Such effects can be attributed to the *Domain Gap*, which we split into three aspects (appearance, content, quality).

leaves us wondering: to which degree is synthetic data from these models ready to replace real samples for image recognition tasks, especially in the case of data scarcity?

Despite the inevitable performance degradation in low data regimes, many researchers have tried to adapt generative models to various domains where real data is hard and time-consuming to acquire (e.g., defective industrial products) [37, 56, 63]. In such cases, images from generative models serve as a kind of data augmentation. For example, by navigating in the learned latent space, most of the generative models can deliver images with novel semantic information according to the guidance [9, 23, 39, 44, 58]. However, the usefulness of synthetic data is highly dependent on use cases as discovered in many works [24, 45, 56, 63, 64]. Tasks such as pose estimation [24] benefit more from synthetic data than object recognition tasks. However, it has been widely observed that in most cases the effectiveness of synthetic images is not on par with that of real ones, and the boost brought by them not only saturates much faster but also diminishes rapidly once more real data are avail-

able [3, 19, 46]. Prior works generally attribute the cause of these effects to the *Domain Gap*.

In this paper, we go beyond the general description of the *Domain Gap* and define a new taxonomy to detail its possible factors into the *Appearance Gap*, the *Content Gap*, and the *Quality Gap* as illustrated in Fig. 1. Then, we endeavor to answer the following questions:

- Does the synthetic data generated by DGMs introduce novel information (rendering DGMs a semantically meaningful data augmentation method), or is its utility only equivalent to resampling a subset of the real dataset the DGM was trained on?
- Which components of the *Domain Gap* contribute most to cases with poor performance?

To address these questions, a novel set of investigations were conducted to examine the effectiveness of synthetic CIFAR-10 [29] images from two popular DGMs [48, 54]. Our results strongly suggest that the *Content Gap* has the highest impact among the factors on the utility of synthetic images in downstream tasks. Inspired by these findings, we further propose two remedies—Pretrained Guidance and Real Guidance—, aiming to boost the effectiveness of synthetic data despite its current shortcomings in downstream recognition tasks.

Extensive results on multiple datasets demonstrate that the proposed Pretrained Guidance, where we use a large-scale pretrained model (e.g., pretrained on ImageNet [10]) as a prior to regularize the distance between features during downstream classifier training, surpass other synthetic-to-real generalization methods. Moreover, together with Real Guidance to mitigate the negative effect of synthetic data, our method outperforms all baselines when mixing with real data, especially in the low data regimes.

2. Related Work

Domain adaptation [14, 51, 53] and generalization [31] have been important research topics in recent studies on deep learning, where the main aim is to recover performance drops caused by the domain gap. Among their sub-disciplines, synthetic-to-real adaption and generalization [13, 40, 42, 52] have been promising research directions for domains where acquiring large amounts of data or labels is either hard or time-consuming (e.g., industrial, and medical). Early works mainly focused on synthetic images from physically based rendering (PBR), where the simulator aims to closely match the physical reality and achieve photo-realism. However, models trained on such kinds of synthetic images often generalize poorly on the target real data despite the relatively low cost of data collection. Many researchers have proposed different methods to mitigate the performance drop caused by the *Appearance Gap*

(e.g., unrealistic textures or over-simplified lighting conditions) such as by utilizing adversarial learning [34, 61] or meta-learning [32]. Lately, Chen et al. proposed to leverage the ImageNet pretrained weights as prior knowledge from real domain and combine it with a continual learning scheme [8] or a contrastive learning framework [7] to implicitly align learned features to their ImageNet counterparts. Unlike their works, we do not impose such feature similarity and instead relax the constraint by encouraging the similarity between the distance distribution of features.

More recently, DGMs have emerged as a new alternative to synthetic data generation [5, 11, 21, 26, 48, 54, 60]. They come with several favorable features over PBR: 1) superior photorealism from training on real-world data; 2) requiring less storage space; 3) applicable for sophisticated cases that cannot be modeled by PBR (e.g., defects in production). Few earlier works [4, 24, 45, 64] have attempted to explore the synthetic data from DGMs for image recognition tasks but mainly focused on GAN-based methods. With the recent bloom of diffusion-based models, some researchers have reported their investigation [19, 46] regarding synthetic data from the state-of-the-art text-to-image generation models and revealed the potential of zero-shot generation, i.e., without seeing any real target data.

However, it remains unclear how the synthetic data from DGMs trained on small-scale datasets (e.g., 100 or 1,000 samples) would affect the downstream recognition tasks. Despite the appealing ability of zero-shot generation from popular text-to-image models, it is non-trivial to adapt it to some specific domains like industrial products, which are critically under-represented online. The plausible solutions therefore lie in low-/few-shot adaption [15, 33, 43, 49, 59, 65] of DGMs. Beckham et al. [3] designed a shallow generative network to study the effectiveness of synthetic data in the few-shot scenario, where they concluded that the benefit of synthetic data was rather subtle in this case. This phenomenon of synthetic data being less effective than real data has been widely observed in previous works [4, 24, 45, 64] and only becomes more pronounced in the low data regime. For scenery datasets, researchers [1, 2] have attributed the drop in performance to a distribution shift in the image content and reported that DGMs tend to omit rare or hard components (e.g., humans). We would consider this one aspect of the *Content Gap* and investigate it for object-centered images. Unlike scenery images, the loss of a mode representing an object with plausible but rare attributes (e.g., a cat in the water) might be easily overlooked. In this work, we study the impact of such mode drop on downstream classification tasks and propose remedies based on our findings.

3. Empirical Investigations

In this section, we first introduce a new taxonomy to detail the potential factors for the *Domain Gap* and then

present an empirical investigation of the two questions stated in Sec. 1, linking the factors of *Domain Gap* to the effectiveness of synthetic data.

3.1. Domain Gap Factors

Researchers often refer to the negative effects, which happen when the data distribution of training and testing sets are different, as *Domain Gap*. However, it remains unclear what exactly are the factors leading to the distribution shift. To dig deeper into multiple constituents of the domain gap, we classify the potential factors into three categories—the *Appearance Gap*, the *Content Gap*, and the *Quality Gap*. The *Appearance Gap* refers to the performance gap induced by artifacts such as unrealistic textures, which is more often observed in the synthetic data acquired from PBR [41]. The *Content Gap* addresses the distribution shift in the composition of a generated image. For example, Bau et al. [2] studied the typical failure cases of GANs by examining the deviation between a specific scene image and its reconstruction, concluding that the models tend to skip difficult subtasks like large human figures. Finally, the *Quality Gap* is stated for the structural distortion in the synthetic images such as two-headed fish or faceless humans.

3.2. Study on Synthetic Datasets

Our exploration on effects of synthetic data was carried out by first assuming full access to a widely used benchmark dataset in image synthesis—CIFAR-10, and checkpoints of popular DGMs trained on CIFAR-10. Then, we sampled synthetic clones of CIFAR-10 from the selected DGMs and analyzed the behavior of classifiers trained on the original and synthetic sets of CIFAR-10 for object-centered image classification. Analogous experiments and effects on ImageNet can be found in Appendix B.4.

Experiment setup. We selected two popular DGMs with state-of-the-art FID [20] at the time—LSGM [54] and StyleGAN-XL (SG-XL) [48]—and sampled images directly from the provided checkpoints by the authors. LSGM is limited to unconditional image synthesis. We therefore leveraged an off-the-shelf Wide Residual Network (WRN) [62] (Top-1 Accuracy on CIFAR-10 classification: 96.21%¹) during the sampling process to obtain labels for the synthetic dataset (with predicted softmax probability > 0.8). To make a fair comparison, we also applied the same filtering procedure on the conditional SG-XL.

We present the results under the setting that all classifiers were initialized with pretrained ImageNet weights and the learning rate was set to 0.0001 (See Appendix B.2 for random initialization). Although the performance drops are extensively observed in different network structures (e.g., EfficientNet [50], ViT [12]), we chose the widely deployed ResNet-50 [18] as backbone for demonstration. We set the

¹<https://github.com/xinntao/pytorch-classification-1>

Table 1. The achieved accuracy of classifiers trained and tested on different sources of images. Note that the results were all acquired from the training sets, therefore the high accuracy in the diagonal line indicates that the classifiers have converged on their own training set.

Source of Test Images	Real	LSGM	SG-XL
Classifier Trained on Real	96.34	93.54	98.68
Classifier Trained on LSGM Images	85.76	97.71	99.11
Classifier Trained on SG-XL Images	78.90	91.14	99.82




Figure 2. The training and validation curves of CIFAR-10 images from different sources over five runs. Standard deviations are plotted as shaded areas. (Zoom in.)

image resolution to 224×224 , the batch size to 128 and trained the models for 200 epochs. To observe the full effect of synthetic data, only random crop and random flip were applied as data augmentation. We evaluated all models on the official CIFAR-10 test set and reported the averaged accuracy over five random runs. Note that we employed a fixed train-val split for the real dataset throughout the paper—500 samples from each class were randomly selected in the beginning as the validation set and the remaining 45,000 images formed the training set. For simplicity, we later denote real images as *Real* and synthetic images from DGMs as *Synthetic*.

Observation #1: The non-mutual performance gap.

It has been widely observed that there exists a performance gap between real and synthetic data. A commonly accepted hypothesis attributes this behavior to the domain gap between the two distributions. Based on this assumption, one would presume that the domain gap between synthetic data and real data is mutual, meaning that the performance drop in the case of synthetic-to-real should be at a similar scale to the opposite case of real-to-synthetic. However, this presumption does not hold as seen in Tab. 1. We can see that the performance of the classifier trained on *Real* is not affected too much when classifying synthetic images, while the classifiers trained on *Synthetic* sources (e.g., SG-XL) exhibit drastic degradation in accuracy. Moreover, SG-XL shows a more severe decrease than LSGM, presumably due to diffu-

sion models being better in mode coverage than GANs.

Observation #2: The training accuracy saturates quickly when training on *Synthetic* sources. We show the training and validation accuracy curves of all image sources in Fig. 2. It can be seen that the training accuracy of *Synthetic* sources grew much faster in the early epochs (< 25) and saturated faster compared to the classifier trained on the *Real*. While the performance gap between the training and validation can be explained by the distribution shift, the peculiar behavior of the training curves of *Synthetic* was puzzling—all the classifiers were tasked with the same objective (i.e., classify the images into ten classes), why would the classifiers trained on the *Synthetic* sources seemed to be solving an easier task? We hypothesize that this behavior can be attributed to the *Content Gap* between synthetic and real data, where the compositions of an object and its rare but possible attributes (e.g., a car in the sea) were either dropped by the models or barely generated. Therefore, the resulting synthetic dataset only contains images associated with frequent attributes, forming a less diverse and simplified training set. This led to the next question: do the current synthetic images from the DGMs just reflect a subset of the real data?

Observation #3: Very small losses are observed for *Synthetic* samples. To validate the hypothesis that synthetic images only cover a subspace of the real dataset, we further examine the training samples from all sources. We propose to measure the importance of a sample by the cross entropy loss it brings to a trained classifier of an opposite source. This means we evaluate all the samples from the *Synthetic* sources by a classifier trained on *Real* samples and vice versa. The concept behind it is that if an incoming sample contains novel information which was not captured during the training process, the loss it simulates should be high. Conversely, if the information a sample can bring is already learned by the model, its loss should be close to zero.

We plot the measured losses in Fig. 3. In each subgraph, the solid bars present the loss distribution in the case where the classifier was trained on *Real* data and used to evaluate *Synthetic* images. On the contrary, the dotted bars show the case of training a classifier on *Synthetic* and test on *Real* images. It can be observed that the synthetic samples (Solid) overall have a lower loss compared to the real ones (Dotted). Compared to the losses brought by the real samples (Dotted), much more synthetic samples (Solid) have a close-to-zero loss and only a small portion possess a value bigger than 1 (See more low- and high-loss images in Appendix B.3). Interestingly, we still observe that utilizing the low-loss synthetic samples for data augmentation is beneficial for downstream performance (see Appendix B.1).

3.3. Conclusion

All three observations emphasize that despite training on a relatively large dataset and having promising FID, images from DGMs still suffer from the domain gap. While the *Appearance Gap* is eliminated by learning directly on the real data, it was unclear whether the synthetic images in our investigations suffered from *Quality* or *Content Gap*. The following two details however suggest that the *Content Gap* is the main issue: (1) The accuracy of *Real* classifiers in **Observation #1** does not show signs of obstruction, as would be assumed from distorted images. (2) The distorted data should also simulate a high loss in **Observation #3**, yet the loss from synthetic samples is generally lower than that of the real ones.

We thus conclude that sampling directly from current DGMs will result in a subset of the real data due to the *Content Gap*, which worsens the long-tailed nature that most real-world data exhibit—rare samples are either omitted by the generative models or barely ever generated. Therefore, downstream feature extractors trained exclusively on such misrepresented distributions tend to show undesirable behavior (e.g., performance drop in classification accuracy) when tested on real data. However, at the end of **Observation #3**, we observe that even adding synthetic samples with lower loss to the real training set still has minor positive effects. We hypothesize that this is because the synthetic images from DGMs are lacking rare samples but do add variations within the covered parts of the training distribution. As a result, we introduce two strategies to boost their effectiveness in the following section.

4. Remedies

The *Content Gap* can also be cast as a mode coverage problem: sampling synthetic data most likely leads to a less diverse and simplified dataset, where rare samples are not represented. The information such a simplified dataset can provide is therefore mostly covered by the original dataset. Unsurprisingly, training recognition models on such a dataset can easily result in degraded performance and factitious-biased representation. One potential way to acquire synthetic data with higher information density is to apply additional techniques during the sampling process [17, 22, 57]. However, this only works under the assumption that the rare samples are still learned by the DGMs and are just rarely generated. The assumption might be violated when the DGMs are trained on small datasets and the rare cases in the training set are omitted by the models. This is a practically relevant task however, because it is often impossible to acquire large-scale data for training a well-behaved DGM from scratch in most real-world scenarios (e.g., industrial). As a result, we propose two novel regularization strategies to improve the effective-

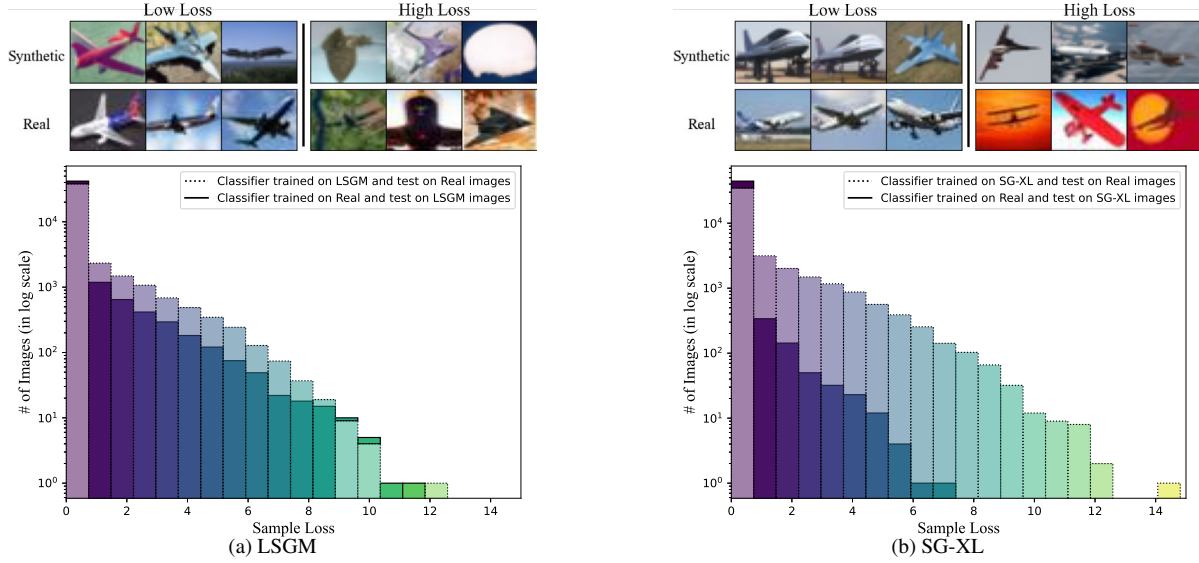


Figure 3. The loss distribution of the samples from different sources. We overlapped the evaluation of *Real* \rightarrow *Synthetic* (Solid) and *Synthetic* \rightarrow *Real* (Dotted) in each subgraph and visualized the associated low-loss and high-loss samples of class Airplane. Note that all models used for evaluation were initialized with ImageNet weights.

ness of the synthetic data even in low data regimes, termed as *Pretrained Guidance* and *Real Guidance*.

Pretrained Guidance (PG). Models pretrained on large-scale datasets like ImageNet often exhibit diverse and rich representations that can be transferred to other tasks. We propose exploiting this trait as external guidance to prevent the models trained on the synthetic dataset from converging into having less distinctive representations. The closest previous works [7, 8] suggested tackling this problem by forcing the new model to have a similar probability distribution to the pretrained ImageNet predictions, implicitly regularizing the new feature extractor to stay close to the ImageNet. But such constraints are not always beneficial and their performance depends heavily on how the model is initialized as we later show in Sec. 5. In contrast, our proposed PG allows the feature extractor to learn freely from the new dataset while maintaining a similar span of data as ImageNet.

The workflow of our PG is illustrated in Fig. 4(a). We denote the model being trained as M_u and utilize a frozen pretrained model M_p (note that M_p can be different to M_u). In each forward pass, a synthetic sample will be given to both networks and obtain the feature representations f_p and f_u , respectively. For a batch of N images, the distance matrices $D_p \in \mathbb{R}^{N \times N}$ and $D_u \in \mathbb{R}^{N \times N}$ are computed with respect to f_p and f_u using a metric δ . We then propose to encourage D_p to possess a similar distance distribution as D_u . In this way, f_u is allowed to learn the representations for the target task freely, while the guidance from the pretrained model helps preserve the span of features. We

formulate the proposed PG regularization as

$$D_p^{ij} = \delta(f_p(i), f_p(j)), \quad D_u^{ij} = \delta(f_u(i), f_u(j)), \quad (1)$$

$$\mathcal{L}_{pg} = \sum_i \sum_j \text{sim} \left\{ D_p^{ij}, D_u^{ij} \right\},$$

where i, j indicate different samples in a batch, $\delta(\cdot)$ stands for any distance metric like the cosine similarity, and $\text{sim}\{\cdot\}$ can be any similarity metric such as L1 or KL-divergence². Note that the effect of PG is independent of initializing the model with pretrained weights because despite providing a good starting point, the initialization does not stop the model from converging into having less distinctive representations.

Real Guidance (RG). Inspired by replay-based continual learning (CL) approaches [6], we adapt the idea of gradient episodic memory to our context to mitigate the domain gap. In contrast to alleviating forgetting, we utilize it as a way to regularize the gradient flow of the synthetic data. Specifically, we assume that a small set of real data (e.g., 10 or 100 samples) is available at training time. It is a more realistic scenario than assuming the sole availability of synthetic data, regarding it would usually require a handful of real data to train or finetune the DGMs in the beginning.

Considering the gradients g_r from real data as a positive influence on the model, the gradients g_f from synthetic data can however be a confounder if the angle between g_r and g_f are more than 90° . We then channel the gradients from the synthetic data so that they are at most perpendicular to the real data, avoiding the misleading signal brought by g_f .

²KL-divergence on distance matrices as derived in [38]

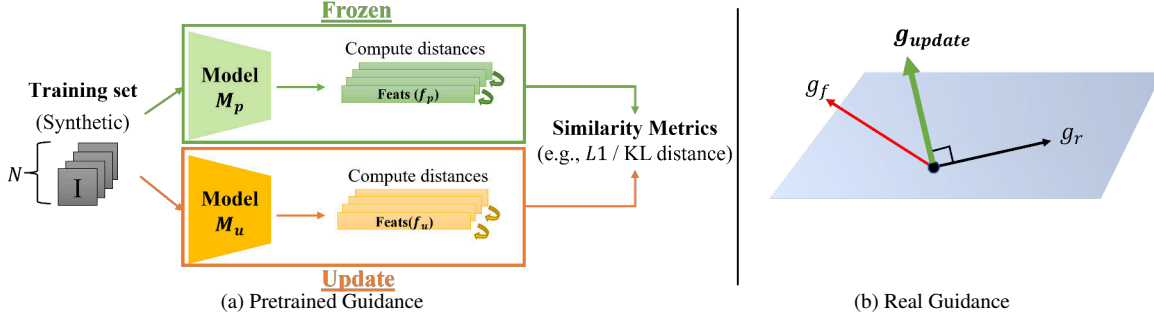


Figure 4. The proposed **Pretrained Guidance** and **Real Guidance**.

This can be formulated as

$$g_{f_{\text{new}}} = \begin{cases} g_f - \frac{g_f^\top g_r}{g_r^\top g_r} \cdot g_r & , \text{if } g_f^\top g_r < 0 \\ g_f & , \text{otherwise} \end{cases}, \quad (2)$$

$$g_{\text{update}} = \lambda_1 \cdot g_r + \lambda_2 \cdot g_{f_{\text{new}}}.$$

In this way, we prevent the synthetic data from overfitting itself while narrowing the domain gap by introducing information from the real data.

In summation, the final objective for a target task is

$$\mathcal{L} = \mathcal{L}_{\text{Task}} + \lambda_3 \cdot \mathcal{L}_{\text{pg}}, \quad (3)$$

where $\mathcal{L}_{\text{Task}}$ can be a cross-entropy loss in case of classification and $\lambda_1, \lambda_2, \lambda_3$ are the hyperparameters to tune the intensity of our proposed regularizations.

5. Experiments

To study the effectiveness of synthetic data from popular DGMs and to evaluate the usefulness of our proposed remedies against the expected content gap, we consider two cases: synthetic-to-real generation and synthetic data as data augmentation. The first case represents a scenario in which no real data is available for downstream tasks, which we also refer to as the **Zero-shot** task. Although it is a theoretical setting ignoring the real samples that were required to train a DGM, the results from this case can serve as a quantitative assessment of the discrepancy between synthetic and real data. The second case presents a scenario assuming at least a handful of real images (e.g., 100 or 1,000) are available, which we describe as the **Low-shot** task. It is more common in real-life settings (e.g., industrial), otherwise training or finetuning a DGM as well as evaluating in downstream tasks would be generally infeasible.

Datasets. Besides the three sets of CIFAR-10 (*Real*, *LSGM* and *SG-XL*) mentioned in Sec. 3, synthetic images of two popular benchmark datasets (CUB Bird [55], Oxford Flower [36]) and one industrial dataset (SDI [56]³) from

³We thank the authors [56] for providing the dataset for our benchmark.

three generative models⁴ (SG-XL, Projected GAN (ProjGAN) [47], and DT-GAN [56]) were used for evaluations. For the first two datasets, we sampled images to meet the size of the original datasets. For the SDI dataset, we sampled images to balance the original SDI dataset and thus, resulted in three sets of synthetic data with size 8,000. Note that due to the mode-dropping issue when training on small datasets, we conducted all classification experiments on a subset of CUB Bird and Oxford Flower. Details on generative model training and sampling are in Appendix A.1.

Experiment Setup. We deployed the same setup for training ResNet-50 classifiers as described in Sec. 3. All following experiments were conducted under the same setting unless otherwise specified. The reported accuracies were averaged over five random runs. We used default hyperparameters from the authors for selected baselines and set λ_1, λ_2 and λ_3 of our method differently for each dataset. See Appendix A.2 for the chosen hyperparameters and Appendix C for more results in different settings.

5.1. Zero-shot Image Classification

We investigate the zero-shot classification performance of the synthetic data from various models. As in Tab. 2, simply training the classifiers on synthetic images leads to a degraded performance compared to the classifiers trained on the same amount of real data. However, has the current synthetic data reached its full potential? We argue that the drop in performance can be reduced based on our insights from Sec. 3 and the proposed PG for regularization as introduced in Sec. 4. It can be seen that methods like [7, 8] highly depend on the initialization of the model. Also, their regularizations are not always beneficial. For example, they often leads to worse results when the number of samples in each class is relatively high (e.g., CIFAR-10). For datasets that do not benefit from the ImageNet initialization even for real data (e.g., industrial products like SDI-A and SDI-C in Tab. 2), these methods often fail to deliver promising results. In contrast, our PG provides better performance in

⁴Diffusion-based models are omitted due to their low performance in the low data regimes. See Appendix A.1 for more information.

Table 2. The achieved accuracy of classifiers trained on different synthetic datasets. (A) The networks were trained from scratch. (B) The networks were initialized by pretrained ImageNet weights. Note that for a fair comparison, we omit the second stage of ASG [8], which uses reinforcement learning for selective parameter update.

	Dataset Size	CIFAR-10 45,000		CUB 4,521		Flower 1,010		SDI-A 8,000	SDI-B 8,000	SDI-C 8,000
		Source	LSGM	SG-XL	Proj-GAN	SG-XL	Proj-GAN	SG-XL	DT-GAN	
(A)	Baseline	85.08	72.22	12.30	10.80	13.54	13.07	64.72	86.80	61.14
	ASG [8]	80.66	71.68	1.64	1.83	5.32	5.12	32.78	69.40	31.43
	CSG [7]	82.47	72.94	1.41	1.86	8.21	8.04	43.94	74.55	32.19
	Ours (L1)	85.58	75.07	11.93	13.94	13.89	13.69	64.73	87.20	64.00
	Ours (KL)	84.76	77.45	17.11	13.45	17.16	15.53	66.36	88.40	62.52
	Real		90.27		18.82		35.59		83.20	87.8
(B)	Baseline	85.57	78.94	33.15	27.20	32.40	28.42	53.64	74.00	52.95
	ASG [8]	82.40	77.86	43.62	39.90	5.63	5.47	67.78	87.95	50.47
	CSG [7]	82.67	79.01	43.62	38.81	5.45	4.33	62.33	84.45	50.86
	Ours (L1)	85.74	79.87	32.89	27.32	30.91	26.00	51.63	73.40	53.52
	Ours (KL)	84.73	80.66	34.50	28.93	33.52	30.01	57.64	85.20	53.90
	Real		91.80		43.15		52.12		66.36	90.00

Table 3. The achieved accuracy of classifiers trained on synthetic CIFAR-10 datasets, where a simplified version of ResNet-50 was used for training images at resolution 32.

Method	LSGM	SG-XL
Baseline	78.55	70.68
ASG [8]	76.63	66.31
Ours (L1)	78.33	70.80
Ours (KL)	80.11	74.21
Real	84.53	

most cases considering both kinds of initialization, suggesting its practical usefulness in real-world industrial cases.

Moreover, our PG requires not the same model architecture for the frozen pretrained model M_p and the new model M_u . Therefore, it can also combine pretrained weights from wide-spread architectures to new models where no pretrained weights are available. As an exemplary case, we set M_u to a simplified ResNet-50⁵, where its network structure was customized to work for CIFAR-10 at resolution 32×32 and its ImageNet pretrained weights were not available. Meanwhile, M_p was set to be the original ResNet-50 with pretrained weights from ImageNet. It should be noted that the feature representations f_p and f_u have different dimensions (1024 vs. 256) due to the structural difference. As shown in Tab. 3, our method again outperforms ASG in this case and delivers the best results when using the KL distance. CSG on the other hand is not feasible in such a scenario due to its design requiring M_p and M_u to have identical structures.

The flexibility of our method also allows choosing the distance and similarity metrics freely. Our results have shown that applying KL distance together with cosine sim-

⁵<https://github.com/kuangliu/pytorch-cifar>

ilarity generally yields better performance in our experiments. All further experiments were therefore conducted with this setting only.

5.2. Low-shot Image Classification

In the low-shot scenario, we investigate our method in the situation where the synthetic data is used as data augmentation to enlarge the real training set. We also compare our results to baseline methods that were designed for domain adaptation, where the synthetic data and real data are used sequentially instead of simultaneously.

We conducted the experiments in the following two settings: (1) Assuming that a large-scale dataset (CIFAR-10) was available for pretraining generative models. For the downstream task, classification in our setting, only very few real samples (< 20) from the large-scale dataset are available. (2) Assuming that a smaller scale dataset (e.g., CUB, Flower, and SDI) is available but has no public accessible pretrained generative models. The training data therefore needs to be used for both to train the generative models from scratch and to train the model for the downstream (classification) task. Note that these settings result in datasets with different proportions of synthetic and real images, which we denote as ‘‘Syn-to-Real Ratio’’.

Theoretically, more settings would be possible such as using pretrained weights from a large-scale dataset for finetuning the generative model on a smaller scale dataset. However, finetuning a DGM requires carefully adaptations and we leave a complete study of all options for future work.

Baselines. For data augmentation, we set the baseline as simply adding the synthetic images to the real training set. In addition, we selected A-GEM [6] as regularization to aid this baseline. Conceptually, this is the closest baseline to our method. For domain adaptation, we chose four domain

Table 4. The achieved accuracy of classifiers trained with different methods under various setting: **RG** indicate Real Guidance, **PG-F** means only apply Pretrained Guidance to the synthetic data, and **PG-R** means only apply Pretrained Guidance to the real data. Note that we denote the baseline two-stage domain adaptation as **Adp.** and baseline data augmentation as **Aug.**.

Dataset	CIFAR-10 (10-shots)		CUB		Flower		SDI-A	SDI-B	SDI-C			
Syn-to-Real Ratio	450:1		1:1		1:1		2:1	2:1	2:1			
Source	RG	PG-R	PG-F	LSGM	SG-XL	Proj-GAN	SG-XL	Proj-GAN	SG-XL	DT-GAN		
Baseline (Real)	-	-	-	91.80		43.15		52.12		66.36	90.00	63.81
Baseline (Fake)	-	-	-	85.57	78.94	33.15	27.20	32.40	28.42	53.64	74.00	52.95
Baseline (Adp.)	-	-	-	84.81	78.38	39.56	37.94	37.71	36.70	81.09	91.40	83.43
Baseline (Aug.)	-	-	-	85.71	79.38	52.94	52.28	60.54	60.61	86.73	90.80	86.86
ADDA [53]	-	-	-	78.73	74.42	23.47	22.07	29.30	27.79	56.72	70.00	57.33
DADA [51]	-	-	-	83.39	76.56	26.42	23.68	20.59	18.93	55.09	74.20	53.33
DANN [14]	-	-	-	86.52	81.01	57.34	55.79	59.88	58.80	77.09	87.60	77.14
LTDA [25]	-	-	-	81.67	73.89	24.72	18.92	20.22	17.39	66.00	83.20	74.29
A-GEM [6]	-	-	-	86.91	80.33	33.38	29.37	33.09	31.88	66.91	91.60	67.81
Ours	-	v	v	86.25	80.96	54.71	53.51	61.47	61.10	89.27	93.00	87.61
	v	-	-	86.91	80.33	59.39	57.87	64.57	63.31	90.73	96.20	91.43
	v	v	-	85.59	81.25	60.65	58.69	66.04	64.88	85.27	95.20	90.48
	v	-	v	86.58	79.70	59.90	58.51	65.00	64.74	89.82	95.60	88.57
	v	v	v	85.94	81.93	60.17	59.11	66.24	65.05	87.64	95.80	89.33

adaptation methods as baselines—DANN [14], ADDA [53], DADA [51] and LTDA [25]. Note that unsupervised methods like ADDA and DADA do not use the labels of real images. Also, all methods except DANN deployed a two-stage training scheme, meaning the model is first pretrained on synthetic data for 150 epochs and then finetuned on real data for 50 epochs instead of directly training for 200 epochs. To incorporate real data during the training process, DANN, ADDA, and DADA aim to learn a domain-agnostic feature extractor via an adversarial scheme. LTDA deploys a meta-learning framework to adapt the model to the target distribution with a small subset from the target domain.

We report the results for different datasets in Tab. 4. It is worth noting that the domain adaptation methods besides DANN drastically underperformed compared to the baseline data augmentation. This is predictable because all the domain adaptation methods except DANN deployed a two-stage training scheme, where the synthetic data used for pre-training are not directly accessible at the finetuning stage. However, it also shows that the current domain adaptation methods, which mainly tackle the significant distribution shift (e.g., cartoon dogs \leftrightarrow dogs in natural images), do not address the fine-grained distribution shift between the synthetic images from a DGM and the real images. We believe that such kind of synthetic-to-real adaptation could be an interesting topic for future work.

However, using synthetic data for data augmentation is a promising approach. Note that the one-stage training scheme allows DANN to access synthetic and real data during the training process directly. Therefore, its performance is closer to the data augmentation baseline. We also compared our results to A-GEM, where the gradients from the real images are only used to constrain the gradients from the

synthetic ones but not used to update the model. We argue that in the case of synthetic data as data augmentation, it is important to also incorporate the real gradients during the training, particularly when the available real data are more than a few shots. It can be observed in Tab. 4 that applying our proposed PG and RG largely improved the classifier performance in most cases, especially in the low data regimes. This showcases the effectiveness of our method in regularizing the distances between data points and eliminating the negative effect of the synthetic data.

6. Conclusion

In this work, we examined to which degree synthetic data from popular DGMs can replace real data in downstream recognition tasks. While a severe performance drop is commonly observed in the context and generally attributed to *Domain Gap*, we further classified the potential factors into three categories. Moreover, we presented a series of observations indicating that the performance degradation can be mainly attributed to the *Content Gap*, where the synthetic data from DGMs can only form a simplified dataset in which rare samples are not represented. Motivated by this hypothesis, we proposed a novel method combining two strategies—Pretrained Guidance and Real Guidance—to regularize the downstream models to keep the span of features despite training on a simplified dataset, even when the regularization comes from a different network architecture. Extensive results on multiple datasets show that our method not only outperformed other synthetic-to-real generalization methods in zero-shot scenarios but also largely improved the effectiveness of the synthetic data when serving as data augmentation in low-shot settings.

References

- [1] David Bau, Jun-Yan Zhu, Hendrik Strobelt, Bolei Zhou, Joshua B. Tenenbaum, William T. Freeman, and Antonio Torralba. Gan dissection: Visualizing and understanding generative adversarial networks. *ArXiv*, abs/1811.10597, 2018. [2](#)
- [2] David Bau, Jun-Yan Zhu, Jonas Wulff, William S. Peebles, Hendrik Strobelt, Bolei Zhou, and Antonio Torralba. Seeing what a gan cannot generate. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 4501–4510, 2019. [2](#), [3](#)
- [3] Christopher Beckham, Issam Hadj Laradji, Pau Rodríguez López, David Vázquez, Derek Nowrouzezahrai, and Christopher Joseph Pal. Challenges in leveraging gans for few-shot data augmentation. In *CoLLAs*, 2022. [2](#)
- [4] Victor Besnier, Himalaya Jain, Andrei Bursuc, Matthieu Cord, and Patrick Pérez. This dataset does not exist: Training models from generated images. *ICASSP 2020 - 2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pages 1–5, 2019. [2](#)
- [5] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *ArXiv*, abs/1809.11096, 2018. [2](#)
- [6] Arslan Chaudhry, Marc’ Aurelio Ranzato, Marcus Rohrbach, and Mohamed Elhoseiny. Efficient lifelong learning with a-gem. *ArXiv*, abs/1812.00420, 2018. [5](#), [7](#), [8](#), [17](#)
- [7] Wuyang Chen, Zhiding Yu, SD Mello, Sifei Liu, Jose M Alvarez, Zhangyang Wang, and Anima Anandkumar. Contrastive syn-to-real generalization. 2021. [2](#), [5](#), [6](#), [7](#)
- [8] Wuyang Chen, Zhiding Yu, Zhangyang Wang, and Anima Anandkumar. Automated synthetic-to-real generalization. In *Proceedings of Machine Learning and Systems 2020*, pages 8272–8282, 2020. [2](#), [5](#), [6](#), [7](#)
- [9] Edo Collins, Raja Bala, Bob Price, and Sabine Süsstrunk. Editing in style: Uncovering the local semantics of gans. *2020 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 5770–5779, 2020. [1](#)
- [10] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *2009 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 248–255. Ieee, 2009. [2](#), [15](#)
- [11] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat gans on image synthesis. *Advances in Neural Information Processing Systems (NeurIPS)*, 34:8780–8794, 2021. [1](#), [2](#), [15](#)
- [12] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, et al. An image is worth 16x16 words: Transformers for image recognition at scale. *ArXiv*, 2020. [3](#)
- [13] Alexey Dosovitskiy, Philipp Fischer, Eddy Ilg, Philip Häusser, Caner Hazirbas, Vladimir Golkov, Patrick van der Smagt, Daniel Cremers, and Thomas Brox. FlowNet: Learning optical flow with convolutional networks. *2015 IEEE International Conference on Computer Vision (ICCV)*, pages 2758–2766, 2015. [2](#)
- [14] Yaroslav Ganin and Victor Lempitsky. Unsupervised domain adaptation by backpropagation, 2015. [2](#), [8](#), [17](#)
- [15] Giorgio Giannone, Didrik Nielsen, and Ole Winther. Few-shot diffusion models. *ArXiv*, abs/2205.15463, 2022. [2](#)
- [16] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative adversarial networks. *Communications of the ACM*, 63(11):139–144, 2020. [1](#)
- [17] Jiyeon Han, Hwanil Choi, Yunjeon Choi, Jae Hyun Kim, Jung-Woo Ha, and Jaesik Choi. Rarity score : A new metric to evaluate the uncommonness of synthesized images. *ArXiv*, abs/2206.08549, 2022. [4](#)
- [18] Kaiming He, X. Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. *2016 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 770–778, 2015. [3](#), [15](#)
- [19] Ruifei He, Shuyang Sun, Xin Yu, Chuhui Xue, Wenqing Zhang, Philip H. S. Torr, Song Bai, and Xiaojuan Qi. Is synthetic data from generative models ready for image recognition? *ArXiv*, abs/2210.07574, 2022. [2](#)
- [20] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and S. Hochreiter. Gans trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2017. [3](#)
- [21] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. *Advances in Neural Information Processing Systems (NeurIPS)*, 33:6840–6851, 2020. [1](#), [2](#)
- [22] Ahmed Imtiaz Humayun, Randall Balestrieri, and Richard Baraniuk. Polarity sampling: Quality and diversity control of pre-trained generative networks via singular values. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10631–10640, 2022. [4](#)
- [23] Ali Jahanian, Lucy Chai, and Phillip Isola. On the” steerability” of generative adversarial networks. *ArXiv*, 2019. [1](#)
- [24] Ali Jahanian, Xavier Puig, Yonglong Tian, and Phillip Isola. Generative models as a data source for multiview representation learning. *ArXiv*, abs/2106.05258, 2021. [1](#), [2](#)
- [25] Muhammad Abdullah Jamal, Matthew Brown, Ming-Hsuan Yang, Liqiang Wang, and Boqing Gong. Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective. *CoRR*, abs/2003.10780, 2020. [8](#), [17](#)
- [26] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. *2019 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 4396–4405, 2018. [2](#)
- [27] 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 (CVPR)*, pages 4401–4410, 2019. [12](#)
- [28] Diederik P. Kingma and Max Welling. Auto-encoding variational bayes. *CoRR*, abs/1312.6114, 2013. [1](#)
- [29] Alex Krizhevsky. Learning multiple layers of features from tiny images. Technical report, 2009. [2](#)

- [30] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. *CoRR*, abs/1904.06991, 2019. **13, 15**
- [31] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Deeper, broader and artier domain generalization. *2017 IEEE International Conference on Computer Vision (ICCV)*, pages 5543–5551, 2017. **2**
- [32] Da Li, Yongxin Yang, Yi-Zhe Song, and Timothy M. Hospedales. Learning to generalize: Meta-learning for domain generalization. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2017. **2**
- [33] Taehong Moon, Moonseok Choi, Gayoung Lee, Jung-Woo Ha, and Juho Lee. Fine-tuning diffusion models with limited data. In *Advances in Neural Information Processing Systems (NeurIPS) 2022 Workshop on Score-Based Methods*, 2022. **2**
- [34] Krikamol Muandet, David Balduzzi, and Bernhard Schölkopf. Domain generalization via invariant feature representation. In *International Conference on Machine Learning (ICML)*, 2013. **2**
- [35] Muhammad Ferjad Naeem, Seong Joon Oh, Youngjung Uh, Yunjey Choi, and Jaejun Yoo. Reliable fidelity and diversity metrics for generative models. 2020. **13, 15**
- [36] M-E. Nilsback and A. Zisserman. Automated flower classification over a large number of classes. In *Proceedings of the Indian Conference on Computer Vision, Graphics and Image Processing*, Dec 2008. **6**
- [37] Shuanlong Niu, Bin Li, Xinggang Wang, and Hui Lin. Defect image sample generation with gan for improving defect recognition. *IEEE Transactions on Automation Science and Engineering*, 17(3):1611–1622, 2020. **1**
- [38] Utkarsh Ojha, Yijun Li, Jingwan Lu, Alexei A Efros, Yong Jae Lee, Eli Shechtman, and Richard Zhang. Few-shot image generation via cross-domain correspondence. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10743–10752, 2021. **5**
- [39] Or Patashnik, Zongze Wu, Eli Shechtman, Daniel Cohen-Or, and Dani Lischinski. Styleclip: Text-driven manipulation of stylegan imagery. In *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2085–2094, October 2021. **1**
- [40] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *ArXiv*, abs/1710.06924, 2017. **2**
- [41] M. Pharr, W. Jakob, and G. Humphreys. *Physically Based Rendering: From Theory to Implementation*. Elsevier Science, 2016. **3**
- [42] Stephan R. Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In Bastian Leibe, Jiri Matas, Nicu Sebe, and Max Welling, editors, *European Conference on Computer Vision (ECCV)*, volume 9906 of *LNCS*, pages 102–118. Springer International Publishing, 2016. **2**
- [43] Esther Robb, Wen-Sheng Chu, Abhishek Kumar, and Jia-Bin Huang. Few-shot adaptation of generative adversarial networks. *ArXiv*, abs/2010.11943, 2020. **2**
- [44] Daniel Roich, Ron Mokady, Amit H Bermano, and Daniel Cohen-Or. Pivotal tuning for latent-based editing of real images. *ACM Trans. Graph.*, 2021. **1**
- [45] Swami Sankaranarayanan, Yogesh Balaji, Arpit Jain, Ser Nam Lim, and Rama Chellappa. Learning from synthetic data: Addressing domain shift for semantic segmentation. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 3752–3761, 2018. **1, 2**
- [46] Mert Bulent Sariyildiz, Alahari Karteek, Diane Larlus, and Yannis Kalantidis. Fake it till you make it: Learning transferable representations from synthetic imagenet clones. 2022. **2**
- [47] Axel Sauer, Kashyap Chitta, Jens Müller, and Andreas Geiger. Projected gans converge faster. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021. **6, 13**
- [48] Axel Sauer, Katja Schwarz, and Andreas Geiger. Stylegan-xl: Scaling stylegan to large diverse datasets. volume abs/2201.00273, 2022. **2, 3, 13, 15**
- [49] Abhishek Sinha, Jiaming Song, Chenlin Meng, and Stefano Ermon. D2c: Diffusion-decoding models for few-shot conditional generation. In M. Ranzato, A. Beygelzimer, Y. Dauphin, P.S. Liang, and J. Wortman Vaughan, editors, *Advances in Neural Information Processing Systems (NeurIPS)*, volume 34, pages 12533–12548. Curran Associates, Inc., 2021. **2**
- [50] Mingxing Tan and Quoc Le. Efficientnet: Rethinking model scaling for convolutional neural networks. In *International Conference on Machine Learning (ICML)*, pages 6105–6114. PMLR, 2019. **3**
- [51] Hui Tang and Kui Jia. Discriminative adversarial domain adaptation. In *Association for the Advancement of Artificial Intelligence (AAAI)*, 2020. **2, 8, 17**
- [52] Joshua Tobin, Rachel Fong, Alex Ray, Jonas Schneider, Wojciech Zaremba, and P. Abbeel. Domain randomization for transferring deep neural networks from simulation to the real world. *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 23–30, 2017. **2**
- [53] Eric Tzeng, Judy Hoffman, Kate Saenko, and Trevor Darrell. Adversarial discriminative domain adaptation. *2017 IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 2962–2971, 2017. **2, 8, 17**
- [54] Arash Vahdat, Karsten Kreis, and Jan Kautz. Score-based generative modeling in latent space. In *Neural Information Processing Systems (NeurIPS)*, 2021. **2, 3, 13**
- [55] C. Wah, S. Branson, P. Welinder, P. Perona, and S. Belongie. The caltech-ucsd birds-200-2011 dataset. Technical Report CNS-TR-2011-001, California Institute of Technology, 2011. **6**
- [56] Ruyi Wang, Sabrina Hoppe, Eduardo Monari, and Marco Huber. Defect transfer gan: Diverse defect synthesis for data augmentation. In *33rd British Machine Vision Conference 2022, BMVC 2022, London, UK, November 21-24, 2022*. BMVA Press, 2022. **1, 6, 12, 13**
- [57] Daniel Watson, William Chan, Jonathan Ho, and Mohammad Norouzi. Learning fast samplers for diffusion models by differentiating through sample quality. *ArXiv*, abs/2202.05830, 2022. **4**

- [58] Zongze Wu, Dani Lischinski, and Eli Shechtman. Stylespace analysis: Disentangled controls for stylegan image generation. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 12858–12867, 2020. [1](#)
- [59] Zongze Wu, Yotam Nitzan, Eli Shechtman, and Dani Lischinski. Stylealign: Analysis and applications of aligned stylegan models. *ArXiv*, 2021. [2](#)
- [60] Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with denoising diffusion gans. *ArXiv*, 2021. [2](#)
- [61] Xiangyu Yue, Yang Zhang, Sicheng Zhao, Alberto L. Sangiovanni-Vincentelli, Kurt Keutzer, and Boqing Gong. Domain randomization and pyramid consistency: Simulation-to-real generalization without accessing target domain data. *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 2100–2110, 2019. [2](#)
- [62] Sergey Zagoruyko and Nikos Komodakis. Wide residual networks. *ArXiv*, abs/1605.07146, 2016. [3](#)
- [63] Gongjie Zhang, Kaiwen Cui, Tzu-Yi Hung, and Shijian Lu. Defect-gan: High-fidelity defect synthesis for automated defect inspection. In *Proceedings of the IEEE/CVF Winter Conference on Applications of Computer Vision (WACV)*, pages 2524–2534, January 2021. [1](#)
- [64] Yuxuan Zhang, Huan Ling, Jun Gao, K. Yin, Jean-Francois Lafleche, Adela Barriuso, Antonio Torralba, and Sanja Fidler. Datasetgan: Efficient labeled data factory with minimal human effort. *2021 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 10140–10150, 2021. [1](#), [2](#)
- [65] Yunqing Zhao, Henghui Ding, Houjing Huang, and Ngai-Man Cheung. A closer look at few-shot image generation. *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pages 9130–9140, 2022. [2](#)