

Will Large-scale Generative Models Corrupt Future Datasets?

Ryuichiro Hataya
 RIKEN ADSP & RIKEN AIP
 ryuichiro.hataya@riken.jp

Han Bao
 Kyoto University
 bao@i.kyoto-u.ac.jp

Hiromi Arai
 RIKEN AIP
 hiromi.arai@riken.jp

Abstract

Recently proposed large-scale text-to-image generative models such as DALL·E 2 [47], Midjourney [42], and StableDiffusion [51] can generate high-quality and realistic images from users’ prompts. Not limited to the research community, ordinary Internet users enjoy these generative models, and consequently, a tremendous amount of generated images have been shared on the Internet. Meanwhile, today’s success of deep learning in the computer vision field owes a lot to images collected from the Internet. These trends lead us to a research question: “**will such generated images impact the quality of future datasets and the performance of computer vision models positively or negatively?**” This paper empirically answers this question by simulating contamination. Namely, we generate ImageNet-scale and COCO-scale datasets using a state-of-the-art generative model and evaluate models trained with “contaminated” datasets on various tasks, including image classification and image generation. Throughout experiments, we conclude that generated images negatively affect downstream performance, while the significance depends on tasks and the amount of generated images. The generated datasets and the codes for experiments will be publicly released for future research. Generated datasets and source codes are available from <https://github.com/moskomule/dataset-contamination>.

1. Introduction

Deep generative models for image generation have advanced progressively since the original GANs [17] and VAEs [32]. Recently, denoising diffusion models [27, 57, 58] have beaten GANs in image quality [11, 28] and become one of the de-facto standard generative models. Among them, some large models trained with billion-scale captioned images collected from the Internet achieved high-fidelity image generation conditioned by users’ prompts [47, 53, 51, 42, 41, 65, 64, 16, 3]. Particularly, DALL·E 2 [47], Midjourney [42], and StableDiffusion [51] have web and smartphone applications, and many

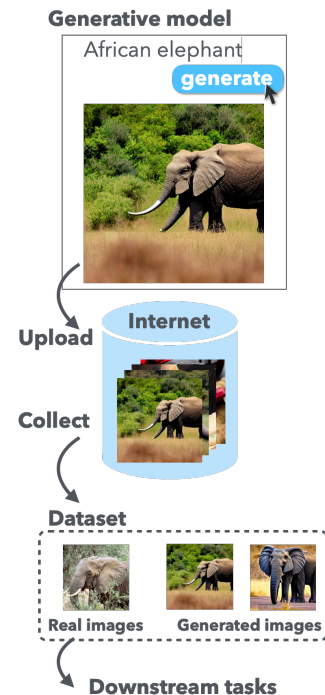


Figure 1: Schematic view of the problem. Some large-scale generative models are public, and many users are playing with them to share generated images on the Internet (top). Dataset collection heavily relies on images on the Internet, which may be contaminated by generated images (bottom). This paper discusses the effect of such dataset corruption.

Internet users enjoy image generation, and consequently, a tremendous amount of generated images have been uploaded to the Internet.¹

At the same time, highly realistic generated images may have potentially significant impacts on society. For example, face images generated by StyleGANs [31] were reportedly used to create fake profiles of SNSs or dating apps to deceive other users [20, 26]. Furthermore, recent text-to-

¹The DALL·E 2 model alone generated two million images per day in September 2022, according to <https://openai.com/blog/dall-e-now-available-without-waitlist/>.

image generative models can generate images that look real at first glance from users’ instruction and are able to support fake news [59]. They also amplify demographic stereotypes [4].

Another concern is that generated images might affect the quality of newly curated image datasets from the Internet in the future, similar to the fact that the outputs of machine translation models degenerate the quality of corpora [55, 48, 12]. Without a doubt, today’s success of deep learning and computer vision, including generative models themselves, largely owes to image datasets collected from the Internet, such as ImageNet [52, 9]. However, when generated images are shared on the Internet, they may contaminate the sources of image datasets (see also Fig. 1).² Based on these backgrounds, our research question in this paper raises: *what will happen if datasets were contaminated by generated images?*

We aim to answer this question through experiments: simulating such contamination by large-scale datasets of generated images and measuring downstream performance trained with them. Specifically, we generate three million images from ImageNet categories and COCO captions using StableDiffusion and emulate the contamination by replacing real images in datasets with generated ones. Then, we measure the performance of models trained with such contaminated datasets in various tasks, namely, image classification, image captioning, and image generation. Throughout experiments, we find that generated images have *negative* effects on downstream performance. We hypothesize that such negative effects are caused by the fact that the generative models capture fewer modes than the actual data, although the existing synthetic experiments have shown high coverage [62].

In summary, our contributions are as follows:

- To simulate the effects of possible contamination, we create large-scale datasets consisting of generated images corresponding to ImageNet and COCO caption (Section 3).
- We conduct experiments over four distinct tasks on the generated datasets and discover negative effects of contamination, which can be partially attributed to fewer modes of generated images than real data (Sections 4 and 5).
- Based on the empirical results, we recommend to researchers how to publish generative models and how to collect datasets (Section 7).

²Although the “official” web applications implant watermarks to generated images, which thus can be filtered, we found that some software has options to disable such functions, e.g., <https://github.com/AUTOMATIC1111/stable-diffusion-webui>.

2. Background and Related Work

A deep generative model aims to approximate the underlying data distribution by neural networks, and the sampled data are expected to be similar to real ones. Since the emergence of GANs [17, 46], the research of deep generative models has advanced progressively. In particular, denoising diffusion models, equivalently, score-based generative models, have achieved high-quality image generation with diversity [27, 57, 58], capturing modes of data distributions faithfully [62]. Text-to-image generative models based on diffusion models can generate high-quality images from users’ text instructions with high fidelity, even for unseen novel combinations of concepts, such as “a photo of an astronaut riding a horse” [47, 53, 51, 42, 41, 65, 64, 16, 3]. Notably, some models have publicly accessible applications [51, 42, 47, 65], and many Internet users are generating images and posting them on the Internet with related texts. Such generated images are sometimes difficult to be distinguished from real ones, and thus some of them are potent to contaminate future datasets collected from the web.

The NLP community has experienced similar problems in the last decade; thanks to the development of NLP technologies, many contents on the Internet have become machine-generated, e.g., by machine translation and optical character recognition systems, but such generated texts have degenerated quality of corpora [55, 12]. As a result, filtering such low-quality samples is essential to maintain downstream performance [48]. Although the NLP community has investigated such issues by generated data, the effects of generated images by text-to-image models on various downstream performances in computer vision have rarely been studied.

The dataset contamination issue in general has been studied from various aspects; including dataset poisoning [8], adversarial training [18], label noise [19], outlier robustness [29], and distribution shifts [56]. The existing studies usually posit an attacker/contamination model that is plausible yet mathematically convenient, such as Huber’s contamination model [29]. By contrast, we are rather interested in realistic contamination of the web images induced by generative models and its potential effects.

Methodologically, our work is related to [49]; they also use accuracy on ImageNet classification, but their purpose is to measure how much a generative model captures the ImageNet data distribution. Differently, our interest is not in evaluation of generative models, but the effect of generated images on various downstream tasks, and the image classification task is used along with others.

3. Dataset Corruption by Generated Images

This paper aims to answer our research question “*will the contamination of generated images perform positively*”

or negatively?” To empirically answer this question, we simulate realistic dataset contamination by generated images and evaluate the quality of contaminated datasets by training commonly-used models with such datasets in several tasks. In this section, we describe the dataset creation.

3.1. Dataset Creation

To simulate image generation by users, we create datasets using a StableDiffusion model [51], a state-of-the-art text-to-image generative model, pre-trained with LAION-2B [54], which include not all but at least some ImageNet and COCO images. These datasets are generated from category names of the ImageNet ILSVRC-2012 classification dataset and captions of the COCO caption dataset, which are referred to as SD-ImageNet and SD-COCO in the remaining text. For generation of both datasets, we disabled the watermarking functionality to trace outputs as generated images and the safety checker to reduce explicit outputs, such as nudes.

SD-ImageNet: The ImageNet ILSVRC-2012 classification dataset [52] is a subset of ImageNet [9] and a dataset for the image classification task. Its training set contains 1.2 million photo images over 1,000 categories selected from synsets of WordNet, *e.g.*, “African elephant”. Using these category names, we prepared prompts like “A photo of African elephant” for each category and generated 1,400 photography-like images for each class.

Figure 2 (left) shows examples from SD-ImageNet. Images are natural at first glance, but contain some flaws. For example, the elephant at the top left has two noses. We will revisit the creation of prompts in Section 5.2.

SD-COCO: The COCO caption dataset [7] is a dataset for the image captioning task. Based on the dataset split in [30], this dataset has 113,000 images with five captions for each image, such as “A small child wearing headphones plays on the computer”. These captions were used as prompts to generate 565,000 images.

Figure 2 (right) presents some examples from SD-COCO with their captions. Similar to examples of SD-ImageNet, the images are apparently faithful to the captions used as prompts, but staring at them reveals unnatural or unfaithful details. For example, the bottom right example fails to produce a “blue and white plate.”

In the remaining text, we call the ILSVRC-2012 dataset as ImageNet and the COCO caption dataset as COCO for simplicity.

3.2. Simulation of Corruption

To simulate possible corruption, we randomly substitute generated images for 20, 40, and 80 % of real images of the original datasets with generated ones without replacement.

We refer to these mixed datasets as IN/SD- $n\%$, where n indicates the ratio of generated data. Similarly to IN/SD- $n\%$, we also created mixtures of COCO and SD-COCO, which are referred to as CO/SD- $n\%$.

In the next section (Section 4), we empirically investigate the effect of corruption using the downstream performance of models trained with these contaminated datasets. However, mixtures of ImageNet and SD-ImageNet, *e.g.*, IN/SD-20%, alone still entangle the effect of artifacts of generated images and the domain shift between generated images and real images. To disentangle the effect of generated images from that of domain shift, we additionally use datasets consisting of real images similar to ImageNet and COCO and compare the downstream performance of models trained with them with that of IN/SDs or IN/COs. As a counterpart of ImageNet, we adopt a subset of the Web-Vision dataset [36], which was collected by querying ImageNet category names to Google and Flickr; then, we mix it with ImageNet. Because this subset is imbalanced, and some of its categories contain fewer images than needed, we then use sampling by replacement to create balanced mixtures. Similar to IN/SD- $n\%$, we refer to these mixed datasets as IN/WV- $n\%$. Correspondingly, as a counterpart of COCO, we use Flickr-30k [45], which contains 32,000 images collected from Flickr with five captions per image. Because its size is much less than COCO, we only prepare CO/FL-40% as a mixture of COCO and Flickr-30k.

4. Experimental Results

In this section, we evaluate the effect of contamination using the datasets created in Section 3 on several downstream tasks.

Shared Experimental Settings

We used neural network models implemented with PyTorch v1.12 [44] and its accompanying torchvision with CUDA 11.3. Experiments including dataset creation described in Section 3 were conducted on NVIDIA V-100 GPUs and NVIDIA A-100 GPUs. Further description of settings and configurations can be found in Appendix A in the Supplementary Material.

4.1. Image Classification

This task classifies images into 1,000 categories of ImageNet. We used ResNet-50 [22], SwinTransformer-S (Swin-S) [38], and ConvNeXt-T [39] in torchvision, training them according to the standardized training protocols³ with ImageNet, SD-ImageNet, WebVision, and their mixtures. ResNet is a convolutional neural network with


³<https://github.com/pytorch/vision/tree/v0.13.1/references/classification>. Exceptionally, ConvNeXt was trained for 300 epochs due to computational resources.

SD-ImageNet

African elephant (n02504458)



Grand piano (n03452741)



SD-COCO



A small child wearing headphones plays on the computer



A commercial stainless kitchen with a pot of food cooking



Red wine is poured into a glass



A dessert with pecans, cherries and nectarines on a blue and white plate.

Figure 2: Randomly selected examples from generated datasets, namely, SD-ImageNet (left) and SD-COCO (right). Images are at a single glance high quality and fidelity to prompts, *i.e.*, category names and captions, while details are unnatural, *e.g.*, a two-nose elephant.

residual connections, SwinTransformer is a variant of Vision Transformer, and ConvNeXt is a CNN inspired by Vision Transformers, which represent modern vision models.

Table 1 shows accuracy on the ImageNet validation set. As can be seen, the performance decreases as the ratio of SD-ImageNet in training data increases. When the ratio of generated images is at most 40%, the performance drops are marginal and may be endurable in most practical scenarios. However, when the ratio is 80%, the performance degeneration is not negligible. Compared to SD-ImageNet, WebVision images have less influence on performance. In the extreme cases, when no ImageNet data are included in training data, *i.e.*, SD-ImageNet and WebVision, this difference in performance is significant, which suggests that the performance drop may not be solely due to the domain gap.

Additionally, Fig. 3 presents confusion matrices of ResNet-50 trained with ImageNet and SD-ImageNet. For clarity, categories are subsampled and rearranged according to 12 superclasses, adopting `big_12` classes from [15]. As can be seen, the mispredictions by ResNet-50 trained with ImageNet mostly fall in the same fine-grained categories, represented by diagonal blocks. Contrarily, we observe that the model trained with SD-ImageNet uniformly misclassifies certain classes, partially because the category names of such classes are ambiguous, and thus, the generated images for such classes are semantically diverse. `Titi` (monkey) is such an example, where it is intended to mean a New World monkey in ImageNet, but it is also a name of people, plants, and places, and thus, generated images are also semantically diverse (see Fig. B.4 in the Supplementary Material).

	ResNet-50	Swin-S	ConvNeXt-T
ImageNet	75.7	83.1	80.8
IN/SD-20%	74.5	82.1	79.7
IN/SD-40%	72.6	81.0	78.3
IN/SD-80%	65.3	74.3	70.8
SD-ImageNet	15.7	19.3	19.6
IN/WV-20%	75.1	82.5	80.0
IN/WV-40%	73.9	81.8	78.8
IN/WV-80%	68.3	NaN	73.9
WebVision	61.3	70.9	66.2

Table 1: Validation accuracy of the image classification task on the ImageNet validation set. We could not stably train Swin-S on IN/WV-80, which resulted in a loss explosion. The performance drop is marginal when the ImageNet images dominate the dataset.

4.2. Image Captioning

Image captioning is a task to generate appropriate captions for given images. We used a pre-trained BLIP model [35], a state-of-the-art vision-language model, and fine-tuned its captioner and filter modules on COCO, SD-COCO, Flickr-30k, and their mixtures for five epochs, following [34].

Table 2 reports the performance in various statistics with the COCO test set when the captions are generated with beam search with a beam size of three. Aligned with the results of image classification, a performance drop by generated images can also be observed. Especially, CO/SD-20% yields comparable or inferior performance to CO/FL-40%, even in metrics for image captioning like SPICE [1]

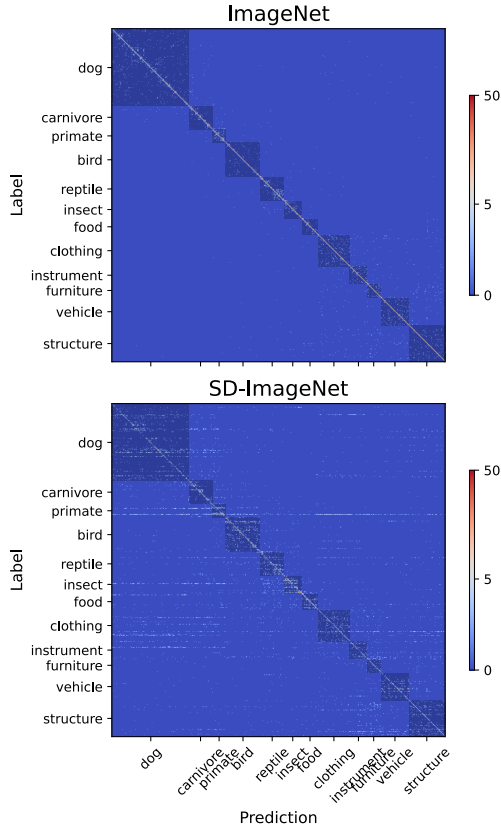


Figure 3: Confusion matrices of ResNet-50 predictions with a subset of ImageNet validation data. Models were trained with ImageNet and SD-ImageNet. Class indices are rearranged. X and Y axes depict superclasses by gathering categories. Colors correspond to the number of data at each pixel in log scale, and block-diagonal components are highlighted. Notice that the model trained with SD-ImageNet misclassifies certain classes uniformly, as illustrated by white horizontal lines.

and CIDEr [61], indicating that generated images cause degeneration of dataset quality. Moreover, a comparison between the results of SD-COCO and Flickr-30k suggest that such performance drops cannot be fully attributed to domain shift.

4.3. Image Generation

Finally, we verify if generated images are useful as training data for the image generation task. We evaluated an improved denoising diffusion probabilistic model (IDDPM) [40] with datasets resized to 64×64 , which we refer to as, for example, ImageNet-64 and IN/SD-20%-64. We trained the model for 1.8×10^6 iterations using the L_{hybrid} objective [40] with a batch size of 512 and generated 5.0×10^4 images with 250 sampling steps.

Table 3 reports the quality of unconditionally generated




Figure 4: Randomly selected generated images by class-unconditional IDDPM [40] using 250 sampling steps.

images in Fréchet Inception Distance (FID) [25], improved precision and recall metrics of 5-nearest neighbors [33] between Inception features of generated images and all validation data from ImageNet and WebVision. Randomly sampled generated images are presented in Fig. 4. The improved precision and recall metrics are computed by estimating the volumes of real and generated images in the embedded space with nearest neighbors [33], from which we can deduce how much two image distributions overlap with each other. From Table 3, we see the trend that the precision and recall increases and decreases, respectively, as the ratio of generated images in training data increases. To put it differently, the heavier contamination results in generated images that are more likely to be in the support of test images, while the test support coverage is worsened. This indicates that generated images may concentrate on a smaller subset of the test support.

5. Analysis

5.1. Possible Cause of Degeneration

Seeing the performance degeneration caused by contamination, we hypothesize that generated images have fewer modes than real ones. To verify this idea, we measure the precision and recall of ImageNet training data and SD-ImageNet compared to ImageNet validation data, presented in Table 4. As can be seen, the recall of SD-ImageNet images is significantly smaller than that of ImageNet images, indicating that SD-ImageNet images cannot cover the modes of ImageNet images, *i.e.*, SD-ImageNet is less diverse than ImageNet. The lack of diversity can also be observed visually, comparing images from ImageNet and SD-ImageNet in Fig. 5 top and middle. This observation can explain the performance degeneration in the main experiments, where contaminated datasets are concentrated in some modes, and thus models trained with them cannot generalize better.

	BLEU-1	BLEU-2	BLEU-3	BLEU-4 [43]	SPICE [1]	METEOR [10]	ROUGE-L [37]	CIDEr [61]
COCO	0.791	0.641	0.508	0.400	0.240	0.310	0.602	1.335
CO/SD-20%	0.787	0.634	0.500	0.391	0.235	0.306	0.596	1.320
CO/SD-40%	0.786	0.632	0.499	0.390	0.236	0.307	0.596	1.319
CO/SD-80%	0.780	0.623	0.486	0.377	0.233	0.300	0.588	1.279
SD-COCO	0.711	0.534	0.394	0.287	0.191	0.252	0.514	1.000
CO/FL-40%	0.787	0.634	0.501	0.393	0.238	0.308	0.598	1.326
Flickr 30k	0.754	0.587	0.439	0.321	0.215	0.275	0.554	1.092
w/o fine-tuning	0.473	0.392	0.308	0.237	0.158	0.212	0.488	0.838

Table 2: Test metrics in image captioning of the BLIP model on the COCO test split. Higher values are better.

	FID ↓	Precision@5 ↑	Recall@5 ↑
ImageNet-64	14.9 / 15.6	0.665 / 0.679	0.644 / 0.653
IN/SD-20%-64	12.6 / 12.7	0.699 / 0.708	0.621 / 0.634
IN/SD-40%-64	11.0 / 10.8	0.730 / 0.739	0.585 / 0.608
IN/SD-80%-64	12.5 / 11.3	0.795 / 0.802	0.490 / 0.512
SD-ImageNet-64	16.9 / 15.4	0.831 / 0.835	0.364 / 0.379

Table 3: Image quality comparison on unconditional ImageNet-64 validation data and WebVision-64 validation data using Inception-V3 features shown in left and right of each cell, respectively.

Dataset	Precision	Recall
ImageNet	0.757	0.791
Original SD-ImageNet	0.836	0.344
Complex SD-ImageNet	0.777	0.450

Table 4: Precision and recall of the ImageNet training data and generated datasets compared with the ImageNet validation data using Inception-V3 features.

5.2. Effects of Prompts

As described in Section 3.2, the images of SD-ImageNet are generated from simple prompts like “a photo of African elephant,” which are unique to each class. To verify the effects of prompt diversity, we create another SD-ImageNet, referred to as complex SD-ImageNet, with more complex prompts mimicking humans’, such as “a monochrome image of African elephant taken with iPhone” and “HDR picture of grand piano outside.” These prompts are programmatically generated from 200 to 1300 variations per class, and the details are explained in Appendix A.5. Figure 5 bottom illustrates samples from complex SD-ImageNet, Table 4 bottom line shows quality of images, Table 5 presents validation accuracy of ResNet-50 trained with the original and complex SD-ImageNets. Although the complex SD-ImageNet’s generation prompts are much more diverse than the original ones, the diversity of generated images is far less than real ones. Consequently, the performance gain

is marginal, indicating that our observation can be applied to predicting the problems caused by images conditionally generated from prompts composed by humans.

	ResNet-50	Swin-S	ConvNeXt-T
Original SD-IN	15.7	19.3	19.6
Complex SD-IN	18.6	21.9	23.1

Table 5: Validation accuracy in classification of ResNet-50 trained with the original SD-ImageNet and complex SD-ImageNet generated from more complex prompts.

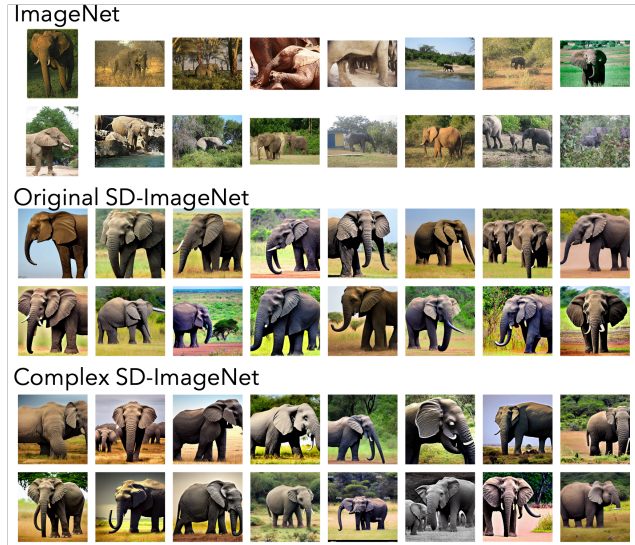


Figure 5: Randomly sampled images from the “African elephant” category of ImageNet, the original SD-ImageNet, and the complex SD-ImageNet. See Figs. B.1 to B.3 in the Supplementary Material for more examples.

5.3. Effects on Robustness

In the main experiments of the classification task, we saw the performance only by accuracy on validation data.

Source	ImageNet Val	ImageNet-A	ImageNet-R
ImageNet	75.7	1.76	36.7
IN/SD-20%	74.5	2.12	35.7
IN/SD-40%	72.6	1.67	35.0
IN/SD-80%	65.3	1.61	30.0
IN/WV-20%	75.1	3.55	40.7
IN/WV-40%	73.9	3.77	41.5
IN/WV-80%	68.3	5.00	40.0

Table 6: Robustness metrics in classification accuracy of ResNet-50 on the ImageNet validation set, ImageNet-A [24], and ImageNet-R [23]. Different from IN/WVs, IN/SDs generally affect robustness to out-of-distribution data.

To further investigate the effect of generated images on learned representation, we measured accuracy on other validation data; namely, ImageNet-A [24] and ImageNet-R [23]. These datasets share the same categories with ImageNet, but are curated independently to measure robustness to out-of-distribution data.

Table 6 summarizes the results, which generally indicate that generated images degenerate the robustness, except for IN/SD-20% on ImageNet-A. Contrarily, WebVision images may consistently enhance the robustness to out-of-distribution data on ImageNet-A and ImageNet-R, meaning that they are from a different distribution than ImageNet but diverse enough (compare, *e.g.*, IN/WV-20% to IN/WV-40%). These results further support the hypothesis that generated images have fewer modes than the real data, and thus, cause the downstream performance drop on test data and out-of-distribution data.

5.4. Comparison with Subsampled and added Data

In the main experiments, we compared the performances between networks trained with the contaminated datasets and with the *full-size* clean datasets. However, one may argue that the performance degradation results from the different amounts of real data in the training set. In this section, we compare the performance of subsampled and added real datasets with contaminated datasets to disentangle the effect of the amount of clean data.

Table 7 shows the validation accuracy of ResNet-50 trained with a 5% subset of ImageNet and IN/SD-95%, which fills 95% of missing data by generated data. Although IN/SD-95% yields 7.4% performance improvement over subsampled ImageNet, this is inferior to the gain by IN/WV-95%. We additionally measured the validation accuracy of ResNet-50 trained with full ImageNet (100%) added with SD-ImageNet with different sizes. Table 8 shows the results of these experiments, indicating that SD-ImageNet images do not contribute to performance improvement even though the total dataset size increases.

ImageNet (5%)	IN/SD-95%	IN/WV-95%
44.7	52.1	61.4

Table 7: Validation accuracy of ResNet-50 trained with a 5% subsampled ImageNet, IN/SD-95%, and IN/WV-95%.

	Accuracy
ImageNet (100%)	75.7
ImageNet (100%) + SD-ImageNet (20%)	75.7
ImageNet (100%) + SD-ImageNet (40%)	74.9
ImageNet (100%) + SD-ImageNet (80%)	75.0

Table 8: Validation accuracy of ResNet-50 when SD-ImageNet images are appended to the original ImageNet dataset.

	BLEU-4 \uparrow	SPICE \uparrow	CIDEr \uparrow
COCO (5%)	0.386	0.233	1.290
CO/SD-95%	0.362	0.226	1.220
COCO (20%)	0.385	0.235	1.305
CO/SD-80%	0.377	0.233	1.279

Table 9: Test metrics of BLIP trained with 5% and 20% COCO subsets and mixtures to complement the missing data. Higher values are better.

In Table 9, test metrics of BLIP trained with 5% and 20% subsets of COCO and their corresponding CO/SDs are presented. In this case, adding generated data affects negatively, even when only 5% of real data are available.

Additionally, Table 10 compares the image quality of generated images by IDDPM trained with the IN/SD-80% and a 20% subset of ImageNet. Aligned with the results in Section 4, the recall metric diminishes with contaminated data, supporting the hypothesis that the modes of generated images are fewer than real ones.

These results emphasize the negative effects of generated data. Additionally, the observations imply that using generated images for data augmentation needs careful consideration. Such an idea has been studied in image classification using conditional GANs [60, 2], particularly in medical imaging [63], but also known to hinder the final performance in large-scale settings [50]. Our results align with the latter that generated images are not always effective in data augmentation.

6. Discussion

6.1. Detection of Generated Images

To avoid the negative effects by generated images, one may want to detect generated images easily. For example, exploiting the differences between real and generated im-

	FID ↓	Precision@5 ↑	Recall@5 ↑
ImageNet-64 (20%)	16.5	0.639	0.646
IN/SD-80%-64	12.5	0.795	0.490

Table 10: Image quality comparison of generated images of IDDPM trained with the IN/SD-80% and a 20% ImageNet subset.

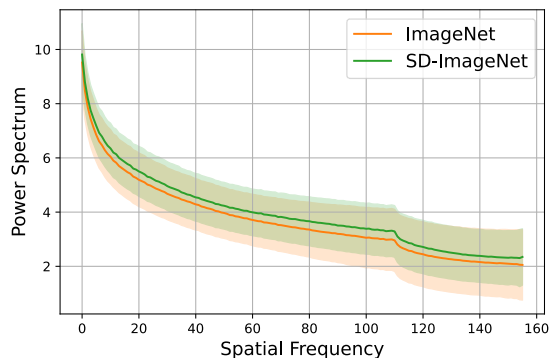


Figure 6: Average and standard deviation of spectra of 10^3 images from ImageNet and SD-ImageNet.

ages in high-frequency spectra is a simple and convincing approach [14]. However, this discrepancy may be caused when an upsampling operation (to decode the original images from the low-dimensional latent representations) includes zero pixel insertion; otherwise, detecting generated images only by frequency spectra is difficult [5]. Probably because StableDiffusion uses upsampling by nearest neighbor rather than zero pixel insertion, distinguishing real and generated images only from frequency information may be difficult. Figure 6 presents the power spectra of 1,000 images from ImageNet and SD-ImageNet, which are highly overlapped in all frequencies, which agrees with [5]. Additionally, we trained a linear classifier and a multi-layer perceptron with ImageNet-pre-trained ResNet features to detect generated images. When trained with 10^4 images from both datasets, the classifiers achieved around 85% test accuracy on 2,000 separated test images, which is still an unsatisfactory detection rate for a binary classification task and indicates the difficulty of detection of generated images.

6.2. Self-supervised Learning as Remedy

The experimental results so far indicate negative effects of generated images, and automatically filtering them out may be difficult (Section 6.1). To alleviate the effect, one may be able to use self-supervised learning that aims to acquire useful representations without using explicit supervision. To verify this idea, we adopt a self-supervised learning method called masked autoencoder (MAE) [21]. We pre-train a Vision Transformer model, specifically ViT-B [13],

	Accuracy
ImageNet	43.9
IN/SD-20%	44.0
IN/SD-40%	44.1
IN/SD-80%	42.5
SD-ImageNet	38.8
IN/WD-20%	43.3
IN/WD-40%	43.3
IN/WD-80%	40.9
WebVision	44.1

Table 11: Validation accuracy of linear probing of MAE [21] on the ImageNet validation set.

as MAE’s encoder for 200 epochs with a mask ratio of 0.75. Table 11 presents validation accuracy after 90 epochs of linear probing that trains only the last classifier layer from the extracted features. Even when the training dataset is fully generated images, the performance degeneration is limited, indicating that self-supervised learning may be a promising way to circumvent the negative effects of generated images.

6.3. Limitations

This paper has revealed the potential effects of generated images on datasets through various experiments. Nevertheless, the discussion has some limitations. Firstly, we could only use StableDiffusion trained with LAION-2B, because models and their pre-trained weights are publicly available, which is important to generate images without identifiable watermarks. Different generative models and source datasets may lead to other conclusions, which are left for future work.

Another limitation is the types of created datasets and tasks of experiments. Specifically, the datasets are created from synthetic prompts, and such a dataset generation scheme may be too simple to approximate possible data generation processes by users’ prompts. In addition, these datasets and tasks may often not be so complex that the insights of this paper would not cover some important aspects of other visual recognition tasks. For example, the object counting task [6] on contaminated data may be challenging because generative models cannot always correctly handle numbers [53]. We leave a further in-depth analysis for future research.

Additionally, our experiments of classification and generation tasks were limited to the “training from scratch” paradigm. Fine-tuning pre-trained models on carefully curated data might effectively circumvent the contamination issues.

7. Conclusion

Recent generative models trained with billion-scale data enable to generate high-quality and high-fidelity images, and many users play with these models to share generated images on the web. Observing such a trend, we questioned if such generated images affect the quality of future datasets collected images from the Internet. To answer this question, we simulated contamination of generated images using a state-of-the-art generative model and conducted experiments on such data in various tasks, namely, image classification, image captioning, and image generation. Throughout experiments, we found that generated images impact negatively on downstream performance, although its extent depends on the ratio of generated images and downstream tasks. Additional analysis revealed that generated images degrade robustness to out-of-distribution data; application of generated images to data augmentation needs careful consideration; easy detection of generated images may not be applicable to up-to-date generative models; and self-supervised learning may be a promising remedy to the problem.

Based on these observations, we recommend that researchers to publish generative models carefully implement watermarks to enable the identification of generated images. As we discussed in this paper, generated images have negative impacts on downstream performance, and their effect on new tasks is immeasurable; thus, publishers of generative models have the responsibility to avoid possible contamination. One simple way to avoid this problem is to implement either identifiable or invisible watermarks, as some publishers have already done, *e.g.*, [47, 51], then dataset curators can easily identify and filter them out. We also suggest that researchers who develop image datasets collected from the Internet should filter out or mark generated images, which may affect final downstream performance, because adding generated images may degenerate performance as shown in Section 5.4.

Another important implication of this paper is further research on the detection methods of generated images, in parallel with the development of generative models. As experimented in Section 6.1, generated images of the latest generative methods cannot be detected by simple methods that once had been effective. Consequently, their development for filtering is essential for the soundness of future research.

Acknowledgement

This work was supported by JST, ACT-X Grant Number JPMJAX210H, Japan. We used computational resources of “mdx: a platform for the data-driven future” and RAIDEN (Riken AIp Deep learning ENvironment). R.H. thanks Kai Katsumata at the University of Tokyo for his suggestions

on image generation experiments. We also appreciate constructive comments from the anonymous reviewers.

References

- [1] Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. SPICE: Semantic propositional image caption evaluation. In *ECCV*, 2016. 4, 6
- [2] Anthreas Antoniou, Amos Storkey, and Harrison Edwards. Data augmentation generative adversarial networks. In *ICLR*, 2018. 7
- [3] Yogesh Balaji, Seungjun Nah, Xun Huang, Arash Vahdat, Jiaming Song, Karsten Kreis, Miika Aittala, Timo Aila, Samuli Laine, Bryan Catanzaro, et al. eDiffi: Text-to-image diffusion models with an ensemble of expert denoisers. *arXiv*, 2022. 1, 2
- [4] Federico Bianchi, Pratyusha Kalluri, Esin Durmus, Faisal Ladhak, Myra Cheng, Debora Nozza, Tatsunori Hashimoto, Dan Jurafsky, James Zou, and Aylin Caliskan. Easily accessible text-to-image generation amplifies demographic stereotypes at large scale, 2022. 2
- [5] Keshigeyan Chandrasegaran, Ngoc-Trung Tran, and Ngai-Man Cheung. A closer look at fourier spectrum discrepancies for cnn-generated images detection. In *CVPR*, 2021. 8
- [6] Prithvijit Chattopadhyay, Ramakrishna Vedantam, Ramprasaath R Selvaraju, Dhruv Batra, and Devi Parikh. Counting everyday objects in everyday scenes. In *CVPR*, 2017. 8
- [7] Xinlei Chen, Hao Fang, Tsung-Yi Lin, Ramakrishna Vedantam, Saurabh Gupta, Piotr Dollar, and C. Lawrence Zitnick. Microsoft coco captions: Data collection and evaluation server. *arXiv*, 2015. 3
- [8] Xinyun Chen, Chang Liu, Bo Li, Kimberly Lu, and Dawn Song. Targeted backdoor attacks on deep learning systems using data poisoning. *arXiv*, 2017. 2
- [9] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009. 2, 3
- [10] Michael Denkowski and Alon Lavie. Meteor universal: Language specific translation evaluation for any target language. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, 2014. 6
- [11] Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat GANs on image synthesis. In *NeurIPS*, 2021. 1
- [12] Jesse Dodge, Maarten Sap, Ana Marasović, William Agnew, Gabriel Ilharco, Dirk Groeneveld, Margaret Mitchell, and Matt Gardner. Documenting large webtext corpora: A case study on the colossal clean crawled corpus. In *EMNLP*, 2021. 2
- [13] Alexey Dosovitskiy, Lucas Beyer, Alexander Kolesnikov, Dirk Weissenborn, Xiaohua Zhai, Thomas Unterthiner, Mostafa Dehghani, Matthias Minderer, Georg Heigold, Sylvain Gelly, Jakob Uszkoreit, and Neil Houlsby. An image is worth 16x16 words: Transformers for image recognition at scale. In *ICLR*, 2021. 8
- [14] Tarik Dzanic, Karan Shah, and Freddie Witherden. Fourier spectrum discrepancies in deep network generated images. In *NeurIPS*, 2020. 8

- [15] Logan Engstrom, Andrew Ilyas, Hadi Salman, Shibani Santurkar, and Dimitris Tsipras. Robustness, 2019. 4
- [16] Oran Gafni, Adam Polyak, Oron Ashual, Shelly Sheynin, Devi Parikh, and Yaniv Taigman. Make-a-scene: Scene-based text-to-image generation with human priors. In *ECCV*, 2022. 1, 2
- [17] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Joshua Bengio. Generative adversarial networks. In *NIPS*, 2014. 1, 2
- [18] Ian J Goodfellow, Jonathon Shlens, and Christian Szegedy. Explaining and harnessing adversarial examples. In *ICLR*, 2015. 2
- [19] Bo Han, Quanming Yao, Tongliang Liu, Gang Niu, Ivor W Tsang, James T Kwok, and Masashi Sugiyama. A survey of label-noise representation learning: Past, present and future. *arXiv*, 2020. 2
- [20] Drew Harwell. Dating apps need women. advertisers need diversity. AI companies offer a solution: Fake people. *Washington Post*, 2020. 1
- [21] Kaiming He, Xinlei Chen, Saining Xie, Yanghao Li, Piotr Dollár, and Ross Girshick. Masked autoencoders are scalable vision learners. In *CVPR*, pages 16000–16009, 2022. 8
- [22] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *CVPR*, 2016. 3
- [23] Dan Hendrycks, Steven Basart, Norman Mu, Saurav Kadavath, Frank Wang, Evan Dorundo, Rahul Desai, Tyler Zhu, Samyak Parajuli, Mike Guo, Dawn Song, Jacob Steinhardt, and Justin Gilmer. The many faces of robustness: A critical analysis of out-of-distribution generalization. In *ICCV*, 2021. 7
- [24] Dan Hendrycks, Kevin Zhao, Steven Basart, Jacob Steinhardt, and Dawn Song. Natural adversarial examples. In *CVPR*, 2021. 7
- [25] Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. GANs trained by a two time-scale update rule converge to a local nash equilibrium. In *NeurIPS*, 2017. 5
- [26] Kashmir Hill and Jeremy White. Designed to deceive: Do these people look real to you? *New York Times*, 2020. 1
- [27] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020. 1, 2
- [28] Jonathan Ho, Chitwan Saharia, William Chan, David J Fleet, Mohammad Norouzi, and Tim Salimans. Cascaded diffusion models for high fidelity image generation. *JMLR*, 23:47–1, 2022. 1
- [29] Peter J Huber. Robust statistics. In *International encyclopedia of statistical science*, pages 1248–1251. Springer, 2011. 2
- [30] Andrej Karpathy and Li Fei-Fei. Deep visual-semantic alignments for generating image descriptions. In *CVPR*, 2015. 3
- [31] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *CVPR*, 2019. 1
- [32] Diederik P Kingma and Max Welling. Auto-encoding variational Bayes. In *ICLR*, 2014. 1
- [33] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved precision and recall metric for assessing generative models. In *NeurIPS*, 2019. 5
- [34] Dongxu Li, Junnan Li, Hung Le, Guangsen Wang, Silvio Savarese, and Steven C. H. Hoi. Lavis: A library for language-vision intelligence. *arXiv*, 2022. 4
- [35] Junnan Li, Dongxu Li, Caiming Xiong, and Steven Hoi. BLIP: Bootstrapping language-image pre-training for unified vision-language understanding and generation. In *ICML*, 2022. 4
- [36] Wen Li, Limin Wang, Wei Li, Eirikur Agustsson, and Luc Van Gool. Webvision database: Visual learning and understanding from web data. *arXiv*, 2017. 3
- [37] Chin-Yew Lin. ROUGE: A package for automatic evaluation of summaries. In *Text Summarization Branches Out*, 2004. 6
- [38] Ze Liu, Yutong Lin, Yue Cao, Han Hu, Yixuan Wei, Zheng Zhang, Stephen Lin, and Baining Guo. Swin Transformer: Hierarchical vision transformer using shifted windows. In *ICCV*, 2021. 3
- [39] Zhuang Liu, Hanzi Mao, Chao-Yuan Wu, Christoph Feichtenhofer, Trevor Darrell, and Saining Xie. A convnet for the 2020s. In *CVPR*, 2022. 3
- [40] Alexander Quinn Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *ICML*, 2021. 5
- [41] Alexander Quinn Nichol, Prafulla Dhariwal, Aditya Ramesh, Pranav Shyam, Pamela Mishkin, Bob McGrew, Ilya Sutskever, and Mark Chen. GLIDE: Towards photorealistic image generation and editing with text-guided diffusion models. In *ICML*, 2022. 1, 2
- [42] Jonas Oppenlaender. The creativity of text-to-image generation. *arXiv*, 2022. 1, 2
- [43] Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. Bleu: a method for automatic evaluation of machine translation. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, 2002. 6
- [44] Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Kopf, Edward Yang, Zachary DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner, Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep learning library. In *NeurIPS*, 2019. 3
- [45] Bryan A. Plummer, Liwei Wang, Christopher M. Cervantes, Juan C. Caicedo, Julia Hockenmaier, and Svetlana Lazebnik. Flickr30k entities: Collecting region-to-phrase correspondences for richer image-to-sentence models. *IJCV*, 123(1):74–93, 2017. 3
- [46] Alec Radford, Luke Metz, and Soumith Chintala. Unsupervised representation learning with deep convolutional generative adversarial networks. In *ICLR*, 2015. 2
- [47] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical text-conditional image generation with clip latents. *arXiv*, 2022. 1, 2, 9
- [48] Spencer Rarrick, Chris Quirk, and Will Lewis. MT detection in web-scraped parallel corpora. In *Proceedings of Machine Translation Summit XIII: Papers*, 2011. 2

- [49] Suman Ravuri and Oriol Vinyals. Classification accuracy score for conditional generative models. In *NeurIPS*, 2019. [2](#)
- [50] Suman Ravuri and Oriol Vinyals. Seeing is not necessarily believing: Limitations of BigGANs for data augmentation. In *ICLR Learning from Limited Labeled Data Workshop*, 2019. [7](#)
- [51] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022. [1](#), [2](#), [3](#), [9](#)
- [52] Olga Russakovsky, Jia Deng, Hao Su, Jonathan Krause, Sanjeev Satheesh, Sean Ma, Zhiheng Huang, Andrej Karpathy, Aditya Khosla, Michael Bernstein, Alexander C. Berg, and Li Fei-Fei. ImageNet Large Scale Visual Recognition Challenge. *IJCV*, 2015. [2](#), [3](#)
- [53] Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily Denton, Seyed Kamyar Seyed Ghasemipour, Burcu Karagol Ayan, S Sara Mahdavi, Rapha Gontijo Lopes, et al. Photorealistic text-to-image diffusion models with deep language understanding. *arXiv*, 2022. [1](#), [2](#), [8](#)
- [54] Christoph Schuhmann, Romain Beaumont, Richard Vencu, Cade W Gordon, Ross Wightman, Mehdi Cherti, Theo Coombes, Aarush Katta, Clayton Mullis, Mitchell Wortsman, Patrick Schramowski, Srivatsa R Kundurthy, Katherine Crowson, Ludwig Schmidt, Robert Kaczmarczyk, and Jenia Jitsev. LAION-5B: An open large-scale dataset for training next generation image-text models. In *NeurIPS*, 2022. [3](#)
- [55] Michel Simard. Clean data for training statistical MT: the case of MT contamination. In *Proceedings of the 11th Conference of the Association for Machine Translation in the Americas: MT Researchers Track*, pages 69–82, 2014. [2](#)
- [56] Aman Sinha, Hongseok Namkoong, and John Duchi. Certifiable distributional robustness with principled adversarial training. In *ICLR*, 2018. [2](#)
- [57] Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *ICML*, 2015. [1](#), [2](#)
- [58] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. Score-based generative modeling through stochastic differential equations. In *ICLR*, 2021. [1](#), [2](#)
- [59] Nitasha Tiku. AI can now create any image in seconds, bringing wonder and danger. *Washington Post*, 2022. [2](#)
- [60] Toan Tran, Trung Pham, Gustavo Carneiro, Lyle Palmer, and Ian Reid. A Bayesian data augmentation approach for learning deep models. In *NIPS*, 2017. [7](#)
- [61] Ramakrishna Vedantam, C Lawrence Zitnick, and Devi Parikh. CIDEr: Consensus-based image description evaluation. In *CVPR*, 2015. [5](#), [6](#)
- [62] Zhisheng Xiao, Karsten Kreis, and Arash Vahdat. Tackling the generative learning trilemma with denoising diffusion GANs. In *ICLR*, 2022. [2](#)
- [63] Xin Yi, Ekta Walia, and Paul Babyn. Generative adversarial network in medical imaging: A review. *Medical Image Analysis*, 58:101552, 2019. [7](#)
- [64] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gungjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, Ben Hutchinson, Wei Han, Zarana Parekh, Xin Li, Han Zhang, Jason Baldridge, and Yonghui Wu. Scaling autoregressive models for content-rich text-to-image generation. *arXiv*, 2022. [1](#), [2](#)
- [65] Han Zhang, Weichong Yin, Yewei Fang, Lanxin Li, Boqiang Duan, Zhihua Wu, Yu Sun, Hao Tian, Hua Wu, and Haifeng Wang. ERNIE-ViLG: Unified generative pre-training for bidirectional vision-language generation. *arXiv*, 2021. [1](#), [2](#)