

# Image retrieval outperforms diffusion models on data augmentation

Anonymous authors

Paper under double-blind review

## Abstract

Diffusion models excel at generating photorealistic images from text-queries. Naturally, many approaches have been proposed to use these generative abilities to augment training datasets for downstream tasks, such as classification. However, diffusion models are themselves trained on large datasets, often with noisy annotations. It is an open question to which extent diffusion models are useful to augment data for improved downstream classification performance. In particular, it is unclear if they generalize enough to improve over directly using the additional data of their pre-training process for augmentation. We perform a systematic evaluation of existing methods to generate images from diffusion models and study new extensions to assess their benefit for data augmentation. We find that personalizing diffusion models towards the target data outperforms simpler prompting strategies. However, using the pre-training data of the diffusion model alone, via a simple nearest neighbor retrieval procedure, leads to even stronger downstream performance. Overall, our study explores the potential of diffusion models in generating new training data and at the same time surprisingly finds that these sophisticated models are not yet able to beat a simple and strong image retrieval baseline on simple downstream vision tasks.

## 1 Introduction

Data augmentation is a key component of training robust and high-performing computer vision models and given its success, it is becoming increasingly sophisticated: From the early simple image transformations (random cropping, flipping, color jittering, and shearing) (Cubuk et al., 2020), over augmenting additional training data by combining pairs of images, such as MixUp (Zhang et al., 2017) and CutMix (Yun et al., 2019), all the way to image augmentations using generative models. Augmentation via image transformations improves robustness towards distortions that resemble the transformation (Wenzel et al., 2022b) and interpolating augmentations are particularly helpful in situations where diverse training data is scarce (Ghiasi et al., 2021). With the success of generative adversarial networks (GANs), generative models finally scaled to high-dimensional domains and allowed the generation of photorealistic images. The idea of using them for data augmentation purposes has been prevalent since their early successes and forms the basis of a new set of data augmentation strategies (Zietlow et al., 2022; Ghosh et al., 2022; Antoniou et al., 2017; Frid-Adar et al., 2018; Esteban et al., 2017; Motamed et al., 2021; Mariani et al., 2018; Ramaswamy et al., 2021).

Given the emergence of diffusion models (DMs) that outperform GANs in terms of visual quality and diversity (Ramesh et al., 2022; Saharia et al., 2022b; Rombach et al., 2022), using them for data augmentation is a natural next step. Unlike GANs, diffusion models can be easily conditioned using text queries, which allows for more controlled data generation. These models are trained on a large dataset of billions of filtered image text pairs retrieved from the internet, enabling them to generate images of unparalleled variety.

We benchmark a variety of existing augmentation techniques based on diffusion models and surprisingly find that these techniques are outperformed by simply retrieving nearest neighbors from the DM’s training dataset with simple prompts (using the same CLIP-like model (Radford et al., 2021) used in the DM). We propose new extensions and variants of diffusion model based strategies each leading to improvements, however, not beating our simple retrieval baseline. We show that simple text prompts based on class




Figure 1: *Are generative methods beneficial for data augmentation?* Each bar shows the accuracy (along with the standard error) of the downstream classifier based on the best augmentation methods within each family. While diffusion model based techniques (orange and green) improve over the baseline of only using the original 10% ImageNet data (blue), a simple, more computationally efficient retrieval method using images from the diffusion model’s pre-training dataset directly (red) performs best. This suggests that the generative capabilities of diffusion models for augmentation have not yet been fully leveraged.

labels suffice for conditioning the DMs to improve the performance of standard classifiers compared to the unaugmented data set. As images generated by simple prompts match the training distribution of the DM and not the training distribution of the classifiers, we test—inspired by related work on personalized DMs—methods that fine-tune the DM conditioning and optionally the DMs denoising model component. These fine-tuned models outperformed the best prompting strategies for their ability to create even better synthetic data for augmentation (Figure 1). Although all investigated methods based on generating synthetic images were sophisticated and compute intensive, none of them reached the performance of simple nearest neighbor retrieval (Figure 1), suggesting that the true potential of DMs for data augmentation is not yet fully realized.

## 2 Related work

**Data augmentation using generative adversarial networks.** Data augmentation is widely used when training deep networks. It overcomes some challenges associated with training on small datasets and improving the generalization of the trained models (Zhang et al., 2017). Manually designed augmentation methods have limited flexibility, and the idea of using ML-generated data for training has attracted attention. Generative adversarial networks (GAN) such as BigGAN (Brock et al., 2018) have been used to synthesize images for ImageNet classes (Besnier et al., 2020; Li et al., 2022). Despite early promising results, the use of GANs to generate synthetic training data has shown limited advantages over traditional data augmentation methods (Zietlow et al., 2022). Diffusion models, on the other hand, might be a better candidate since they are more flexible via general-purpose text-conditioning and exhibit a larger diversity and better image quality (Sohl-Dickstein et al., 2015; Ramesh et al., 2022; Saharia et al., 2022b; Rombach et al., 2022). Hence, in this work, we focus on diffusion models.

**Data augmentation using diffusion models.** Recently, diffusion models showed astonishing results for synthesizing images (Sohl-Dickstein et al., 2015; Ramesh et al., 2022; Saharia et al., 2022b; Rombach et al., 2022). Numerous approaches have been published that adapt diffusion models to better fit new images and can be used for augmentation. We evaluate methods that employ diffusion models in a guidance-free manner or prompt them without adapting the model: Luzi et al. (2022) generate variations of a given dataset by first adding noise to the images and then denoising them again, and Sariyildiz et al. (2022) generate a synthetic ImageNet clone only using the class names of the target dataset. We also evaluate methods for specializing (AKA personalizing) diffusion models into our study. Given a few images of the same object (or concept),

Gal et al. (2022) learn a joint word embedding (pseudo word) that reflects the subject and can be used to synthesize new variations of it (e.g., in different styles) and Gal et al. (2023) recently extended their method to significantly reduce the number of required training steps. Kavar et al. (2022) follow a similar approach and propose a method for text-conditioned image editing by fine-tuning the diffusion model and learning a new word embedding that aligns with the input image and the target text. We investigate the usefulness of “personalization” for data augmentation, which was not conclusively addressed in the original papers, and propose and evaluate extensions of these methods tailored to data augmentation and provide a thorough evaluation in a unified setting. Additionally, we benchmark all of these approaches against our suggested retrieval baseline.

Given the fast moving field, there have been multiple concurrently proposed methods to synthesize images for various downstream tasks that we could not include in our evaluation of their helpfulness for data augmentation to train an image classifier. Some methods focus on fine-tuning the diffusion model learning a unique identifier to personalize the DM to the given subject (Ruiz et al., 2022), Shipard et al. (2023) create synthetic clones of CIFAR (Krizhevsky et al., 2009) and EuroSAT (Helber et al., 2019) for zero-shot classification, employ alternative losses for image generation to improve few-shot learning (Roy et al., 2022), and Ghalebikesabi et al. (2023) create synthetic, privacy-preserving clones of medical data. Other methods optimize the features of the embedded images to augment specifically small datasets (Zhang et al., 2022) and use alternative prompts to the diffusion model with class descriptions generated by a language model (He et al., 2022). Trabucco et al. (2023) additionally employ image synthesis, image editing (Meng et al., 2021) and in painting (Lugmayr et al., 2022; Saharia et al., 2022a), Azizi et al. (2023) combine frozen prompt-conditioning with fine-tuning only the diffusion model, Bansal & Grover (2023) generate new images conditioning the DM by prompts and example images leading to degraded performance when used to augment ImageNet, and other work focuses on medical data and sample curation (Akrouit et al., 2023).

### 3 Experimental protocol

We augment using a wide-range of generative and retrieval based techniques, evaluating performance on a downstream classification task (Figure 2A).

**Dataset.** *ImageNet.* Unless explicitly stated, we simulate training data in a low-data regime, we sample 10% of the ImageNet Large Scale Visual Recognition Challenge 2012 (ILSVRC2012) (Russakovsky et al., 2015) training split retaining class imbalance. As the retrieval method did not return sufficient samples for 10 of the 1,000 ImageNet classes, they were excluded from augmentation, model training and evaluation. We additionally sample a disjoint set of images of the same size from the training split for hyperparameter optimization and model selection, which we refer to as the validation split. For results on the full ImageNet data set, we include all but our validation split to the training set. All trained classification models are evaluated on the original ILSVRC2012 validation split. *Caltech256.* We additionally verify our findings on Caltech256 (Griffin et al., 2007) excluding 5 of the 256 classes for which retrieval did not return sufficient samples. We randomly sampled 80% of the data as train split and randomly partitioned the remaining 20% of the samples into equally sized disjoint validation and test splits.

**Data augmentation and classifier training.** For each augmentation strategy on ImageNet, we generate 390 samples per class ensuring that the number of images at least tripled per class, as in our subsampled dataset each class contains between 74 and 130 images. On Caltech256 we generate three-times as many images per class as in the training set. We resample the additional data into 5 sets containing the same number of samples per class as our target dataset and train a ResNet-50 classifier (He et al., 2016) on each to derive variance estimates. During training, the samples are further augmented by random resizing and cropping as is typical in ImageNet training. On ImageNet we trained the classifier with batch size 256 and a learning rate of 0.1, on Caltech256 we used a batch size of 1024 and after an initial learning rate optimization sweep we set the learning rate to 0.003. We divided the learning rate by 10 when validation accuracy failed to improve for 10 epochs. We stopped training when the validation accuracy did not improve for 20 epochs or after at most 270 epochs and used the highest validation accuracy checkpoint for final scoring. Each model was trained on 8 NVIDIA T4 or V100 GPUs using distributed data parallel training.

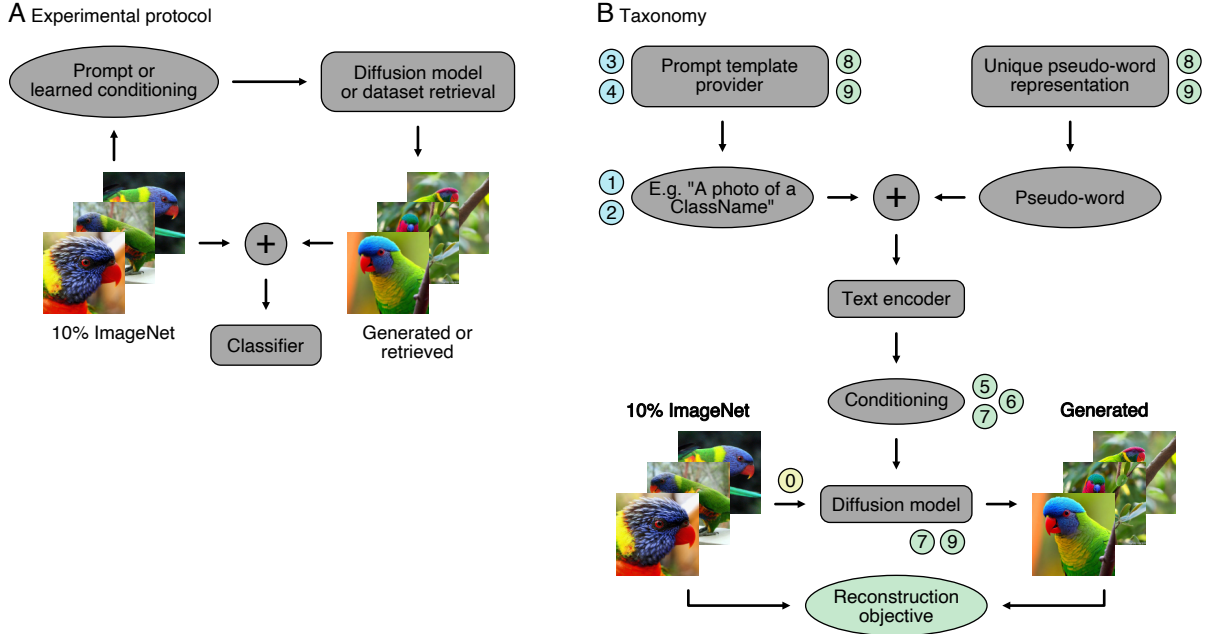


Figure 2: *Experimental protocol and taxonomy of diffusion model based augmentation methods.* **A.** We generate images by guiding the Stable Diffusion model by text prompts or by learned conditioning, or retrieving images by a nearest neighbor search in the CLIP embedding space of the DM’s training data. We then train the downstream classifier on the original 10% ImageNet data augmented by the additional data and evaluate on the original ImageNet validation split. **B.** All considered methods adapt different components of the prompting, conditioning mechanisms, and the fine-tuning of the diffusion model. We reference each method by a circled number (see section 3.1 for details). Some methods edit the prompts while keeping the DM frozen: ①, ② use a single prompt for each class and ③, ④ use multiple prompts from a set of templates. Another family of methods optimize the conditioning vectors for the given images: ⑤, ⑥ only optimize the conditioning vector keeping the DM frozen, while ⑦ also jointly fine-tunes the DM. Instead of optimizing the conditioning vector, ⑧ learns a pseudo-word description of the class using multiple prompts keeping the DM frozen, while ⑨ additionally fine-tunes the DM. ⑩ does not adapt any component of the DM and relies on encoding and decoding to create variations of the given images.

### 3.1 Augmentation methods

The benchmarked augmentation methods can be grouped into four categories: (1) guidance-free diffusion model sampling, (2) simple conditioning techniques with prompts based on the objects’ class label, and (3) personalization techniques that fine-tune the diffusion model conditioning and optionally the diffusion model itself to the classifier’s data domain. We compare diffusion model approaches to a simple baseline (4) using images retrieved from the dataset that the diffusion model was trained on.

**Unconditional generation.** Following the procedure of ⑩ BOOMERANG (Luzi et al., 2022), we investigated a guidance-free method that does not require conditioning or updating the diffusion model. Instead, the approach adds noise to individual samples before denoising them.

**Prompt conditioning.** We explore several prompt-based methods of guiding the DM to produce samples for a specific class (Figure 2B). ① SIMPLE PROMPT: We condition the model by simple prompts containing the object’s class name  $n$ , prompting the DM with “A photo of  $n$ .” and a version of it, ② SIMPLE PROMPT (NO WS), stripping whitespace,  $w(\cdot)$ , from class names, “A photo of  $w(n)$ .” ③ CLIP PROMPTS: We add sampling prompts from the set of CLIP (Radford et al., 2021) text-encoder templates, e.g. “a photo of many



$w(n)$ ,” “a black and white photo of the  $w(n)$ ,” etc. and ④ SARIYILDIZ ET AL. PROMPTS: a set of templates proposed to create a synthetic ImageNet clone (Sariyildiz et al., 2022).

**Fine-tuning the diffusion model.** We explore various methods for fine-tuning a DM for class-personalized sampling to improve reconstruction of the classification dataset (Figure 2B). ⑤ FT CONDITIONING: freezing the DM and optimizing one conditioning per class. ⑥ FT CLUSTER CONDITIONING: optimizing multiple conditionings per class instead of just one. ⑦ inspired by IMAGIC (Kawar et al., 2022), jointly fine-tuning the conditioning and the DM’s denoising component. ⑧ TEXTUAL INVERSION (Gal et al., 2022): instead of fine-tuning the conditioning, sampling prompt templates and combining them with optimizing a pseudo-word representing the class-concept. ⑨ PSEUDOWORD+DM: combining the previous approach with optimizing the DM’s denoising component.

**Laion nearest neighbor retrieval.** As a baseline comparison, we propose using ⑩ RETRIEVAL to select images from the Laion dataset used to train the diffusion model. This method finds nearest neighbor images to the SIMPLE PROMPT (NO WS) class name prompts in the CLIP embedding space.

### 3.2 Implementation and training details

**Diffusion model backbone.** We used the pretrained Stable Diffusion v1.4 network, based on a latent diffusion architecture (Rombach et al., 2022). We discarded generated images if marked as NSFW by the provided safety checker, replacing them with new samples. As Stable Diffusion was trained on image sizes of 512px, we kept this resolution for all methods.

**Prompt generation.** For prompt-based sampling methods, we generated prompts based on the ImageNet class names, defined by WordNet (Miller, 1995) synsets representing distinct entities in the WordNet graph. Each synset consists of one or multiple lemmas describing the class, where each lemma can consist of multiple words, e.g., “Tiger shark, Galeocerdo Cuvieri”. We link each class via its synset to its class name. If a synset consists of multiple lemmas, we separate them by a comma, resulting in prompts like “A photo of tiger shark, Galeocerdo Cuvieri,” as we found that providing multiple lemmas led to better performance than using only the first lemma of a synset. Whenever methods inserted the class name into prompt templates, we sampled the templates randomly with replacement. Sariyildiz et al. (Sariyildiz et al., 2022) provided multiple categories of prompt templates (e.g. class name only, class name with WordNet hypernyms, additionally combined with “multiple” and “multiple different” specifications, class name with WordNet definition, and class name with hypernyms and randomly sampled backgrounds from the places dataset (Zhou et al., 2017)). Here, we sampled for each category the same number of images and randomly across the background templates. For Caltech256 we used the class names as provided by the data set.

**Fine-tuning.** For all methods that require additional fine-tuning, we trained the model with the default Stable Diffusion optimization objective (Rombach et al., 2022) until the validation loss stagnated or increased – no model was trained for more than 40 epochs. When fine-tuning the diffusion models on full ImageNet, the number of training samples increased 10-fold, thus, to keep the computational budget comparable across modalities, we stopped training after 3 epochs. We set hyperparameters in accordance with published works (Gal et al., 2022; Kawar et al., 2022; Luzi et al., 2022) or existing code where available. Scaling to larger batch sizes was implemented by square-root scaling the learning rate to ensure constant gradient variance, and we performed a fine-grained grid-search for the optimal learning rate. For the FT CLUSTER CONDITIONING models, we clustered the training images within a class using k-means on the inception v3 embeddings. The conditioning vectors were initialized by encoding the SIMPLE PROMPT (NO WS) prompts with the DM’s text encoder. For textual inversion (Gal et al., 2022) we fine-tuned the text embedding vector corresponding to the introduced pseudoword and sample from the provided text templates (Gal et al., 2022) to optimize the reconstruction objective for our ImageNet subset, freezing all other model components. We initialized the pseudo-word embedding with the final word in the first synset lemma (i.e., for the class “tiger shark” we used “shark”). Where this initialization resulted in multiple initial tokens, we initialized with the mean of the embeddings. For each fine-tuning run we stored checkpoints with the best train and validation loss and those corresponding to 1, 2 and 3 epochs of training. In the case of IMAGIC (Kawar et al.,

Table 1: *Overall performance of the data augmentation methods.* We report the top-1 accuracy and standard error of the downstream classifier trained on the augmented 10% ImageNet subsplit. The methods are described in section 3.1 and are grouped into families as described in the beginning of section 4. The circles refer to the taxonomy in Figure 2B.

Augmentation method	Accuracy (%)
<b>10% ImageNet</b>	<b>57.2 <math>\pm</math> 0.2</b>
<b>20% ImageNet</b>	<b>70.2 <math>\pm</math> 0.3</b>
<b>Boomerang</b> (Luzi et al., 2022) ①	<b>56.3 <math>\pm</math> 0.3</b>
Simple prompt (no ws) ②	60.0 $\pm$ 0.3
Simple prompt ①	60.1 $\pm$ 0.2
Sariyildiz et al. prompts (Sariyildiz et al., 2022) ④	60.8 $\pm$ 0.2
<b>CLIP prompts</b> ③	<b>60.9 <math>\pm</math> 0.2</b>
FT conditioning ⑤	60.8 $\pm$ 0.1
FT cluster conditioning ⑥	60.9 $\pm$ 0.2
Imagic (conditioning & DM; Kawar et al. (2022)) ⑦	61.0 $\pm$ 0.3
Textual inversion (pseudoword; Gal et al. (2022)) ⑧	61.0 $\pm$ 0.4
<b>Pseudoword+DM</b> (combining Gal et al. (2022); Kawar et al. (2022)) ⑨	<b>61.2 <math>\pm</math> 0.2</b>
<b>Retrieval (our suggested baseline)</b> ⑩	<b>62.6 <math>\pm</math> 0.1</b>

(2022) we found that the validation loss was still decreasing after 40 epochs, however, as the quality of the reconstructed images significantly deteriorated after 8 epochs, we instead stored checkpoints for the first 8 epochs. This is similar to the Imagic training scheme, which optimized the embedding for 100 steps and the U-Net for 1,500 steps on a single image. Although we follow existing procedure as closely as possible, image quality might improve for longer model training runs. We selected the final model checkpoint based on the validation accuracy of a ResNet-50 classifier trained on the union of augmented samples and the target data set.

**Nearest neighbor retrieval.** For RETRIEVAL, we used Laion 5b (Schuhmann et al., 2022), a publicly available dataset of 5 billion image-caption pairs extracted via web-crawler. The data was then filtered by only retaining images where the CLIP image embedding was consistent with the caption embedding. This filtering acts as a weak form of supervision, that retains those images where CLIP is more likely to work.

The dataset provides a CLIP embedding nearest neighbors search index for each instance and an official package<sup>1</sup> allows for fast search and retrieval. We used this to retrieve 130 images per class for ImageNet and the same number as samples per class for Caltech256. Images with a Laion aesthetics score of less than 5 were discarded to allow a fair comparison with Stable Diffusion 1.4 which was trained on this subset. For a fair comparison to our generative augmentation methods, we used the same safety checker model to discard images that were marked as NSFW. Due to changes in the availability of the images at the URLs in the dataset and the described filtering steps, it is often necessary to retrieve more than the desired number of images, which we do by increasing the number of nearest neighbors gradually from 1.4·130 to 10·1.4·130 when not enough samples were found. To avoid using the same image multiple times, we applied the duplicate detector of the clip-retrieval package, however, this does not detect all near duplicate images and some duplicates are still used.

## 4 Results

Table 1 provides a results summary of the augmented ImageNet subsplit in which each block corresponds to the groups of augmentation methods introduced in Section 3.1. In the following, section 4.1 introduces the

<sup>1</sup><https://github.com/rom1504/clip-retrieval/releases/tag/2.35.1>




Figure 3: *Example images generated by BOOMERANG* (Luzi et al., 2022). The generated images lack diversity and effectively only add noise to the original ones. **A.** Example images from 10% ImageNet. **B.** The augmentations produced by Boomerang. **C.** There is only a small difference between the original and augmented images (best viewed on screen).

baselines, section 4.2 benchmarks the previously suggested unconditional DM augmentation method and a method proposed to create a synthetic ImageNet clone (Sariyildiz et al., 2022) against our simple RETRIEVAL baseline method, section 4.3 discusses improvements of conditioning the DM with text prompts, and section 4.4 discusses personalization approaches to diffusion models. Finally, section 4.5 provides additional insights.

#### 4.1 Baselines

To establish the baseline classifier performance, we train on our 10% ImageNet subset (no added data), achieving a top-1 accuracy of  $57.2 \pm 0.2$ . Since all augmentation methods double the size of the dataset, we consider an upper bound performance by using 20% of original ImageNet (no added data) with a classifier accuracy of  $70.2 \pm 0.3$ .

#### 4.2 Retrieval outperforms previously suggested diffusion-based augmentation

**Unconditional generation.** BOOMERANG (Luzi et al., 2022) generates augmented images by sequentially adding Gaussian noise and uses a diffusion model to denoise elements of the target dataset, without altering the diffusion model or prompts. This results in a lower top-1 accuracy of  $56.3 \pm 0.3$  compared to the original subset of ImageNet. Figure 3 shows that this method does not induce significant diversity in the dataset and only slightly distorts the original images, leading to an overall decrease in performance. However, since this method is directly applied on the target dataset samples, it does not suffer from domain shift or class ambiguity.

**Prompt-conditioning.** Sariyildiz et al. (2022) generate augmented training examples with a diffusion model adapted to the target dataset via conditioning mechanisms. Their method guides image generation by various prompts based on the class name (see section 3.2 for details) and achieves  $60.8 \pm 0.2$  classification accuracy, performing better than the 10% ImageNet baseline and worse than the 20% ImageNet upper bound (Table 1).

**Retrieving from the diffusion model pre-training data.** We now established that diffusion models are helpful for creating augmentation data, however, it is unclear how much value the DM’s generative capabilities add compared to simply using their pre-training data directly for augmentation. To answer this question, we propose a simple RETRIEVAL method fetching images from the DM’s pre-training data that are semantically closest to the SIMPLE PROMPT (NO WS) prompts (see section 3.2 for details). Augmenting with this data outperformed the sophisticated diffusion model based approaches at  $62.6 \pm 0.1$  top-1 accuracy, while being computationally less demanding.

**Summary.** Using diffusion models off the shelf, i.e., without adaptation or conditioning using the target dataset, is not beneficial and can even deteriorate the classifier performance compared to the ImageNet baseline. Augmenting with generated samples adapted to the target dataset by prompt-conditioning performs better than the unaugmented baseline. However, simply augmenting with images retrieved from the diffusion



Figure 4: *Example images obtained from the investigated augmentation methods.* **A.** 10% ImageNet original images. **B.-D.** Images generated by our prompt-based sampling techniques: (B) SIMPLE PROMPT, (C) SIMPLE PROMPT (NO WS) and (D) CLIP PROMPTS, respectively, where in (D) we show images to the prompts “a bad photo of a  $w(n)$ ”, “a black and white photo of the  $w(n)$ ”, “a cartoon  $w(n)$ ”, and “a photo of many  $w(n)$ ”. **E.-H.** Images generated by the fine-tuned diffusion models, specifically (E) FT CONDITIONING, (F) TEXTUAL INVERSION, (G) IMAGIC, and (H) PSEUDOWORD+DM. **I.** Examples of RETRIEVAL from the diffusion model’s training set. Best viewed when zoomed in.

model’s training data outperforms the DM based approaches, indicating that DM generated images have significant shortcomings. As retrieved images are real-world images, they show good photorealism, variety, and detail. However, retrieval underperforms the 20% ImageNet upper bound performance, which might be caused by retrieval suffering from class ambiguity (“papillon” and “mailbag, postbag” in Figure 4I), reflecting a mismatch of concepts from CLIP latent space to ImageNet classes as we measured semantic similarity as distance of CLIP embeddings.

### 4.3 Advanced prompt-conditioning improves but is not competitive with retrieval

Then, we embarked on a quest asking if we can improve diffusion model based methods over the strong RETRIEVAL baseline. To do so, we first explored several variations of text prompt conditioning strategies.

**Simple prompt conditioning.** We investigate a SIMPLE PROMPT conditioning method that uses the prompt “A photo of  $n$ .”, where  $n$  is replaced by the class name, e.g., “A photo of tiger shark, Galeocerdo cuvieri.” This improved over the 10% ImageNet baseline to a top-1 classification accuracy of  $60.1 \pm 0.2$ , but Figure 4B shows clear problems with class ambiguity and a lack of diversity.

**Tackling class ambiguity by white space removal.** While generally mitigating class ambiguity is a hard problem, we focus on the ambiguity introduced by class names composed of multiple words. To this end, we investigated the variant SIMPLE PROMPT (NO WS) (no white space) that removed the white space from the class name used by SIMPLE PROMPT, e.g., “A photo of desktopcomputer.”. While class ambiguity reduced for some classes (e.g., desktop computer; Figure 4C) it was ineffective for others (e.g., papillon)



and overall resulted in slightly degraded performance  $60.0 \pm 0.3$  compared to keeping the white space. We explore more advanced methods in the following sections.

**Improving sample diversity by diverse prompt templates.** To increase the diversity of the samples, we made use of multiple prompt templates. We use CLIP PROMPTS which randomly selects one of the text templates provided by CLIP (Radford et al., 2021) (see Figure 4D for examples and Radford et al. (2021) for a full list), increasing classification accuracy to  $60.9 \pm 0.2$ . This method performed best among all prompt-based techniques. Interestingly, the overall performance increased even though some prompts lead to synthetic images that did not match the style of ImageNet samples (e.g., “a cartoon photo of the papillon.”) or images with texture-like contents instead of objects (e.g., papillon, image to the bottom right in Figure 4D). Surprisingly, these slightly more elaborate templates just adding few more words compared to SIMPLE PROMPT (NO WS) (e.g., “a bad photo of a papillon.”) improved class ambiguity in some cases (e.g., papillon; Figure 4D).

**Summary.** Augmenting a dataset with images sampled by prompt-based conditioning techniques improves the downstream classifier performance beyond the unaugmented datasets, however, none of these methods improve over RETRIEVAL. Inspecting the generated samples, we find that various challenges remain: ImageNet is more than a decade old, and some images included in the dataset are older and of different style when compared to the images created by the recently trained Stable Diffusion model (e.g., desktop computer; Figure 4A-D). In other cases, generated images do not match the domain of ImageNet samples, for instance because they are too artificial, sometimes even like a computer rendering (e.g., motor scooter; Figure 4A-D), their texture does not match (e.g., papillon; Figure 4A,D) or the prompt does not match the desired style (e.g., cartoon; Figure 4A,D). This tendency may have been amplified by Stable Diffusion v1.4 being trained on a subset of Laion, which was filtered to contain only aesthetic images.<sup>2</sup>

#### 4.4 Personalizing the diffusion model improves further but does not outperform retrieval

In the previous section, we found that adapting the diffusion models by only editing the prompt offers limited improvement when used for augmentation. On our quest to improve diffusion models to narrow down the performance gap to the RETRIEVAL baseline, this section explores a more advanced set of methods that additionally fine-tune parts of the diffusion model to “personalize” it to the 10% ImageNet training images by optimizing the DM reconstruction objective (c.f. Figure 2B). Originally, these methods were proposed in the context of personalizing the model to a specific object or concept, i.e., a single or only a few images.

**Fine-tuning the conditioning vectors.** The general diffusion model architecture has multiple components that can be fine-tuned (see Figure 2B). We start by fine-tuning one conditioning vector per class to optimize the reconstruction objective for our ImageNet subsample, while keeping the other model weights fixed. This method, dubbed FT CONDITIONING (fine-tune conditioning), achieves an augmentation accuracy of  $60.8 \pm 0.1$  and is on par with the best prompt editing method CLIP PROMPTS. Interestingly, while this method achieves good augmentation performance, the generated images do not look photorealistic, and suffer from noise and missing backgrounds (Figure 4E).

**Fine-tuning clusters of conditioning vectors.** Using a single conditioning vector for all images from one class might be insufficient to capture the full class variability. To see if this is the case, we explore FT CLUSTER CONDITIONING. We generate  $k$  clusters of images per class, before fitting a conditioning vector to each individual cluster (see section 3.2 for details). We find that using  $k = 5$  clusters slightly improves the accuracy by 0.1 percentage points over using  $k = 1$ . For larger  $k = 10$  and  $k = 15$ , performance decreased ( $60.8 \pm 0.2$  and  $60.6 \pm 0.3$  accuracy). This might be due to the smaller number of images per cluster, making the fine-tuning more prone to over-fitting.

**Textual inversion.** To reduce the number of unrealistic images, we explored a variant of TEXTUAL INVERSION (Gal et al., 2022). This method learns a pseudoword  $n$  representing the class concept combined

<sup>2</sup><https://github.com/CompVis/stable-diffusion>

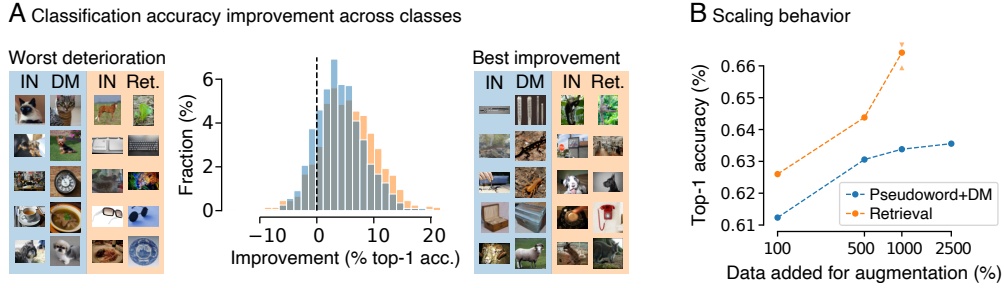


Figure 5: *Retrieval and diffusion model performance across classes and scaling behavior.* **A.** The distribution of classifier improvement for each class is shown for the best-performing DM-based method PSEUDOWORD+DM (blue) and RETRIEVAL (orange). For each class the improvement is computed by comparing the downstream classification performance using augmentations compared to only using the original 10% ImageNet samples. Images to the left and right show examples of classes with worst deterioration and greatest improvement. **B.** RETRIEVAL outperforms synthetic images across augmentation ranges. For 1000% augmentation we could not retrieve sufficient samples for 7 additional classes. Assuming classifier performance of 0% or 100% on these classes leads to best- ( $\blacktriangledown$ ) and worst-case ( $\blacktriangle$ ) estimates.

with a randomly drawn textual description of the generated image style (e.g., “a photo of a  $n$ ”, “a rendering of a  $n$ ”, etc.; see Gal et al. (2022) for a full list). While this method improves the photorealism of the generated samples (Figure 4F), it only slightly improves the augmentation accuracy over the simple fine-tuning of conditioning vectors (FT CONDITIONING) by 0.1 percentage points (Table 1).

**Fine-tuning the denoising.** We now explore fine-tuning the DMs denoising module jointly with the conditioning vector. This idea stems from IMAGIC (Kawar et al. 2022) and results in an augmentation accuracy of  $61.0 \pm 0.3$  (examples in Figure 4G), which is on par with TEXTUAL INVERSION.

Finally, we combine the best-performing methods: jointly optimizing a pseudo-word per class (TEXTUAL INVERSION (Gal et al. 2022)) and the DMs denoising module (IMAGIC (Kawar et al. 2022)). We denote this method by PSEUDOWORD+DM and it resulted in an accuracy of  $61.2 \pm 0.2$ , outperforming all other DM based techniques investigated. The generated images also show improved photorealism over TEXTUAL INVERSION and IMAGIC (Figure 4H).

**Summary.** Personalizing diffusion models improves in matching the domain of the 10% ImageNet images better (e.g., generated images for “desktop computer” now look similarly old as in ImageNet), improve upon prompt-based techniques to reduce class-ambiguity (e.g. papillon and desktop computer showed ambiguity in Figure 4B but not in panels E-H), show good sample variety and the best-performing PSEUDOWORD+DM method enhances photorealism and reduces the artistic style of generated images. This model performs best across all generative augmentation techniques investigated, and suggests that we can leverage personalization techniques to combine the DM’s knowledge of billions of annotated images with learning the domain distribution of the data we want to augment. At the same time, none of these sophisticated methods outperformed the simpler and computationally cheaper RETRIEVAL baseline.

#### 4.5 When does retrieval help?

**Retrieval helps for most classes.** To better understand the failure cases of the augmentation methods, we check for each class in ImageNet if the augmentation method is beneficial. To this end, for each class, we compute the performance improvement of the downstream classifier using the augmentation method compared to only using the original ImageNet samples. Figure 5A shows the distribution of improvements for the RETRIEVAL augmentation method and the best DM-based method PSEUDOWORD+DM. Both methods improve the performance for most of the classes, however, there are still many classes where the performance is decreased up to 10%. The distribution of improvements for RETRIEVAL is similarly shaped as for PSEU-



DOWORD+DM, but significantly pushed to the right. Systematically investigating these failure cases might be a fruitful avenue to further improve generative augmentation.

**Retrieval has better scaling behavior.** As diffusion models can generate arbitrarily many samples, we explored our best performing diffusion model based method’s (PSEUDOWORD+DM) scaling behavior for higher augmentation ratios. To compare its performance to RETRIEVAL, we retrieved abundant samples and augmented with those images which CLIP embeddings were most similar to our retrieval prompt. While for DM generated images performance saturated for high augmentation ratios, retrieval still substantially increased performance, outperforming the best diffusion model based approach (Figure 5B).

**Retrieval works on full data sets.** So far, all presented results address the case of a small data set, simulated by drawing a subsample from ImageNet. Hence, we asked how well the best prompt-conditioned and personalized DM augmentation techniques would perform compared to retrieval on large data sets. On full ImageNet, we found consistent with the results on subsampled ImageNet that RETRIEVAL performed best ( $79.0 \pm 0.2$ ), however, the performance gap to PSEUDOWORD+DM ( $78.8 \pm 0.2$ ) and CLIP PROMPTS ( $78.7 \pm 0.1$ ) narrowed down compared to the small data set case we investigated before. All three approaches outperform the ResNet-50 baseline (76.1). Additionally, we verified these results on Caltech256 (Griffin et al., 2007) (RETRIEVAL:  $68.4 \pm 0.4$ , PSEUDOWORD+DM:  $68.1 \pm 0.3$ , CLIP PROMPTS:  $68.2 \pm 0.2$ , no augmentation:  $59.3 \pm 0.7$ ) showing that diffusion models do not outperform our simple RETRIEVAL baseline across data sets.

## 5 Discussion

Diffusion models have shown their effectiveness in many application areas, and using them for data augmentation is an intriguing research direction. We have evaluated multiple methods to prompt and personalize diffusion models on their usefulness for data augmentation, showing that none of them beat the simple baseline of retrieving images from the diffusion model’s pre-training dataset. Why is the retrieval baseline so hard to beat? One reason might be that retrieval potentially accesses more information, as the pre-training dataset is usually much larger than the weights of the generative models trained on it. Although it has been argued that diffusion models might partially compress the training data (Somepalli et al., 2022; Carlini et al., 2023), it is still unclear if the generative model captured all relevant information. However, diffusion models could possibly improve upon the so-far superior retrieval by generating a large number of additional data and more diverse and compositionally novel images, for instance by generating out-of-domain samples (e.g., “a photograph of an astronaut riding a horse”) (Ramesh et al., 2022; Saharia et al., 2022b; Rombach et al., 2022). Furthermore, diffusion models allow, in principle, for more controlled adaptation than retrieval methods. We showed that personalization methods are a good step in this direction, however, they typically focus only on creating variants of specific given images. To unlock the true potential of diffusion models for data augmentation, new methods that capture the target dataset manifold as a whole, are needed.

**Limitations and future work.** It would be interesting to extend our analysis to other datasets which might not be abundantly available for RETRIEVAL (e.g. medical images) and to explore the investigated methods for out-of-distribution generalization (Qiu et al., 2022; Wenzel et al., 2022a). Diffusion models are a very active research area, and new methods and applications are published on a daily basis. Hence, our experiments could only capture a subset of possible methods and newly proposed extensions of them (in total, eleven methods). We have shown that simple retrieval is a very strong competitor for data augmentation. Further improvements could be introduced by diversifying the retrieval set (Wenzel et al., 2020; Yue & Guestrin, 2011) or retrieving images using linear combinations of inputs (Zietlow et al., 2022).

**Conclusion.** We showed that a simple retrieval baseline outperforms a wide range of diffusion model based augmentations. However, given the fast rate of progress in this field it is not possible to definitively say that retrieval can not be beaten, and we believe that diffusion models have the potential to improve over this baseline. Nonetheless, the strength of retrieval’s performance makes it clear that future works using diffusion models for augmentation should also compare against this baseline. We hope that our paper provides ground for researchers to benchmark generative augmentation methods and assess their benefit by comparing them

with retrieval baselines. In the longer-term, we hope that new methods can be developed which combine retrieval and generation, leading to greater improvements in the diversity and quality of augmented images.

## References

- Mohamed Akrouf, Bálint Gyepesi, Péter Holló, Adrienn Poór, Blága Kincső, Stephen Solis, Katrina Cirone, Jeremy Kawahara, Dekker Slade, Latif Abid, Máté Kovács, and István Fazekas. Diffusion-based Data Augmentation for Skin Disease Classification: Impact Across Original Medical Datasets to Fully Synthetic Images, January 2023. URL <http://arxiv.org/abs/2301.04802>. arXiv:2301.04802 [cs, eess].
- Antreas Antoniou, Amos Storkey, and Harrison Edwards. Data augmentation generative adversarial networks. *arXiv preprint arXiv:1711.04340*, 2017.
- Shekoofeh Azizi, Simon Kornblith, Chitwan Saharia, Mohammad Norouzi, and David J. Fleet. Synthetic Data from Diffusion Models Improves ImageNet Classification, April 2023. URL <http://arxiv.org/abs/2304.08466>. arXiv:2304.08466 [cs].
- Hritik Bansal and Aditya Grover. Leaving Reality to Imagination: Robust Classification via Generated Datasets, February 2023. URL <http://arxiv.org/abs/2302.02503>. arXiv:2302.02503 [cs].
- Victor Besnier, Himalaya Jain, Andrei Bursuc, Matthieu Cord, and Patrick Pérez. This dataset does not exist: training models from generated images. In *ICASSP 2020-2020 IEEE International Conference on Acoustics, Speech and Signal Processing (ICASSP)*, pp. 1–5. IEEE, 2020.
- Andrew Brock, Jeff Donahue, and Karen Simonyan. Large scale gan training for high fidelity natural image synthesis. *arXiv preprint arXiv:1809.11096*, 2018.
- Nicholas Carlini, Jamie Hayes, Milad Nasr, Matthew Jagielski, Vikash Sehwal, Florian Tramèr, Borja Balle, Daphne Ippolito, and Eric Wallace. Extracting Training Data from Diffusion Models, January 2023. URL <http://arxiv.org/abs/2301.13188>. arXiv:2301.13188 [cs].
- Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition workshops*, pp. 702–703, 2020.
- Cristóbal Esteban, Stephanie L Hyland, and Gunnar Rätsch. Real-valued (medical) time series generation with recurrent conditional gans. *arXiv preprint arXiv:1706.02633*, 2017.
- Maayan Frid-Adar, Idit Diamant, Eyal Klang, Michal Amitai, Jacob Goldberger, and Hayit Greenspan. Gan-based synthetic medical image augmentation for increased cnn performance in liver lesion classification. *Neurocomputing*, 321:321–331, 2018.
- Rinon Gal, Yuval Alaluf, Yuval Atzmon, Or Patashnik, Amit H. Bermano, Gal Chechik, and Daniel Cohen-Or. An Image is Worth One Word: Personalizing Text-to-Image Generation using Textual Inversion, August 2022. URL <http://arxiv.org/abs/2208.01618>. arXiv:2208.01618 [cs].
- Rinon Gal, Moab Arar, Yuval Atzmon, Amit H. Bermano, Gal Chechik, and Daniel Cohen-Or. Designing an Encoder for Fast Personalization of Text-to-Image Models, February 2023. URL <http://arxiv.org/abs/2302.12228>. arXiv:2302.12228 [cs].
- Sahra Ghalebikesabi, Leonard Berrada, Sven Goyal, Ira Ktena, Robert Stanforth, Jamie Hayes, Soham De, Samuel L. Smith, Olivia Wiles, and Borja Balle. Differentially Private Diffusion Models Generate Useful Synthetic Images, February 2023. URL <http://arxiv.org/abs/2302.13861>. arXiv:2302.13861 [cs, stat].
- Golnaz Ghiasi, Yin Cui, Aravind Srinivas, Rui Qian, Tsung-Yi Lin, Ekin D Cubuk, Quoc V Le, and Barret Zoph. Simple copy-paste is a strong data augmentation method for instance segmentation. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 2918–2928, 2021.

- Partha Ghosh, Dominik Zietlow, Michael J Black, Larry S Davis, and Xiaochen Hu. Invgan: invertible gans. In *Pattern Recognition: 44th DAGM German Conference, DAGM GCPR 2022, Konstanz, Germany, September 27–30, 2022, Proceedings*, pp. 3–19. Springer, 2022.
- Gregory Griffin, Alex Holub, and Pietro Perona. Caltech-256 object category dataset. 2007.
- Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep residual learning for image recognition. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, pp. 770–778, 2016.
- Ruifei He, Shuyang Sun, Xin Yu, Chuhui Xue, Wenqing Zhang, Philip Torr, Song Bai, and Xiaojuan Qi. Is synthetic data from generative models ready for image recognition?, October 2022. URL <http://arxiv.org/abs/2210.07574>. arXiv:2210.07574 [cs].
- Patrick Helber, Benjamin Bischke, Andreas Dengel, and Damian Borth. Eurosat: A novel dataset and deep learning benchmark for land use and land cover classification. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing*, 12(7):2217–2226, 2019.
- Bahjat Kawar, Shiran Zada, Oran Lang, Omer Tov, Huiwen Chang, Tali Dekel, Inbar Mosseri, and Michal Irani. Imagic: Text-Based Real Image Editing with Diffusion Models, October 2022. URL <http://arxiv.org/abs/2210.09276>. arXiv:2210.09276 [cs].
- Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- Daiqing Li, Huan Ling, Seung Wook Kim, Karsten Kreis, Sanja Fidler, and Antonio Torralba. Bigdataset-gan: Synthesizing imagenet with pixel-wise annotations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 21330–21340, 2022.
- Andreas Lugmayr, Martin Danelljan, Andres Romero, Fisher Yu, Radu Timofte, and Luc Van Gool. RePaint: Inpainting using Denoising Diffusion Probabilistic Models. In *2022 IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*, pp. 11451–11461, June 2022. doi: 10.1109/CVPR52688.2022.01117. ISSN: 2575-7075.
- Lorenzo Luzi, Ali Siahkoohi, Paul M. Mayer, Josue Casco-Rodriguez, and Richard Baraniuk. Boomerang: Local sampling on image manifolds using diffusion models, October 2022. URL <http://arxiv.org/abs/2210.12100>. arXiv:2210.12100 [cs, stat].
- Giovanni Mariani, Florian Scheidegger, Roxana Istrate, Costas Bekas, and Cristiano Malossi. Bagan: Data augmentation with balancing gan. *arXiv preprint arXiv:1803.09655*, 2018.
- Chenlin Meng, Yutong He, Yang Song, Jiaming Song, Jiajun Wu, Jun-Yan Zhu, and Stefano Ermon. Sdedit: Guided image synthesis and editing with stochastic differential equations. In *International Conference on Learning Representations*, 2021.
- George A Miller. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41, 1995.
- Saman Motamed, Patrik Rogalla, and Farzad Khalvati. Data augmentation using generative adversarial networks (gans) for gan-based detection of pneumonia and covid-19 in chest x-ray images. *Informatics in Medicine Unlocked*, 27:100779, 2021.
- Jielin Qiu, Yi Zhu, Xingjian Shi, Florian Wenzel, Zhiqiang Tang, Ding Zhao, Bo Li, and Mu Li. Are multimodal models robust to image and text perturbations? *arXiv preprint arXiv:2212.08044*, 2022.
- Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever. Learning Transferable Visual Models From Natural Language Supervision. *arXiv:2103.00020 [cs]*, February 2021. URL <http://arxiv.org/abs/2103.00020>. arXiv: 2103.00020.
- Vikram V Ramaswamy, Sunnie SY Kim, and Olga Russakovsky. Fair attribute classification through latent space de-biasing. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*, pp. 9301–9310, 2021.

- Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical Text-Conditional Image Generation with CLIP Latents, April 2022. URL <http://arxiv.org/abs/2204.06125>. arXiv:2204.06125 [cs].
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10684–10695, 2022.
- Aniket Roy, Anshul Shah, Ketul Shah, Anirban Roy, and Rama Chellappa. DiffAlign : Few-shot learning using diffusion based synthesis and alignment, December 2022. URL <http://arxiv.org/abs/2212.05404>. arXiv:2212.05404 [cs].
- Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dream-Booth: Fine Tuning Text-to-Image Diffusion Models for Subject-Driven Generation, August 2022. URL <http://arxiv.org/abs/2208.12242>. arXiv:2208.12242 [cs].
- 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. *International Journal of Computer Vision (IJCV)*, 115(3):211–252, 2015. doi: 10.1007/s11263-015-0816-y.
- Chitwan Saharia, William Chan, Huiwen Chang, Chris Lee, Jonathan Ho, Tim Salimans, David Fleet, and Mohammad Norouzi. Palette: Image-to-Image Diffusion Models. In *ACM SIGGRAPH 2022 Conference Proceedings*, SIGGRAPH ’22, pp. 1–10, New York, NY, USA, July 2022a. Association for Computing Machinery. ISBN 978-1-4503-9337-9. doi: 10.1145/3528233.3530757. URL <https://doi.org/10.1145/3528233.3530757>.
- 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, Tim Salimans, Jonathan Ho, David J. Fleet, and Mohammad Norouzi. Photorealistic Text-to-Image Diffusion Models with Deep Language Understanding, May 2022b. URL <http://arxiv.org/abs/2205.11487>. arXiv:2205.11487 [cs].
- Mert Bulent Sariyildiz, Kartteek Alahari, Diane Larlus, and Yannis Kalantidis. Fake it till you make it: Learning(s) from a synthetic ImageNet clone, December 2022. URL <http://arxiv.org/abs/2212.08420>. arXiv:2212.08420 [cs].
- 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 *Thirty-sixth Conference on Neural Information Processing Systems Datasets and Benchmarks Track*, 2022. URL <https://openreview.net/forum?id=M3Y74vmsMcY>.
- Jordan Shipard, Arnold Wiliem, Kien Nguyen Thanh, Wei Xiang, and Clinton Fookes. Diversity is Definitely Needed: Improving Model-Agnostic Zero-shot Classification via Stable Diffusion, April 2023. URL <http://arxiv.org/abs/2302.03298>. arXiv:2302.03298 [cs].
- Jascha Sohl-Dickstein, Eric Weiss, Niru Maheswaranathan, and Surya Ganguli. Deep unsupervised learning using nonequilibrium thermodynamics. In *International Conference on Machine Learning*, pp. 2256–2265. PMLR, 2015.
- Gowthami Somepalli, Vasu Singla, Micah Goldblum, Jonas Geiping, and Tom Goldstein. Diffusion Art or Digital Forgery? Investigating Data Replication in Diffusion Models, December 2022. URL <http://arxiv.org/abs/2212.03860>. arXiv:2212.03860 [cs].
- Brandon Trabucco, Kyle Doherty, Max Gurinas, and Ruslan Salakhutdinov. Effective Data Augmentation With Diffusion Models, February 2023. URL <http://arxiv.org/abs/2302.07944>. arXiv:2302.07944 [cs].

- Florian Wenzel, Jasper Snoek, Dustin Tran, and Rodolphe Jenatton. Hyperparameter ensembles for robustness and uncertainty quantification. In *Advances in Neural Information Processing Systems*, 2020.
- Florian Wenzel, Andrea Dittadi, Peter Vincent Gehler, Carl-Johann Simon-Gabriel, Max Horn, Dominik Zietlow, David Kernert, Chris Russell, Thomas Brox, Bernt Schiele, Bernhard Schölkopf, and Francesco Locatello. Assaying Out-Of-Distribution Generalization in Transfer Learning, October 2022a. URL <http://arxiv.org/abs/2207.09239>. arXiv:2207.09239 [cs, stat].
- Florian Wenzel, Andrea Dittadi, Peter Vincent Gehler, Carl-Johann Simon-Gabriel, Max Horn, Dominik Zietlow, David Kernert, Chris Russell, Thomas Brox, Bernt Schiele, et al. Assaying out-of-distribution generalization in transfer learning. *arXiv preprint arXiv:2207.09239*, 2022b.
- Yisong Yue and Carlos Guestrin. Linear submodular bandits and their application to diversified retrieval. In *Advances in Neural Information Processing Systems*, 2011. URL <https://proceedings.neurips.cc/paper/2011/file/33ebd5b07dc7e407752fe773eed20635-Paper.pdf>.
- Sangdoo Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *Proceedings of the IEEE/CVF international conference on computer vision*, pp. 6023–6032, 2019.
- Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical risk minimization. *arXiv preprint arXiv:1710.09412*, 2017.
- Yifan Zhang, Daquan Zhou, Bryan Hooi, Kai Wang, and Jiashi Feng. Expanding Small-Scale Datasets with Guided Imagination, November 2022. URL <http://arxiv.org/abs/2211.13976>. arXiv:2211.13976 [cs].
- Bolei Zhou, Agata Lapedriza, Aditya Khosla, Aude Oliva, and Antonio Torralba. Places: A 10 million image database for scene recognition. *IEEE transactions on pattern analysis and machine intelligence*, 40(6): 1452–1464, 2017.
- Dominik Zietlow, Michael Lohaus, Guha Balakrishnan, Matthäus Kleindessner, Francesco Locatello, Bernhard Schölkopf, and Chris Russell. Leveling down in computer vision: Pareto inefficiencies in fair deep classifiers. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 10410–10421, 2022.