
Fill-Up: Balancing Long-Tailed Data with Generative Models

Joonghyuk Shin

Minguk Kang

Jaesik Park

Pohang University of Science and Technology (POSTECH), South Korea
{jhshin4727,mgkang,jaesik.park}@postech.ac.kr

Abstract

Modern text-to-image synthesis models have achieved an exceptional level of photorealism, generating high-quality images from arbitrary text descriptions. In light of the impressive synthesis ability, several studies have exhibited promising results in exploiting generated data for image recognition. However, directly supplementing data-hungry situations in the real-world (*e.g.* few-shot or long-tailed scenarios) with existing approaches result in marginal performance gains, as they suffer to thoroughly reflect the distribution of the real data. Through extensive experiments, this paper proposes a new image synthesis pipeline for long-tailed situations using Textual Inversion. The study demonstrates that generated images from textual-inverted text tokens effectively aligns with the real domain, significantly enhancing the recognition ability of a standard ResNet50 backbone. We also show that real-world data imbalance scenarios can be successfully mitigated by filling up the imbalanced data with synthetic images. In conjunction with techniques in the area of long-tailed recognition, our method achieves state-of-the-art results on standard long-tailed benchmarks when trained from scratch.

1 Introduction

The intricate landscape of real-world data often reveals an inherent imbalance, where the abundance of common samples dominate the scarce presence of their rare counterparts [1, 2, 3, 4, 5, 6, 7, 8]. Long-tailed recognition is an area that aims to acquire balanced knowledge in this lopsided situation, irrespective of different sample sizes for each class. Several works tried to mitigate this class imbalance challenge through diverse approaches such as class-balanced sampling [9, 10, 11, 12], loss re-weighting [13, 14, 15, 16, 17] and data augmentation [9, 11, 18, 19, 20, 21, 22, 23, 24, 25]. Recent approaches also include decoupling methods [26, 27, 28], and the integration of mixture-of-experts techniques [29, 30, 31, 32].

Given that the underlying motivation of these approaches revolves around re-balancing data distributions, an intuitive approach would be to fill up the imbalanced data using synthetic samples. However, training traditional class conditional generative models under severe data imbalance is known to be highly challenging [33, 34, 35, 36, 37], hindering the practical usages of well-established generative models such as VAE [38] and GAN [39].

The rapid development of generative models [40, 41, 42, 43, 44, 45, 46, 47], has ushered in advancements across diverse areas of machine learning. Synthetic data generation is one area where these photorealistic images with delicate details exhibit promising results. Multiple studies [48, 49, 50, 51] have proposed efficient methods for zero/few-shot classification, elucidating the robustness of features acquired from synthetic data and their efficacy in transfer learning scenarios. However, these zero-shot methods, predominantly reliant on prompting techniques, often exhibit marginal performance gains in the presence of real data, constraining their practical usages in real-world long-tailed scenarios.

In this paper, with the help of recent large-scale text-to-image synthesis models, we effectively address this issue by adopting per-class optimization of randomly initialized text tokens. This work starts by evaluating a range of novel and established strategies devised for synthetic image generation. After a thorough examination, we empirically discover that our new approach based on Textual Inversion [52] exhibits a remarkable ability to generate diverse images that align with the real domain without using class-related text information. Notably, by exclusively leveraging information within the image domain, our method showcases its capability to generate classes that were underrepresented by the original text-conditioned pretrained model. In conjunction with the simple yet effective Balanced Softmax loss [17] and the proposed two-stage training procedure, our method attains state-of-the-art results on standard long-tailed classification benchmarks [4, 3] when trained from scratch. Our method especially demonstrates significantly higher accuracies in few-shot scenarios involving classes with fewer than 20 samples.

Overall, our work makes the following key contributions: (1) We offer a comprehensive evaluation of various generation strategies based on text-to-image synthesis models, incorporating diverse ablations and providing valuable insights. (2) We propose a novel image generation method tailored for real-world’s data imbalance scenarios involving Textual Inversion, which outperforms existing generation methods. (3) In combination with our suggested two-stage training procedure and the Balanced Softmax loss, our method achieves state-of-the-art results on standard long-tailed benchmarks when trained from scratch.

2 Related work

2.1 Large scale text-to-image synthesis models

Recent advances in diffusion models [40, 53, 54, 55] and multimodal learning [56, 57] have led to significant progress in the area of text-conditioned image synthesis, which was considered to be extremely challenging with traditional conditional GANs [39, 58, 59, 60, 61]. Some notable examples of diffusion models include GLIDE [41], Dall-E 2 [44], Imagen [45], Stable Diffusion [43], and eDiff-I [46]. Autoregressive and GAN-based models such as Dall·E [42], Parti [62], GigaGAN [47], and StyleGAN-T [63] have also shown very promising results. In conjunction with the advancements of generative models, there has been extensive research on methods to personalize synthesized outputs. Representative works include Textual Inversion [52], DreamBooth [64], and Custom Diffusion [65], each focusing on fine-tuning different parts of diffusion models.

Among several options, we utilize Stable Diffusion, a text-conditioned variant of the Latent Diffusion Model [43] (LDM), trained on LAION-5B [66]. LDM-based models leverage pretrained encoder and decoder networks to perform forward and reverse diffusion processes in the latent space, reducing the expensive compute overhead of the diffusion family. In this paper, we utilize Stable Diffusion v1.5 [67], unless stated otherwise.

2.2 Learning from synthetic data

Building real-world datasets is often challenging due to the difficulty of collecting a balanced set of images, as well as the inherent privacy and bias concerns involved. Simulation pipelines from 3D engines [68, 69] or pretrained class-conditional GANs have been utilized as data sources. For instance, GANs can generate pixel-wise annotations for semantic segmentation [70, 71] and improve representation learning through different view generation [72, 73]. Nonetheless, synthetic data generated from 3D engines may have domain gaps with real data and show limited diversity. Additionally, training generative models from scratch usually requires a large amount of well-curated data, further restricting the applicability of existing methods in data-scarce situations.

More recent approaches [48, 49, 50, 74, 51] incorporate general-purpose text-conditioned diffusion models. He *et al.* [48] employed GLIDE to generate synthetic images and used them to fine-tune CLIP [56]. Sariyildiz *et al.* [49] focused more on ImageNet [75] and trained models from scratch using various prompting methods. Their common findings suggest that features learned from the synthetic data can be robustly transferred to downstream tasks. These results highlight the crucial roles of classifier-free guidance scale and the number of images in achieving high diversity and classification accuracy. Different from the above approaches, Li *et al.* [76] directly turned pretrained diffusion models into zero-shot classifiers, and Azizi *et al.* [74] fine-tuned Imagen [45] with ImageNet [75] to

produce class-conditional model that can potentially improve ImageNet classification. The use of synthetic data from diffusion models is also becoming prevalent in other computer vision domains, such as semantic segmentation [77], object detection [78], and semi-supervised learning [79].

However, the synthetic images generated through the aforementioned methods, which target zero-shot scenarios, present an issue of distribution misalignment [48]. Consequently, the direct application of these zero-shot methods in the presence of real data (*e.g.* long-tailed or few-shot) leads to the generation of noisy samples that deviate from the distribution of the real data. As a result, despite the availability of valuable real samples, these approaches struggle to further improve the performance, as discussed in Sec. 4.3.

Considering the challenge of learning from imbalanced data for generative models [35, 36, 37], simple strategies such as training or fine-tuning the entire generative model under imbalanced data, are also not practical. Moreover, employing a naïve per-class fine-tuning approach would result in considerable inefficiency, as it necessitates a substantial amount of model weights to be saved with an increasing number of classes. In this study, we effectively circumvent these issues by adopting an efficient per-class tuning strategy for text tokens.

2.3 Long-tailed recognition

Re-sampling and re-weighting. The classical approach for handling long-tailed data involves re-sampling and re-weighting. Re-sampling methods aim to generate class-balanced data by over-sampling of minority classes [12, 80] and under-sampling of majority classes [12, 81]. However, it is known that over-sampling can lead to overfitting on tail classes [82], while under-sampling may result in losing valuable information from head classes [24]. Re-weighting methods tackle this issue by assigning different weights on either the class level [14, 83, 84, 16, 85] or the instance level [86, 87], resulting in distinct losses for each class. Nevertheless, it is known that re-weighting can introduce instability during the training, particularly when dealing with the extreme class imbalance in large-scale datasets.

Data augmentation. Augmenting real samples [19, 20, 21] or introducing synthetic samples [23, 88, 89] is also promising direction. Modern augmentation strategies involve transferring information from the majority classes to the minority classes [24]. In the domain of synthetic sample generation, some studies [9, 25] have shown successful results by generating samples at the feature level. Closer to our work, GAMO [88] adopts a three-player adversarial game to generate feature-level samples and uses VAE [38] to generate image-level samples. However, given the easy-to-collapse nature of the adversarial game and the difficulty of learning imbalanced distributions, GAMO struggles to generate high-quality images, resulting in a marginal performance gain that is further constrained to smaller datasets. To the best of our knowledge, our approach is the first to generate realistic image-level samples that can actually perform on par or even surpass other state-of-the-art long-tailed methods.

Other long-tailed methods. Recently, decoupling the learning of representation and classifier, following [26], has become one of the mainstream approaches [27, 28]. Another advantage of the decoupling method is that it can be readily combined with other techniques, such as re-sampling or re-weighting, further amplifying its potential. Meta-learning [17, 90, 91, 92], contrastive learning [93, 94], and ensembling [29, 30, 31, 32] based methods are additional sources of recent successes. The most recent approaches [95, 96, 97] involve using a powerful multimodal model, CLIP [56]. By incorporating additional text data acquired from the web (*e.g.*, *Wikipedia*) [96] or additional prompt-tuning [97], recent works further improved the performance of the pretrained CLIP. However, these methods heavily rely on the strong zero-shot performance of CLIP, which already surpasses the performances of most state-of-the-art methods. In this work, we mainly focus on learning imbalanced distribution from scratch, with a specific emphasis on acquiring robust representation exclusively in the domain of images.

3 Filling up long-tailed data with generative models

3.1 Synthetic image generation strategies

Our study begins by evaluating diverse synthetic data generation approaches and selecting the best fit for our data-scarce scenario. As pointed out by previous work [49, 50, 74], the problem of diversity

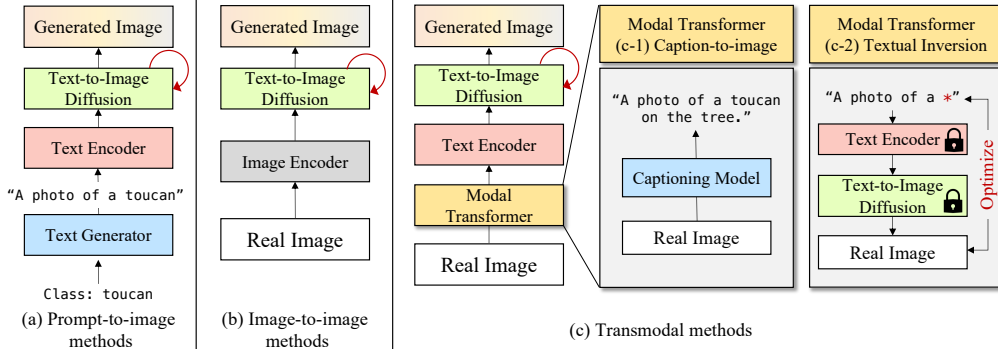


Figure 1: Illustration of our different generation strategies: prompt-to-image methods (*Single template*, *CLIP template*, *T5* [100], and *Flan-T5 XXL* [101]), image-to-image methods (*Image variation* [102] and *Reimagine* [103]), and transmodal methods (*Captioning* [104] and *Textual Inversion* [52]).

has been problematic in utilizing text-to-image synthesis models for image recognition. Since cutting-edge text-to-image synthesis models prioritize aligning the visual appearance with input text prompts, they tend to generate high-fidelity but low-diversity images. A straightforward approach to balance fidelity and diversity is adjusting the classifier-free guidance scale [98] as shown in Fig. 2a.

Recent models that employ classifier-free guidance, such as Stable Diffusion, provide the capability to generate more diverse but slightly noisier samples by decreasing the default guidance scale w , which is set to 7.5 in the case of Stable Diffusion. Eq. (1) demonstrates the process of classifier-free guidance in LDM variants [43], where ϵ represents a denoising U-Net, t denotes a timestamp, z refers to the latent variable, and c is a conditioning vector. Higher guidance enforces stronger conditioning of input prompts, and it is worth noting that classifier-free guidance is no longer adopted when the scale is set to 1.0, reducing the necessity of an additional reverse process for null-conditioning.

$$\tilde{\epsilon}_t = \epsilon_\theta(z_t) + w(\epsilon_\theta(z_t, c) - \epsilon_\theta(z_t)). \quad (1)$$

To effectively address the issue of diversity, we thoroughly evaluate a range of novel and established methods designed to generate feature-rich images and extensively compare their performances at different guidance scales. For evaluation of the mentioned methods, we use IN100 and IN100-LT from Sec. 4.1, generating 1,300 images per class using DDIM sampler [99] with 50 diffusion steps. We present comprehensive qualitative results in Appendix.

Prompt-to-image based methods. The most straightforward method to utilize text-to-image synthesis models is to use naïve prompts (e.g. "a photo of a {CLS}"). However, this method, namely, *single template*, is known to produce images that substantially lack diversity. He *et al.* [48] and Yuan *et al.* [105] introduce the use of large-scale language models and CLIP templates to alleviate this issue. He *et al.* use T5 [100] fine-tuned on CommonGen [106] combined with CLIP filtering to remove noisy generated sentences. Building upon previous work, we experiment with the same T5 [107], a larger language model (Flan-T5 XXL [101] fine-tuned in Alpaca style [108, 109]) and a community-famous prompt-extender [110] for Stable Diffusion which is a GPT-2 model fine-tuned on prompts from DiffusionDB [111]. Table 1 presents the results for four settings: *single template*, *CLIP templates*, *T5*, and *Flan-T5 XXL*. Additional details (e.g. decoding strategies) and the performance for *Prompt-Extend* can be found in Appendix. We generally observe enhancements with more diverse prompts when the scale is $w = 7.5$, but *single template* performs the best as the scale becomes $w = 1.0$. With more noise engaged, we believe promoting diversity via complex prompts harms precision by compromising important details of the original class.

Image-to-image based methods. While the above methods can generally be considered as zero-shot methods, the approaches from here require real samples. For image-to-image based methods, we utilize the state-of-the-art Stable Diffusion Reimagine [103], which replaces the CLIP text encoder with the CLIP image encoder. Since the source images are fully encoded and the reverse process does not originate from the noised source images, this method can produce images that differ in both overall structure and high-frequency details. For a fair comparison, we also test a similar variant called Image Variation [102] based on Stable Diffusion v1.3, as Reimagine is based on Stable




Figure 2: (a) Impact of classifier-free guidance scale on top-1 and top-5 accuracy. (b) Improvements in top-1 accuracy when scaling number of training samples per class for different generation strategies. (c, d) Performance of *Textual Inversion* based method with varying optimization steps used for generation and increasing size of optimizable text embedding.

Table 1: Comparison of accuracies among diverse generation methods for IN100. Each method generated 130K images (1,300 images per class) and synthetic images were exclusively used to train ResNet50 [112]. Low guidance indicates a scale of 1.0 and high guidance denotes 7.5, except for *Image Variation* [102] (3.0) and *Reimagine* [103] (10.0).

Method (Top-1/Top-5)	Prompt-to-Image				Image-to-Image		Transmodal	
	Single Template	CLIP Templates	T5	Flan-T5 -XXL	Image Variation	Reimagine	Image Captioning	Textual Inversion
Low Guidance	59.9 / 83.7	57.2 / 82.2	45.8 / 71.5	51.6 / 77.3	47.2 / 77.1	54.2 / 80.7	55.2 / 83.7	62.8 / 86.2
High Guidance	41.5 / 67.5	48.7 / 73.6	36.0 / 61.0	41.3 / 67.3	42.2 / 70.9	41.5 / 68.7	45.6 / 73.6	49.2 / 74.3

Diffusion v2.1. We adopt real samples of IN100-LT described in Sec 4.1 as source images to evaluate the robustness of these methods in real-world data-scarce scenarios. Experimental results of these two settings (*Reimagine* and *Image Variation*) are elaborated in Table 1. Broadly speaking, these methods perform worse than *single template*, and through qualitative review, we attribute this to specific circumstances where a "reimagined" object replaces the crucial component of an image containing class-specific information. This observation aligns with the fact that this method achieves a relatively high top-5 accuracy considering its performance on top-1. Technical details and qualitative results are provided in Appendix.

Transmodal methods. Here, we propose two novel methods that first generate prompts from real samples, then subsequently synthesize images from these prompts. We refer to this process as a "Transmodal" or simply "image-to-prompt-to-image" approach. The first method utilizes the state-of-the-art captioning model BLIP2 [104] to generate prompts and uses these prompts for sampling. The second method utilizes Textual Inversion [52]. Using few sets of images, we optimize a text token "*" per every class and generate images based on the prompt "a photo of a *". The results of *Captioning* and *Textual Inversion* can be found in Table 1, while additional details are provided in Appendix. We observe that *Textual Inversion* consistently outperforms all other methods by a notable margin.

Textual Inversion details. Textual Inversion [52] is an optimization-based personalization method designed to effectively learn the optimal text vector that best explains the concept of given samples. Keeping Stable Diffusion’s encoder (\mathcal{E}), denoising U-net (ϵ_θ), and CLIP text encoder (c_θ) frozen, Textual Inversion aims to optimize the text embedding (v_*), which is a continuous vector representation of the text token (y), using the following objective.

$$v_* = \arg \min_v \mathbb{E}_{z \sim \mathcal{E}(x), y, \epsilon \sim \mathcal{N}(0,1), t} \left[\|\epsilon - \epsilon_\theta(z_t, t, c_\theta(y))\|_2^2 \right]. \quad (2)$$

Our empirical results demonstrate that Textual Inversion can efficiently learn class-level concepts given numerous images (*i.e.* head classes) and also effectively capture decisive features given few number of images (*i.e.* tail classes), which is evident considering its initial tailoring for personalization tasks involving 3-5 images. Furthermore, the experiment in Fig. 2c reveals that class-level information is already well-learned in the early stages (500 steps) of optimization, which reduces the need for prolonged training. From this observation, we adopt a simple heuristic based on the number of images that restricts excessive training time and leverage checkpoints from various steps to enhance diversity. Lastly, we observe from Fig. 2d that enlarging the token embedding’s capacity from a single vector

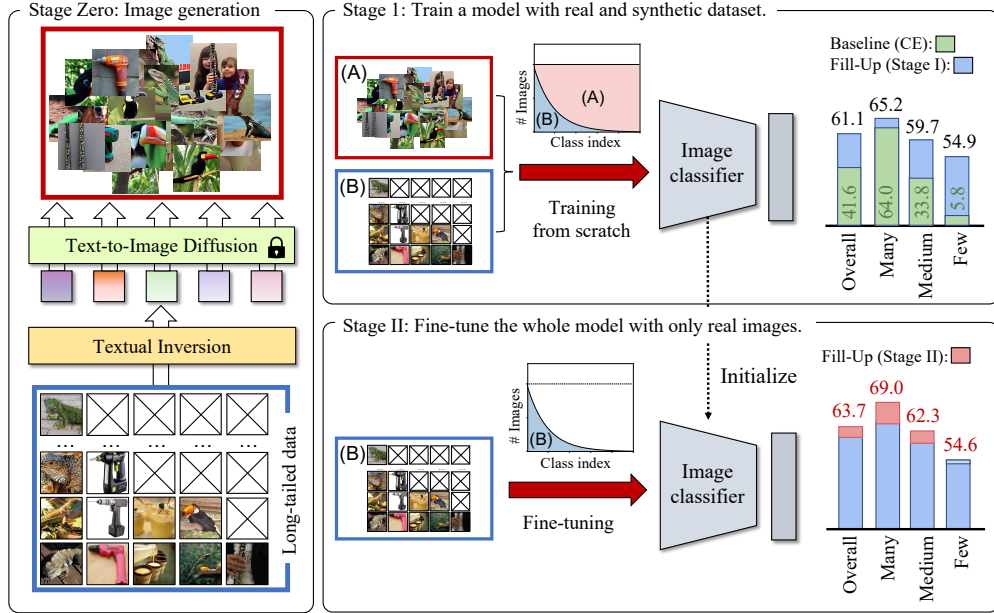


Figure 3: Overview of our proposed long-tailed image recognition pipeline on ImageNet-LT. During the Stage I, the image classifier is trained on balanced dataset, obtained by filling up long-tailed data with synthetic samples. Stage II fine-tunes the image classifier with long-tailed real dataset along with Balanced Softmax loss [17].

of 1×768 to 4×768 significantly improves top-1 accuracy, allowing finer details to be captured [113]. Yet, further increasing the size of embedding (e.g. 10×768) did not yield additional gains, suggesting that a text embedding of shape 4×768 can adequately represent most classes.

Starting from a random text token "*", Textual Inversion consistently proves to be effective for long-tailed data by generating images that align well with the target domain, regardless of the sample size, without relying on any class-related text information. Especially, the high scaling efficiency of this method described in Fig. 2b, provides compelling evidence of the strong potential inherent in textual-inverted tokens. Our qualitative results even suggest its capability to generate classes that were previously underrepresented by conventional text-conditioning. More information on the mentioned heuristic and training, as well as supplementary qualitative results, can be found in Appendix.

3.2 Two-stage training procedure

One of the mainstream approaches in long-tailed recognition is to decouple representation learning and classifier learning as suggested by Kang *et al.* [26]. Motivated by this idea, we adopt a two-stage training procedure. During the first stage, the entire model is trained on balanced dataset, obtained by filling up long-tailed data with synthetic samples.

In the second stage, we fine-tune the model with only real samples on top of robust feature learned from the first stage [48, 49]. Interestingly, while cRT [26] only retrains the classifier (*i.e.* linear layer), fine-tuning the whole model with a reduced learning rate produced superior results in our case. For both stages, we adopt Balanced Softmax loss [17], which effectively handles the discrepancy between the posterior distributions of training and test set. Given the number of samples in each class i (denoted as n_i), and the model's output (denoted as η), Balanced Softmax ($\hat{\phi}$) can be expressed as $\hat{\phi}_j = \frac{n_j e^{\eta_j}}{\sum_{i=1}^k n_i e^{\eta_i}}$. In both stages, only real samples are accounted for Balanced Softmax and no other re-sampling or re-weighting technique is adopted. The effectiveness of Balanced Softmax loss, which harmoniously blends with our method, is demonstrated through ablation experiments in Sec. 4.5.

3.3 Comparison with existing methods

Here, we clarify the differences between our method and previous works [48, 49, 74] on synthetic data generation. He *et al.* [48] and Sariyildiz *et al.* [49] mainly focus on the zero-shot learning and

transfer learning of the robust representation acquired from synthetic images. The key difference between these two is that Sariyildiz *et al.* emphasize training the network from scratch, whereas He *et al.* utilize synthetic data to fine-tune a pretrained CLIP model. Another notable work by Azizi *et al.* [74] focus on tuning text-conditioned Imagen [45] to a class conditional diffusion model and highlight their method using a metric called Classification Accuracy Score [114]. We gently note that comparing the performance of this model, trained on full data, with the aforementioned zero-shot or few-shot methods may result in misleading conclusions. Our work primarily focuses on long-tailed scenarios through per-class optimization of random initialized text tokens.

4 Experiments

4.1 Datasets

To validate our method, we conduct experiments on commonly used long-tailed datasets: ImageNet-LT [4], Places-LT [4], and iNaturalist2018 [3]. We utilize ImageNet-100 (IN100) and ImageNet-100-LT (IN100-LT) for our ablation studies. IN100 is a randomly selected subset of ImageNet-1K consisting of 115.8K training images, and IN100-LT is its long-tailed version with an exponentially decaying number of samples. Further details are summarized in Table 2. Our findings align with previous studies [115, 49], as we observe similar trends across ImageNet, ImageNet-LT, and their downscaled versions, IN100 and IN100-LT. It is important to note that direct comparison of IN100 results with previous papers [49, 51] could lead to misleading results, as each study adopts a distinct set of randomly selected classes.

Table 2: Overview of long-tailed datasets

Dataset	# Classes	# Images	Imbalance Factor
ImageNet-LT [4]	1000	115.8K	256
Places-LT [4]	365	62.5K	996
iNaturalist-LT [3]	8142	437K	500
IN100-LT	100	19K	200

Table 3: Classification accuracy (%) on ImageNet-LT and iNaturalist2018 using ResNet50 backbone. * denotes models trained with longer epochs (over 200). Our Fill-Up adopts longer training only for iNaturalist2018. We separately categorize ensembling-based methods involving multiple experts.

Method	ImageNet-LT [4]				iNaturalist2018 [3]			
	Overall	Many	Medium	Few	Overall	Many	Medium	Few
Baseline (CE) [24]	41.6	64.0	33.8	5.8	61.0	73.9	63.5	55.5
Decouple-cRT [26]	47.3	58.8	44.0	26.1	68.2	73.2	68.8	66.1
Decouple-LWS [26]	47.7	57.1	45.2	29.3	69.5	71.0	69.8	68.8
Remix [22]	48.6	60.4	46.9	30.7	70.5	-	-	-
LDAM-DRW [16]	49.8	60.4	46.9	30.7	66.1	-	-	-
BS [17]	51.0	60.9	48.8	32.1	70.0	70.0	70.2	69.9
PaCo* [93]	57.0	65.0	55.7	38.2	73.2	70.3	73.2	73.6
BS + CMO* [24]	58.0	67.0	55.0	44.2	74.0	71.9	74.2	74.2
RIDE (4 experts) [29]	55.4	66.2	52.3	36.5	72.6	70.9	72.4	73.1
NCL (3 experts)* [30]	59.5	-	-	-	74.9	72.7	75.6	74.5
BalPOE (3 experts)* [31]	60.8	67.8	59.2	46.5	76.9	75.0	77.4	76.9
Fill-Up (Stage I)	61.1	65.2	59.7	54.9	64.4	58.7	63.9	66.3
Fill-Up (Stage II)(*)	63.7	69.0	62.3	54.6	74.7	71.7	74.8	75.3

4.2 Implementation details

For generating synthetic samples, we adhere to our best setting from Sec. 3.1. To facilitate comparison, we utilize the widely adopted ResNet50 [112] for all experiments. For Places-LT, we additionally report scores initiated from pretrained ResNet152 [112], following the conventional approaches [7, 117]. With RandAugment [21], our models are trained for 100 epochs in Stage I and fine-tuned for 30 epochs in the Stage II. However, we make an exception for the fine-grained iNaturalist2018 dataset, extending the Stage II to 400 epochs, based on the notable performance increase observed with longer training schedules, as recommended by [93]. After evaluating the models on the corresponding balanced test/validation sets, we report overall top-1 accuracy as well as accuracies for many-shot (more than 100 images), medium-shot (20-100 images), and few-shot (less than 20 images) classes. Further details are available in Appendix.

Table 4: Classification accuracy (%) on Places-LT trained using pretrained ResNet152. † denotes models trained from scratch using ResNet50 backbone.

Method	Places-LT [4]			
	Overall	Many	Medium	Few
Baseline (CE)†	18.7	37.4	12.0	0.5
Baseline (CE) [93]	30.2	45.7	27.3	8.2
Decouple-LWS [26]	37.6	40.6	39.1	28.6
BS [17]	38.6	42.0	39.3	30.5
ResLT [116]	39.8	39.8	43.6	31.4
MiSLAS [28]	40.4	39.6	43.3	36.1
PaCo [93]	41.2	37.5	47.2	33.9
NCL (3 experts) [30]	41.8	-	-	-
SHIKE (3 experts) [32]	41.9	43.6	39.2	44.8
Fill-Up (Stage I)†	40.0	41.5	41.3	34.5
Fill-Up (Stage II)†	42.0	45.0	43.5	34.1
Fill-Up (Stage I)	41.5	42.5	43.2	36.2
Fill-Up (Stage II)	42.6	45.7	43.7	35.1

Table 5: Classification accuracy (%) on ImageNet-LT test set using different synthetic data generation methods.

Method	ImageNet-LT [4]	
	Guidance	Top-1
ImageNet-LT	-	41.6
Sariyildiz <i>et al.</i> [49]	2.0	42.9
Bansal <i>et al.</i> [51]	7.5	30.1
Single Template	7.5	26.6
Single Template + Real Data	1.0	41.9
	-	53.3
Textual Inversion (1.3K)	1.0	49.9
+ Real Data	-	55.6
+ BS + RandAug	-	61.8
Textual Inversion (2.6K)	1.0	53.5
+ Real Data	-	57.1
+ BS + RandAug	-	63.7
ImageNet (Full)	-	76.1

4.3 Efficiency of Textual Inversion generated images

Before delving into long-tailed benchmarks, we first evaluate the efficacy of the Textual Inversion-based generation strategy on the scale of ImageNet. Table 5 presents the results of Stage I training with synthetic data targeted for ImageNet-LT. Following standard ImageNet data, which contains 1,300 images for each class, we generate either 1,300 or 2,600 images, resulting in a total of 1.3M and 2.6M images, respectively. If real data is provided, we balance it by filling it up to contain 1,300 or 2,600 images per class accordingly. Aligning with our observations on IN100 and IN100-LT, generating images from textual-inverted tokens yields 8% higher top-1 accuracy than using single template generated data when trained solely with the synthetic data, and 2.3% higher top-1 accuracy when trained with filled-up real data. We also observe that these numbers scale as the amount of synthetic samples grow. Although it is important to mention that our method uses long-tailed real samples to learn text-tokens, it is noteworthy that this approach outperforms the previous state-of-the-art generation method from Sariyildiz *et al.* [49] which employs DINO style data augmentation, by 7% without employing any data augmentation, using the same number (1,300) of images.

4.4 Benchmarks on long-tailed datasets

We compare the proposed two-stage method based on Textual Inversion with the state-of-the-art methods. The benchmark results for ImageNet-LT and iNaturalist2018 can be found in Table 3, and results for Places-LT are displayed in Table 4. The benchmark results from ImageNet-LT demonstrate that our method effectively learns a balanced representation in Stage I, and achieves a notable performance improvement by learning decisive features from the real domain in Stage II. Especially, our method exhibits a significantly higher top-1 accuracy for few-shot classes, demonstrating more than 16% performance gain compared to PaCo [93], which previously achieved the highest accuracy for few-shot classes.

iNaturalist2018 dataset is not only known for its imbalanced data distribution, but also for its extremely fine-grained classes [7]. The relatively modest performance gain observed during our Stage I training can be attributed to the noise in the data generated using a low classifier-free guidance of scale 1.0. However, despite the potential hindrance caused by noisy samples in learning fine-grained classes, our Stage II model, trained on top of balanced feature space derived from Stage I, exhibits a substantial performance improvement while retaining balanced knowledge. We also note remarkable performance gain from 70.4% to 74.7% with PaCo-style longer training schedules.

Places-LT benchmark, as displayed in Table 4, is a very challenging dataset where training of ResNet50 from scratch only yields 18.7% overall accuracy and an exceptionally low 0.5% accuracy for few-shot classes. As a result, adopting an ImageNet pretrained ResNet152 is a common practice. However, our proposed method not only outperforms existing state-of-the-art approaches with the standard protocol of ResNet152 but also surpasses them when trained from scratch using ResNet50.

Figure 4: An illustration of different Fill-Up methods. Long-tailed data represents the actual class distribution of IN100-LT.




Table 6: Classification accuracy (%) on IN100-LT using ResNet50. Diverse Fill-Up approaches are denoted in accordance with Figure 4.

Method	IN100-LT			
	Overall	Many	Medium	Few
Baseline (Full Data)	82.3	81.3	84.1	81.8
Baseline (LT data)	26.4	50.8	13.5	1.9
+ BS	29.3	45.1	27.5	6.1
+ BS + RandAug	27.9	43.3	26.7	4.7
Fake Only	66.1	67.0	71.5	56.7
Stage I				
+ A (Under-balance)	56.0	70.7	53.6	35.0
+ B (Balance)	65.8	77.3	65.9	47.2
+ C (Over-balance)	69.5	80.1	70.2	52.0
+ BS	71.2	73.0	73.1	66.1
+ BS + RandAug	72.0	73.8	72.9	68.2
+ D (Addon)	70.1	81.1	70.9	51.7
Stage II				
Stage I (C+ BS+ RandAug)	69.5	80.1	70.2	52.0
+ Naïve	67.8	78.8	66.4	51.7
+ Class-balanced sampling	64.4	75.3	67.0	44.1
+ cRT	72.7	79.7	76.9	56.8
+ BS	74.0	79.1	75.9	63.8
+ RandAug	74.5	79.8	77.0	63.1

4.5 Ablation studies

Here, we perform extensive ablations using IN100 and IN100-LT datasets, focusing on the impact of different Fill-Up strategies and the effectiveness of our design choice in two-stage training. As depicted in Fig 4, we adopt four distinct strategies to fill the long-tailed data with synthetic images. In line with the previous observation from Fig 2b, performance generally tends to scale with an increasing number of images used for Fill-Up. However, experimental results suggest that naïvely adding all 1,300 images for every class (referred to as "D - addon") without considering the relative balance among classes leads to only marginal performance gains. It is also important to highlight the powerful synergy between our method and the Balanced Softmax loss. Compared to the baseline, the accuracy of the few-shot classes rises from 6.1% to an impressive accuracy of 66.1%.

For Stage II training, we begin from our best setting in Stage I. As suggested by previous works, naïve fine-tuning or class-balanced sampling degrades the balanced features learned from the first stage under extremely long-tailed data. cRT-style fine-tuning of the final classifier brings 3.2% performance gains on top of the balanced features. In our case, fine-tuning the whole model with Balanced Softmax loss resulted in a stronger performance. We believe, this unbiased extension of the softmax function, effectively mitigates the challenges posed by highly imbalanced data, allowing the entire set of model parameters to be efficiently tuned without collapsing. Given the additional 0.5% gain from RandAugment, our method achieves a remarkable accuracy of 74.5% which significantly reduces the gap with real data (82.3%) compared to the long-tailed baseline (26.4%).

5 Conclusion

This paper introduces the concept of filling-up long-tailed data with synthetic samples from recent large-scale text-to-image synthesis models. We provide valuable insights into different methods of synthetic data generation through detailed experiments and propose a novel approach leveraging Textual Inversion that outperforms other methods. Carefully designed with real-world data imbalance scenarios in mind, our approach, which comprises per-class optimization of text tokens and two-stage training, displays state-of-the-art results on standard long-tailed benchmarks, with particularly notable performance in few-shot classes accuracy.

Limitations and future work. While synthetic data generation demonstrates highly scaled performance with an increasing number of generated images, vast number of sampling through diffusion process demands massive computational resources. Recent advancements in one-step generative models (*e.g.* GAN [47, 63], and consistency models [118]) offer potential solutions. Furthermore, the ability to generate unlimited data comes at the cost of increased computation for training. While our approach shows higher efficiency in generating feature-rich samples, the challenge of generating images with strong features on par with real samples remains as another core problem for future.

References

- [1] Tsung-Yi Lin, Michael Maire, Serge Belongie, James Hays, Pietro Perona, Deva Ramanan, Piotr Dollár, and C Lawrence Zitnick. Microsoft coco: Common objects in context. In *European Conference on Computer Vision (ECCV)*, 2014.
- [2] Ziwei Liu, Ping Luo, Xiaogang Wang, and Xiaoou Tang. Deep learning face attributes in the wild. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2015.
- [3] Grant Van Horn, Oisin Mac Aodha, Yang Song, Yin Cui, Chen Sun, Alex Shepard, Hartwig Adam, Pietro Perona, and Serge Belongie. The inaturalist species classification and detection dataset. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [4] Ziwei Liu, Zhongqi Miao, Xiaohang Zhan, Jiayun Wang, Boqing Gong, and Stella X Yu. Large-scale long-tailed recognition in an open world. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [5] Justin M Johnson and Taghi M Khoshgoftaar. Survey on deep learning with class imbalance. *Journal of Big Data*, 2019.
- [6] Michael A Hedderich, Lukas Lange, Heike Adel, Jannik Strötgen, and Dietrich Klakow. A Survey on Recent Approaches for Natural Language Processing in Low-Resource Scenarios. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2021.
- [7] Lu Yang, He Jiang, Qing Song, and Jun Guo. A survey on long-tailed visual recognition. *International Journal of Computer Vision (IJCV)*, 2022.
- [8] Sophie Henning, William H Beluch, Alexander Fraser, and Annemarie Friedrich. A Survey of Methods for Addressing Class Imbalance in Deep-Learning Based Natural Language Processing. *arXiv preprint arXiv:2210.04675*, 2022.
- [9] Nitesh V Chawla, Kevin W Bowyer, Lawrence O Hall, and W Philip Kegelmeyer. SMOTE: synthetic minority over-sampling technique. *Journal of artificial intelligence research*, 2002.
- [10] Chris Drummond, Robert C Holte, et al. C4. 5, class imbalance, and cost sensitivity: why under-sampling beats over-sampling. In *Workshop on learning from imbalanced datasets II*, 2003.
- [11] Hui Han, Wen-Yuan Wang, and Bing-Huan Mao. Borderline-SMOTE: a new over-sampling method in imbalanced data sets learning. In *Advances in Intelligent Computing: International Conference on Intelligent Computing*, 2005.
- [12] Jason Van Hulse, Taghi M Khoshgoftaar, and Amri Napolitano. Experimental perspectives on learning from imbalanced data. In *International Conference on Machine Learning (ICML)*, 2007.
- [13] Chen Huang, Yining Li, Chen Change Loy, and Xiaoou Tang. Learning deep representation for imbalanced classification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [14] Salman H Khan, Munawar Hayat, Mohammed Bennamoun, Ferdous A Sohel, and Roberto Togneri. Cost-sensitive learning of deep feature representations from imbalanced data. *IEEE transactions on neural networks and learning systems*, 2017.
- [15] Yin Cui, Menglin Jia, Tsung-Yi Lin, Yang Song, and Serge Belongie. Class-balanced loss based on effective number of samples. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [16] Kaidi Cao, Colin Wei, Adrien Gaidon, Nikos Arechiga, and Tengyu Ma. Learning imbalanced datasets with label-distribution-aware margin loss. *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [17] Jiawei Ren, Cunjun Yu, Xiao Ma, Haiyu Zhao, Shuai Yi, et al. Balanced meta-softmax for long-tailed visual recognition. *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [18] Haibo He, Yang Bai, Eduardo A Garcia, and Shutao Li. ADASYN: Adaptive synthetic sampling approach for imbalanced learning. In *IEEE international joint conference on neural networks*, 2008.
- [19] Hongyi Zhang, Moustapha Cisse, Yann N. Dauphin, and David Lopez-Paz. mixup: Beyond Empirical Risk Minimization. In *International Conference on Learning Representations (ICLR)*, 2018.

- [20] Sangdoon Yun, Dongyoon Han, Seong Joon Oh, Sanghyuk Chun, Junsuk Choe, and Youngjoon Yoo. Cutmix: Regularization strategy to train strong classifiers with localizable features. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [21] Ekin D Cubuk, Barret Zoph, Jonathon Shlens, and Quoc V Le. Randaugment: Practical automated data augmentation with a reduced search space. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR) Workshop*, 2020.
- [22] Hsin-Ping Chou, Shih-Chieh Chang, Jia-Yu Pan, Wei Wei, and Da-Cheng Juan. Remix: rebalanced mixup. In *European Conference on Computer Vision (ECCV) Workshop*, 2020.
- [23] Jaehyung Kim, Jongheon Jeong, and Jinwoo Shin. M2m: Imbalanced classification via major-to-minor translation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [24] Seulki Park, Youngkyu Hong, Byeongho Heo, Sangdoon Yun, and Jin Young Choi. The Majority Can Help The Minority: Context-rich Minority Oversampling for Long-tailed Classification. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [25] Yuhang Zang, Chen Huang, and Chen Change Loy. Fasa: Feature augmentation and sampling adaptation for long-tailed instance segmentation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [26] Bingyi Kang, Saining Xie, Marcus Rohrbach, Zhicheng Yan, Albert Gordo, Jiashi Feng, and Yannis Kalantidis. Decoupling Representation and Classifier for Long-Tailed Recognition. In *International Conference on Learning Representations (ICLR)*, 2020.
- [27] Yongshun Zhang, Xiu-Shen Wei, Boyan Zhou, and Jianxin Wu. Bag of tricks for long-tailed visual recognition with deep convolutional neural networks. In *Proceedings of the AAAI conference on artificial intelligence*, 2021.
- [28] Zhisheng Zhong, Jiequan Cui, Shu Liu, and Jiaya Jia. Improving calibration for long-tailed recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [29] Xudong Wang, Long Lian, Zhongqi Miao, Ziwei Liu, and Stella Yu. Long-tailed Recognition by Routing Diverse Distribution-Aware Experts. In *International Conference on Learning Representations (ICLR)*, 2021.
- [30] Jun Li, Zichang Tan, Jun Wan, Zhen Lei, and Guodong Guo. Nested collaborative learning for long-tailed visual recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [31] Emanuel Sanchez Aimar, Arvi Jonnarth, Michael Felsberg, and Marco Kuhlmann. Balanced Product of Experts for Long-Tailed Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [32] Yan Jin, Mengke Li, Yang Lu, Yiu-ming Cheung, and Hanzi Wang. Long-Tailed Visual Recognition via Self-Heterogeneous Integration with Knowledge Excavation. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [33] Shengyu Zhao, Zhijian Liu, Ji Lin, Jun-Yan Zhu, and Song Han. Differentiable Augmentation for Data-Efficient GAN Training. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [34] Tero Karras, Miika Aittala, Janne Hellsten, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Training Generative Adversarial Networks with Limited Data. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [35] Mohamad Shahbazi, Martin Danelljan, Danda Pani Paudel, and Luc Van Gool. Collapse by Conditioning: Training Class-conditional GANs with Limited Data. In *International Conference on Learning Representations (ICLR)*, 2022.
- [36] Harsh Rangwani, Naman Jaswani, Tejan Karmali, Varun Jampani, and R Venkatesh Babu. Improving GANs for Long-Tailed Data Through Group Spectral Regularization. In *European Conference on Computer Vision (ECCV)*, 2022.
- [37] Harsh Rangwani, Lavish Bansal, Kartik Sharma, Tejan Karmali, Varun Jampani, and R Venkatesh Babu. NoisyTwins: Class-Consistent and Diverse Image Generation through StyleGANs. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [38] Diederik P Kingma and Max Welling. Auto-encoding variational bayes. *arXiv preprint arXiv:1312.6114*, 2013.

- [39] Ian Goodfellow, Jean Pouget-Abadie, Mehdi Mirza, Bing Xu, David Warde-Farley, Sherjil Ozair, Aaron Courville, and Yoshua Bengio. Generative Adversarial Nets. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2014.
- [40] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [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 *International Conference on Machine Learning (ICML)*, 2022.
- [42] Aditya Ramesh, Mikhail Pavlov, Gabriel Goh, Scott Gray, Chelsea Voss, Alec Radford, Mark Chen, and Ilya Sutskever. Zero-shot text-to-image generation. In *International Conference on Machine Learning (ICML)*, 2021.
- [43] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [44] Aditya Ramesh, Prafulla Dhariwal, Alex Nichol, Casey Chu, and Mark Chen. Hierarchical Text-Conditional Image Generation with CLIP Latents. *arXiv preprint arXiv:2204.06125*, 2022.
- [45] 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 preprint arXiv:2205.11487*, 2022.
- [46] 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 preprint arXiv:2211.01324*, 2022.
- [47] Minguk Kang, Jun-Yan Zhu, Richard Zhang, Jaesik Park, Eli Shechtman, Sylvain Paris, and Taesung Park. Scaling up GANs for Text-to-Image Synthesis. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [48] 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? In *International Conference on Learning Representations (ICLR)*, 2023.
- [49] Mert Bulent Sariyildiz, Karteek Alahari, Diane Larlus, and Yannis Kalantidis. Fake it till you make it: Learning transferable representations from synthetic ImageNet clones. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2023.
- [50] 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, 2023.
- [51] Hritik Bansal and Aditya Grover. Leaving Reality to Imagination: Robust Classification via Generated Datasets. In *International Conference on Learning Representations (ICLR) Workshop*, 2023.
- [52] 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. *arXiv preprint arXiv:2208.01618*, 2022.
- [53] Prafulla Dhariwal and Alex Nichol. Diffusion models beat gans on image synthesis. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2021.
- [54] Alex Nichol and Prafulla Dhariwal. Improved denoising diffusion probabilistic models. In *International Conference on Machine Learning (ICML)*, 2021.
- [55] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations (ICLR)*, 2021.
- [56] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al. Learning transferable visual models from natural language supervision. In *International Conference on Machine Learning (ICML)*, 2021.
- [57] Chao Jia, Yinfei Yang, Ye Xia, Yi-Ting Chen, Zarana Parekh, Hieu Pham, Quoc Le, Yun-Hsuan Sung, Zhen Li, and Tom Duerig. Scaling up visual and vision-language representation learning with noisy text supervision. In *International Conference on Machine Learning (ICML)*, 2021.

- [58] Andrew Brock, Jeff Donahue, and Karen Simonyan. Large Scale GAN Training for High Fidelity Natural Image Synthesis. In *International Conference on Learning Representations (ICLR)*, 2019.
- [59] Tero Karras, Samuli Laine, and Timo Aila. A style-based generator architecture for generative adversarial networks. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [60] Axel Sauer, Katja Schwarz, and Andreas Geiger. Stylegan-xl: Scaling stylegan to large diverse datasets. In *ACM SIGGRAPH*, 2022.
- [61] Minguk Kang, Joonghyuk Shin, and Jaesik Park. Studiogan: A taxonomy and benchmark of gans for image synthesis. *arXiv preprint arXiv:2206.09479*, 2022.
- [62] Jiahui Yu, Yuanzhong Xu, Jing Yu Koh, Thang Luong, Gunjan Baid, Zirui Wang, Vijay Vasudevan, Alexander Ku, Yinfei Yang, Burcu Karagol Ayan, et al. Scaling autoregressive models for content-rich text-to-image generation. *arXiv preprint arXiv:2206.10789*, 2022.
- [63] Axel Sauer, Tero Karras, Samuli Laine, Andreas Geiger, and Timo Aila. Stylegan-t: Unlocking the power of gans for fast large-scale text-to-image synthesis. *arXiv preprint arXiv:2301.09515*, 2023.
- [64] Nataniel Ruiz, Yuanzhen Li, Varun Jampani, Yael Pritch, Michael Rubinstein, and Kfir Aberman. Dreambooth: Fine tuning text-to-image diffusion models for subject-driven generation. *arXiv preprint arXiv:2208.12242*, 2022.
- [65] Nupur Kumari, Bingliang Zhang, Richard Zhang, Eli Shechtman, and Jun-Yan Zhu. Multi-Concept Customization of Text-to-Image Diffusion. *arXiv preprint arXiv:2212.04488*, 2022.
- [66] Christoph et. al. Schuhmann. LAION-5B: An open large-scale dataset for training next generation image-text models. *arXiv preprint arXiv:2210.08402*, 2022.
- [67] Runwayml. stable-diffusion. <https://github.com/runwayml/stable-diffusion>, 2022.
- [68] Stephan R Richter, Vibhav Vineet, Stefan Roth, and Vladlen Koltun. Playing for data: Ground truth from computer games. In *European Conference on Computer Vision (ECCV)*, 2016.
- [69] Xingchao Peng, Ben Usman, Neela Kaushik, Judy Hoffman, Dequan Wang, and Kate Saenko. Visda: The visual domain adaptation challenge. *arXiv preprint arXiv:1710.06924*, 2017.
- [70] Yuxuan Zhang, Huan Ling, Jun Gao, Kangxue Yin, Jean-Francois Lafleche, Adela Barriuso, Antonio Torralba, and Sanja Fidler. Datasetgan: Efficient labeled data factory with minimal human effort. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [71] Daiqing Li, Huan Ling, Seung Wook Kim, Karsten Kreis, Sanja Fidler, and Antonio Torralba. BigDatasetGAN: Synthesizing ImageNet with pixel-wise annotations. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [72] Lucy Chai, Jun-Yan Zhu, Eli Shechtman, Phillip Isola, and Richard Zhang. Ensembling with deep generative views. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [73] Ali Jahanian, Xavier Puig, Yonglong Tian, and Phillip Isola. Generative Models as a Data Source for Multiview Representation Learning. In *International Conference on Learning Representations (ICLR)*, 2022.
- [74] Shekoofeh Azizi, Simon Kornblith, Chitwan Saharia, Mohammad Norouzi, and David J Fleet. Synthetic data from diffusion models improves imagenet classification. *arXiv preprint arXiv:2304.08466*, 2023.
- [75] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Fei-Fei Li. ImageNet: A large-scale hierarchical image database. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2009.
- [76] Alexander C Li, Mihir Prabhudesai, Shivam Duggal, Ellis Brown, and Deepak Pathak. Your Diffusion Model is Secretly a Zero-Shot Classifier. *arXiv preprint arXiv:2303.16203*, 2023.
- [77] Weijia Wu, Yuzhong Zhao, Mike Zheng Shou, Hong Zhou, and Chunhua Shen. Diffumask: Synthesizing images with pixel-level annotations for semantic segmentation using diffusion models. *arXiv preprint arXiv:2303.11681*, 2023.
- [78] Shaobo Lin, Kun Wang, Xingyu Zeng, and Rui Zhao. Explore the Power of Synthetic Data on Few-shot Object Detection. *arXiv preprint arXiv:2303.13221*, 2023.

- [79] Zebin You, Yong Zhong, Fan Bao, Jiacheng Sun, Chongxuan Li, and Jun Zhu. Diffusion Models and Semi-Supervised Learners Benefit Mutually with Few Labels. *arXiv preprint arXiv:2302.10586*, 2023.
- [80] Mateusz Buda, Atsuto Maki, and Maciej A Mazurowski. A systematic study of the class imbalance problem in convolutional neural networks. *Neural networks*, 2018.
- [81] Inderjeet Mani and I Zhang. kNN approach to unbalanced data distributions: a case study involving information extraction. In *International Conference on Machine Learning (ICML) Workshop*, 2003.
- [82] Nikolaos Sarafianos, Xiang Xu, and Ioannis A Kakadiaris. Deep imbalanced attribute classification using visual attention aggregation. In *European Conference on Computer Vision (ECCV)*, 2018.
- [83] Yin Cui, Yang Song, Chen Sun, Andrew Howard, and Serge Belongie. Large scale fine-grained categorization and domain-specific transfer learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2018.
- [84] Mengye Ren, Wenyuan Zeng, Bin Yang, and Raquel Urtasun. Learning to reweight examples for robust deep learning. In *International Conference on Machine Learning (ICML)*, 2018.
- [85] Salman Khan, Munawar Hayat, Syed Waqas Zamir, Jianbing Shen, and Ling Shao. Striking the right balance with uncertainty. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2019.
- [86] Tsung-Yi Lin, Priya Goyal, Ross Girshick, Kaiming He, and Piotr Dollár. Focal loss for dense object detection. In *IEEE International Conference on Computer Vision (ICCV)*, 2017.
- [87] Jun Shu, Qi Xie, Lixuan Yi, Qian Zhao, Sanping Zhou, Zongben Xu, and Deyu Meng. Meta-weight-net: Learning an explicit mapping for sample weighting. *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [88] Sankha Subhra Mullick, Shounak Datta, and Swagatam Das. Generative adversarial minority oversampling. In *IEEE International Conference on Computer Vision (ICCV)*, 2019.
- [89] Jędrzej Kozerawski, Victor Fragoso, Nikolaos Karianakis, Gaurav Mittal, Matthew Turk, and Mei Chen. Blt: Balancing long-tailed datasets with adversarially-perturbed images. In *Proceedings of the Asian Conference on Computer Vision*, 2020.
- [90] Muhammad Abdullah Jamal, Matthew Brown, Ming-Hsuan Yang, Liqiang Wang, and Boqing Gong. Rethinking class-balanced methods for long-tailed visual recognition from a domain adaptation perspective. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2020.
- [91] Renzhen Wang, Kaiqin Hu, Yanwen Zhu, Jun Shu, Qian Zhao, and Deyu Meng. Meta feature modulator for long-tailed recognition. *arXiv preprint arXiv:2008.03428*, 2020.
- [92] Shuang Li, Kaixiong Gong, Chi Harold Liu, Yulin Wang, Feng Qiao, and Xinjing Cheng. Metasaug: Meta semantic augmentation for long-tailed visual recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [93] Jiequan Cui, Zhisheng Zhong, Shu Liu, Bei Yu, and Jiaya Jia. Parametric contrastive learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [94] Min-Kook Suh and Seung-Woo Seo. Long-Tailed Recognition by Mutual Information Maximization between Latent Features and Ground-Truth Labels. *arXiv preprint arXiv:2305.01160*, 2023.
- [95] Teli Ma, Shijie Geng, Mengmeng Wang, Jing Shao, Jiasen Lu, Hongsheng Li, Peng Gao, and Yu Qiao. A simple long-tailed recognition baseline via vision-language model. *arXiv preprint arXiv:2111.14745*, 2021.
- [96] Changyao Tian, Wenhai Wang, Xizhou Zhu, Jifeng Dai, and Yu Qiao. VI-ltr: Learning class-wise visual-linguistic representation for long-tailed visual recognition. In *European Conference on Computer Vision (ECCV)*, 2022.
- [97] Bowen Dong, Pan Zhou, Shuicheng Yan, and Wangmeng Zuo. LPT: Long-tailed Prompt Tuning for Image Classification. In *International Conference on Learning Representations (ICLR)*, 2023.
- [98] Jonathan Ho and Tim Salimans. Classifier-Free Diffusion Guidance. In *Conference on Neural Information Processing Systems (NeurIPS) Workshop*, 2021.
- [99] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising Diffusion Implicit Models. In *International Conference on Learning Representations (ICLR)*, 2020.

- [100] Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J Liu. Exploring the limits of transfer learning with a unified text-to-text transformer. *The Journal of Machine Learning Research*, 2020.
- [101] Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. Scaling instruction-finetuned language models. *arXiv preprint arXiv:2210.11416*, 2022.
- [102] LambdaLabsML. lambda-diffusers. <https://github.com/LambdaLabsML/lambda-diffusers>, 2022.
- [103] Stability-AI. stablediffusion. <https://github.com/Stability-AI/stablediffusion>, 2023.
- [104] Junnan Li, Dongxu Li, Silvio Savarese, and Steven Hoi. Blip-2: Bootstrapping language-image pre-training with frozen image encoders and large language models. *arXiv preprint arXiv:2301.12597*, 2023.
- [105] Jianhao Yuan, Francesco Pinto, Adam Davies, Aarushi Gupta, and Philip Torr. Not Just Pretty Pictures: Text-to-Image Generators Enable Interpretable Interventions for Robust Representations. *arXiv preprint arXiv:2212.11237*, 2022.
- [106] Bill Yuchen Lin, Wangchunshu Zhou, Ming Shen, Pei Zhou, Chandra Bhagavatula, Yejin Choi, and Xiang Ren. CommonGen: A Constrained Text Generation Challenge for Generative Commonsense Reasoning. In *Findings of the Association for Computational Linguistics: EMNLP*, 2020.
- [107] Manuel Romero. t5-base-finetuned-commongen. https://huggingface.co/mrm8488/t5-base-finetuned-common_gen, 2020.
- [108] Rohan Taori, Ishaan Gulrajani, Tianyi Zhang, Yann Dubois, Xuechen Li, Carlos Guestrin, Percy Liang, and Tatsunori B. Hashimoto. Stanford alpaca: An instruction-following llama model. https://github.com/tatsu-lab/stanford_alpaca, 2023.
- [109] DeCLaRe Lab. flan-alpaca-xxl. <https://huggingface.co/declare-lab/flan-alpaca-xxl>, 2023.
- [110] Partho Das. prompt-extend. <https://huggingface.co/daspartho/prompt-extend>, 2022.
- [111] Zijie J Wang, Evan Montoya, David Munechika, Haoyang Yang, Benjamin Hoover, and Duen Horng Chau. DiffusionDB: A Large-scale Prompt Gallery Dataset for Text-to-Image Generative Models. *arXiv preprint arXiv:2210.14896*, 2022.
- [112] Kaiming He, Xiangyu Zhang, Shaoqing Ren, and Jian Sun. Deep Residual Learning for Image Recognition. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [113] Rinon Gal. textual inversion issue 35. https://github.com/rinongal/textual_inversion/issues/35, 2022.
- [114] Suman Ravuri and Oriol Vinyals. Classification accuracy score for conditional generative models. *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.
- [115] Aniruddha Saha, Ajinkya Tejankar, Soroush Abbasi Koohpayegani, and Hamed Pirsiavash. Backdoor Attacks on Self-Supervised Learning. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2022.
- [116] Jiequan Cui, Shu Liu, Zhuotao Tian, Zhisheng Zhong, and Jiaya Jia. Reslt: Residual learning for long-tailed recognition. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2022.
- [117] Yifan Zhang, Bingyi Kang, Bryan Hooi, Shuicheng Yan, and Jiashi Feng. Deep long-tailed learning: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence (TPAMI)*, 2023.
- [118] Yang Song, Prafulla Dhariwal, Mark Chen, and Ilya Sutskever. Consistency models. *arXiv preprint arXiv:2303.01469*, 2023.
- [119] Jonathan Ho, Ajay Jain, and P. Abbeel. Denoising Diffusion Probabilistic Models. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2020.
- [120] Aaron Van Den Oord, Oriol Vinyals, et al. Neural discrete representation learning. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2017.

- [121] Patrick Esser, Robin Rombach, and Bjorn Ommer. Taming transformers for high-resolution image synthesis. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2021.
- [122] 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 *Conference on Neural Information Processing Systems (NeurIPS)*, 2017.
- [123] Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking the Inception Architecture for Computer Vision. In *IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*, 2016.
- [124] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. Improved Precision and Recall Metric for Assessing Generative Models. In *Conference on Neural Information Processing Systems (NeurIPS)*, 2019.

Appendices

A Diffusion models

A.1 Denoising diffusion probabilistic model

In this section, we provide a basic explanation of Denoising Diffusion Probabilistic Model (DDPM) [119] and Latent Diffusion Model (LDM) [43]. Diffusion model involves thermodynamics-inspired processes, namely the forward and reverse processes, which are associated with a Markov chain with Gaussian transition to gradually add and remove noises. Specifically, each step of the forward process (also called the diffusion process) is defined by a conditional Gaussian distribution where β_t controls mean and variance of output noises. It is demonstrated that $q(x_T)$ follows an isotropic standard Gaussian distribution under reasonable T and β_t constants. The reverse process, on the other hand, aims at recovering the original data distribution $q(x_0)$ from the isotropic standard Gaussian distribution $\mathcal{N}(0, \mathbf{I})$. Each reverse process is defined by a conditional Gaussian distribution $p_\theta(x_{t-1}|x_t)$, where the mean and variance are determined by a neural network parameterized by θ and time-dependent constants σ_t , respectively. Mathematically, each step of the forward and the reverse processes can be expressed as follows:

$$q(x_t|x_{t-1}) := \mathcal{N}(x_t; \sqrt{1 - \beta_t}x_{t-1}, \beta_t\mathbf{I}), \quad (\text{A1})$$

$$p_\theta(x_{t-1}|x_t) := \mathcal{N}(x_{t-1}; \mu_\theta(x_t, t), \sigma_t^2). \quad (\text{A2})$$

Since the reverse process (Eq. (A2)) is not analytically derived from the forward process (Eq. (A1)), current diffusion models approximate the true conditional distribution for reverse step $q(x_{t-1}|x_t)$ by plugging in $p_\theta(x_{t-1}|x_t)$ instead of $q(x_{t-1}|x_t)$ and minimizing negative log-likelihood for the true data distribution. Since calculating the exact log-likelihood is intractable, Ho *et al.* [119] optimize a variational lower bound \mathcal{L}_{VLB} on the negative log-likelihood. The variational lower bound is expressed in Eq. (A3).

$$\begin{aligned} \mathcal{L}_{\text{VLB}} &:= \mathbb{E}_q \left[-\log \frac{p_\theta(x_{0:T})}{q(x_{1:T}|x_0)} \right], \quad (\text{A3}) \\ &= \mathbb{E}_q \left[\underbrace{\text{KL}(q(x_T|x_0) \| p(x_T))}_{\mathcal{L}_T} + \sum_{t>1} \underbrace{\text{KL}(q(x_{t-1}|x_t, x_0) \| p_\theta(x_{t-1}|x_t))}_{\mathcal{L}_{t-1}} - \log p_\theta(x_0|x_1) \right]_{\mathcal{L}_0}. \end{aligned}$$

The above \mathcal{L}_{VLB} consists of three sub losses: 1) loss for the forward process \mathcal{L}_T , 2) loss for the backward process \mathcal{L}_{t-1} , and 3) discrete decoder loss \mathcal{L}_0 . Although the above variational lower bound can be computed analytically (KL divergence between two Gaussian distributions has an exact solution), the authors of DDPM utilize the practical loss function below and demonstrate the successful generation of realistic images from the real-world domain.

$$\mathcal{L}_{\text{simple}}(\theta) := \mathbb{E}_{x_0 \sim q(x), \epsilon \sim \mathcal{N}(0, \mathbf{I}), t} \left[\left\| \epsilon - \epsilon_\theta(\sqrt{\bar{\alpha}_t}x_0 + \sqrt{1 - \bar{\alpha}_t}\epsilon, t) \right\|_2^2 \right], \quad (\text{A4})$$

where $\bar{\alpha}_t = \prod_{s=1}^t \alpha_s$.

Finally, image sampling is conducted by recursively iterating the following equation:

$$x_{t-1} = \frac{1}{\sqrt{\alpha_t}} \left(x_t - \frac{1 - \alpha_t}{\sqrt{1 - \bar{\alpha}_t}} \epsilon_\theta(x_t, t) \right) + \sigma_t z, \quad \text{where } z \sim \mathcal{N}(0, 1). \quad (\text{A5})$$

A.2 Latent diffusion model

Latent Diffusion Model (LDM) reduces the excessive training and inference time of DDPM by performing forward and reverse processes in the latent space. Specifically, an encoder \mathcal{E} learns to map an arbitrary image $x \sim q(x)$ into a latent code $z = \mathcal{E}(x)$, and a decoder \mathcal{D} learns to map such latent back into the image space ($\mathcal{D}(\mathcal{E}(x)) \approx x$). Regularization by either KL-divergence loss [120] or VQGAN-like [121] vector quantization is adopted to avoid high-variance latent spaces. Given time step t , diffused latent variable z_t , unscaled noise sample ϵ , denoising network ϵ_θ , conditioning network c_θ , and conditional input y , the LDM loss is given as follows:

$$\mathcal{L}_{\text{LDM}} := \mathbb{E}_{z \sim \mathcal{E}(x), y, \epsilon \sim \mathcal{N}(0, 1), t} \left[\left\| \epsilon - \epsilon_\theta(z_t, t, c_\theta(y)) \right\|_2^2 \right] \quad (\text{A6})$$

Stable Diffusion is a specific text-conditioned variant of the LDM trained on LAION-5B [66] dataset which utilizes CLIP text encoder [56] as a conditional network c_θ for text input y . In LDM variants, classifier-free guidance [98] simultaneously guides the reverse process by calculating the difference between the text- and null-conditioned score functions. The mathematical formulation of classifier-free guidance can be written as follows:

$$\tilde{\epsilon}_t = \epsilon_\theta(z_t) + w(\epsilon_\theta(z_t, c) - \epsilon_\theta(z_t)), \quad (\text{A7})$$

where $\tilde{\epsilon}_t$ is a guided output, and w is a guidance scale.

B Additional experimental details

B.1 Details on prompt-to-image based methods

For prompt-to-image-based generation, we employ five different methods: *single template*, *CLIP templates*, *T5*, *Flan-T5 XXL*, and *Prompt-Extend*. In the previous work by He *et al.* [48], they apply CLIP filtering to refine the outputs generated by the T5 language model, aiming to remove noisy prompts. However, we believe it is worth considering the noisy prompts as they are an inherent part of the language model and have potential value in evaluating the model’s performance. Hence, we refrain from applying additional filters to the generated prompts.

Table A7: Effect of different decoding strategies on Top-1 and Top-5 accuracy.

Decoding Method	T5	Flan-T5 XXL	Prompt-Extend
Beam Search	34.2 / 61.7	37.0 / 65.7	27.1 / 50.9
Top-p Sampling	36.0 / 61.0	41.3 / 67.3	35.0 / 58.8

Among a wide range of T5 [100] variants fine-tuned on different datasets, we specifically utilize T5 model fine-tuned on CommonGen dataset [106, 107]. Consistent with the findings from He *et al.* [48], we observe prompts generated from this variant tend to generate natural descriptions of plausible real-world objects. We supply the model with the class name and let it generate a sentence up to 60 tokens long.

For a larger model, we adopt Flan-T5 XXL [101], which is an instruction-tuned variant of T5 XXL. Specifically, we utilize the model fine-tuned through the recently popular Stanford Alpaca [108, 109]. Our specific instruction for the model is as follows: "Generate a prompt-style sentence for image generation including {CLS}". In general, this approach led to the generation of significantly more plausible outputs with reduced noise.

We observe that not all generated prompts can be effectively described by the Stable Diffusion due to limitations in its expressiveness. Motivated by this observation, we employ a widely used prompt-extender [110] for Stable Diffusion, which is a GPT-2 model fine-tuned on the DiffusionDB [111] dataset. Like T5, we supply the model with the class name and let the model extend prompts in a more favorable way for Stable Diffusion.

For language models, improved decoding methods also play a crucial role in enhancing the quality of the outputs. To evaluate their performances, we conduct experiments using two widely used decoding techniques: beam search and top-p nucleus sampling. To obtain n prompts, we use $4 \times n$ beams for beam search and use $k = 50$ and $p = 0.95$ for sampling. We generate 200 and 1,000 prompts for beam search and top-p nucleus sampling, respectively. Prompts are evenly utilized to generate 1,300 images per class. Following previous experiments, we generate 130K images for IN100 with a scale of 7.5. As seen in Table A7, we generally observe better performance with top-p sampling. The final table for prompt-to-image-based methods, covering scales of 1.0 and 7.5, is given in Table A8.

Table A8: Full comparison of accuracies among different prompt-to-image-based generation methods for IN100. Low Guidance indicates a scale of 1.0, and high guidance denotes 7.5.

Method (Top-1/Top-5)	Prompt-to-Image				
	Single Template	CLIP Templates	T5	Flan-T5 XXL	Prompt-Extend
Low Guidance	59.9 / 83.7	57.2 / 82.2	45.8 / 71.5	51.6 / 77.3	29.7 / 54.7
High Guidance	41.5 / 67.5	48.7 / 73.6	36.0 / 61.0	41.3 / 67.3	35.0 / 58.8

For higher guidance, which better reflects the complexity of given prompts, incorporating diversity in prompts seems to improve accuracy. However, additional complexity from fancy prompts tends

to deteriorate the recognition ability in lower guidance situations with more noise engaged. This motivates our intuition of using Textual Inversion and adopting a simple prompt ("a photo of *") with a more context-rich token ("*"). Furthermore, using specific domain-favored prompts, such as Prompt-Extend favored by Stable Diffusion, can lead to greater deviation from the original distribution. Qualitative results for all generated prompts and images can be found in Sec. D.1 and Sec. D.2.

B.2 Details on image-to-image based methods

We directly utilize models from [102, 103] for image-to-image-based methods. Since these two have different training setups, they have distinct default guidance scales, which are 3.0 for *Image Variation* and 10.0 for *Reimagine*. Reimagine also has a parameter called noise level, which controls the amount of noise added to the image embedding. Nonetheless, the impact of this parameter is minimal in top-1 accuracy (41.6% with noise level 0 and 41.3% with noise level 100 in a scale of 10.0), so we keep the noise level to 0. Moreover, considering that the Stable Diffusion 2 Reimagine model produces images with a resolution of 768 pixels, we note a potential advantage of this variant over all other approaches that generate images with a resolution of 512 pixels. When generating synthetic images, we try to maintain a balanced utilization of the given source images. For example, to generate 1,300 images using 5 real samples in the tail class, we utilize each source image 260 times for image-to-image generation. As discussed in the main paper, image-to-image-based methods sometimes tend to replace a crucial class-related component of an image with a "reimagined" object. Qualitative results of image-to-image based methods, including mentioned circumstances, can be found in Sec. D.2.

B.3 Details on transmodal methods

For transmodal methods, we introduce two novel approaches: *Captioning and Textual Inversion*. For captioning model, we utilize 6.7B parameter variant of BLIP2 [104]. We extract a single prompt using a default greedy decoding method for every single image in the long-tailed data. Unlike previous prompt-to-image-based approaches, we observe prompts diverging with more advanced decoding methods, failing to describe the images accurately. We evenly distribute prompts and use them to generate 1,300 images per class.

In the case of the *Textual Inversion* method, we optimize a single token "*" of embedding shape 4×768 for every class unless otherwise specified. To avoid manual handcrafted optimization, we adopt a unified learning rate of 5.0×10^{-3} , batch size of 2, and a simple heuristic for determining the total training steps, given by $(\min(\max(\text{number of images} \times 100, 2000), 10000))$. We save checkpoints every 100 steps and evenly utilize all of them for generating images. From experimental results, we believe that incorporating coarse-to-fine details during the optimization process contributes to learning robust class representations. In terms of storage, even for the head classes, the file sizes strictly remain below 1.6MB per class due to our limitation on the maximum training steps.

B.4 Details on training classifiers

For all Stage I training, we follow the standard ResNet50 training protocol for ImageNet. We randomly crop the images to a size of 224×224 and perform random horizontal flipping and color jittering. We train for 100 epochs with a batch size of 256 and an initial learning rate of 0.1, which is decayed by a factor of 0.1 every 30 epochs. If noted with RandAug or BS, standard RandAugment [21] or Balanced Softmax loss [17] was adopted following Park *et al.* [24]. For Stage II training, we continue training for 30 epochs from the last checkpoint in Stage I with only real data. During the second stage, the initial learning rate of 1×10^{-3} decayed by a factor of 0.1 every ten epochs is adopted with five warm-up epochs. However, as noted in the main paper, we adopt the PaCo [93] style (400 epochs and initial learning rate of 0.01 decayed by cosine schedule) Stage II training for iNaturalist2018 as it outperforms our basic 30 epoch training by a huge margin. Lastly, to accelerate training for iNaturalist2018, which becomes a huge dataset containing over 8M images with our balanced Fill-Up strategy, we specifically utilize a batch size of 2,048 for iNaturalist2018.

C Extended analysis

C.1 Alignment of generated data with real data

Here, we discuss the alignment between generated data and real data. One of the most famous metrics to quantitatively measure the quality of generated images in relation to the true distribution is the Fréchet Inception Distance(FID) [122]. FID calculates the distance between two distributions in the feature space provided by a pre-trained Inception-V3 [123] network. We also utilize Precision & Recall [124], which are widely adopted in the field of generative modeling, to quantify the fidelity and the diversity of the generated images separately. Precision is defined as the portion of generated images that falls within the support of the real data distribution, while recall is defined as the portion of real images falling within the support of the fake distribution [61]. We adopt Inception-V3 network as a feature extractor for Precision & Recall. In line with previous experiments, we generate a total of 130K images for IN100, with 1,300 images per class. Moreover, we maintain the practice of utilizing IN100-LT exclusively for learning text tokens and do not mix real samples during the evaluation.

As indicated in Table A9, we generally observe better FID and Precision & Recall values with Textual Inversion-generated images. A qualitative comparison between these two methods is depicted in Fig. A5. In general, textual-inverted tokens exhibit key features from the real domain, enabling the generation of images with some diversity similar to real samples. Moreover, we confirm that classifier-free guidance plays a crucial role in maintaining a balance between Precision & Recall. Although it is still very difficult to directly correlate these metrics with Top-1 accuracy, we suspect that achieving a high Recall value is important for improving accuracy. More qualitative analysis from Sec. D.2 supports the notion that noisier samples, generated with a low guidance scale, sparsely contain feature-rich instances that aid in the gradual expansion of the decision boundary.




Figure A5: Randomly selected images from IN100 classes generated via Single Template and Textual Inversion methods. For better visualization, images were generated using a guidance scale of 7.5. Class names are as follows: {A: soccer ball, B: traffic light, C: frying pan, D: monastery, E: airship} In general, Textual Inversion produces more diverse images.

Table A9: Comparison of generated images via single template method and Textual Inversion method on IN100. Real samples from IN100-LT are used to optimize text tokens. FID and Precision & Recall are calculated against 130K real images from IN100.

Scale	Single Template				Textual Inversion			
	Top-1 (%) \uparrow	FID \downarrow	Precision \uparrow	Recall \uparrow	Top-1 (%) \uparrow	FID \downarrow	Precision \uparrow	Recall \uparrow
0.0	1.3	130.2	0.292	0.632	1.3	130.2	0.292	0.632
0.5	37.0	130.3	0.206	0.751	40.5	133.3	0.210	0.751
1.0	59.9	75.2	0.329	0.670	66.1	77.5	0.338	0.691
1.1	57.8	65.9	0.362	0.658	66.4	68.2	0.378	0.670
1.5	61.4	42.4	0.492	0.577	65.9	41.7	0.503	0.607
2.0	58.0	29.3	0.598	0.498	66.6	26.0	0.600	0.536
3.0	55.1	21.6	0.703	0.386	63.1	15.0	0.695	0.450
5.0	45.4	21.1	0.764	0.254	54.9	12.3	0.743	0.350
7.5	41.5	22.8	0.771	0.180	47.2	13.4	0.744	0.290
9.0	37.5	23.9	0.765	0.160	47.6	14.0	0.737	0.266




Figure A6: Comparison of per-class accuracies using different methods. Colored lines indicate polynomial regression graphs of order 3, with bands indicating the 95% confidence interval. For ImageNet-LT, colored dots represent the actual values. We do not include dots for Places-LT and iNaturalist2018 as their evaluation set consists of only 20 and 3 images per class, respectively, resulting in a flat line of no significant variation. Fill-Up generally presents well-balanced per class accuracies, exhibiting notable improvements over PaCo.

C.2 Per class accuracies of Fill-Up

Sometimes, the standard protocol of reporting many-shot (more than 100 images), medium-shot (20-100 images), and few-shot (less than 20 images) accuracies may not be sufficient to fully represent the effectiveness of long-tailed methods, especially when some extreme tail classes exhibit significantly low accuracies. In Fig. A6, we present detailed per-class accuracies of our method in three datasets. We compare our method with the widely recognized PaCo method on ImageNet-LT dataset, as only ImageNet-LT pre-trained weights are available for comparison.

C.3 Impact of initial word choice for Textual Inversion

One major advantage of using Textual Inversion is that prior knowledge related to the classes is not required. In Fig. A7, we illustrate the optimization process of Textual Inversion with different

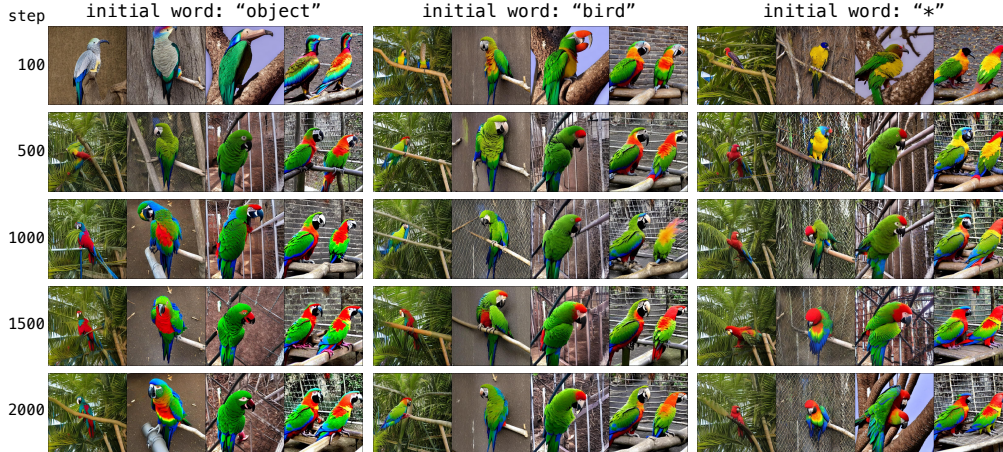


Figure A7: Optimization process of Textual Inversion with varying initialization of the token "*". We demonstrate the Textual Inversion training on six images from the macaw class of IN100-LT, displaying the intermediate images at each step of the optimization. We note that the optimization process of Textual Inversion is robust to the choice of initial words.

initialization words. While it is true that choosing an initialization word closer to the concept of the actual class can accelerate convergence, we also observe reasonably fast and accurate convergence with more general initialization words such as "object" and "*". Hence, we initialize the token "*" with the CLIP embedding of "object" and continue optimization. Our method can effectively capture the image distribution without any prior knowledge from the text domain by solely relying on the image domain during the optimization process.

C.4 Impact of varying diffusion parameters

While our main focus is on discussing more complex generation strategies, it is important to note that changing basic parameters of the diffusion process, such as scale, can have a dramatic impact on the results. Our findings mostly align with the previous research by Sariyildiz *et al.* [49]. As suggested in Fig. A8, the number of reverse process steps does not introduce significant performance gains when exceeding the default step of 50. Thus, regardless of the generation strategy, we keep our generation rule to 50 diffusion steps. We further explored adopting random steps between 0 and 100, but this didn't bring significant changes.

As discussed earlier in diffusion communities, we observe that it is very important to fix at least one dimension of generated images' resolution to 512. Empirically, we notice that excessive scalings or scalings that do not maintain at least one dimension at 512 results in a notable degradation in image quality. Therefore, in our work, we adhere to the default image resolution of Stable Diffusion, which is 512×512 .

C.5 Impact of the different prefix for single template

Throughout our paper, we employ a single template (*i.e.* "a photo of a {CLS}") as a strong baseline since it surprisingly outperforms most other methods, especially in lower guidance scales. Previous research by Sariyildiz *et al.* [49] suggests that the use of a simpler prompt, specifically "{CLS}", resulted in higher performance. However, due to the variations in experimental setups, particularly regarding the guidance scale and the version of Stable Diffusion used, we actually observed better results when using the prompt "a photo of a {CLS}". This prompt achieved a score of 59.9%, whereas "{CLS}" achieved a score of 56.9% in the scale 1.0 setting.

C.6 Performance difference between Stable Diffusion v1 and v2


As noted in the main paper, our experiments primarily utilize Stable Diffusion v1.5. Using a single template prompting strategy, we provide a comparison between v1 and v2 on different guidance

scales in Table A10. Following previous experiments, IN100 with 130K generated images were used. We adopt Stable Diffusion v2.1 base model in 512×512 resolution for a fair comparison.

Table A10: Comparison of top-1 accuracies (%) between Stable Diffusion v1 and v2 on IN100 with single template strategy ("a photo of a {CLS}"). It is important to note that scale 7.5 and 9.0 are the basic guidance scales of Stable Diffusion v1 and v2, respectively.

Scale	IN100 Top-1 Accuracy (%)	
	Stable Diffusion v1	Stable Diffusion v2
1.0	59.9	59.2
7.5	41.5	33.8
9.0	37.5	30.9

Figure A8: Impact of diffusion steps on top-1 accuracy.



D Qualitative results

D.1 Generated prompts

We first present generated prompts using the methods mentioned above. Specifically, prompts for CLIP templates (Fig. A9), T5 (Fig. A10), Flan-T5 XXL (Fig. A12), Prompt-Extend (Fig. A13), and the caption generated by BLIP2 (Fig. A11) are provided. For language models, we randomly select prompts for green iguana and African bush elephant classes.

D.2 Generated images

We present images generated by prompt-to-image based methods in scale 7.5 (Fig. A14) and scale 1.0 (Fig. A15), and image-to-image and transmodal methods in scale 7.5 (Fig. A16) and scale 1.0 (Fig. A17). We also provide samples for iNaturalist2018 in Fig. A18, and Places-LT in Fig. A19. Overall, images with higher guidance scale of 7.5 show better visual results when compared to those with lower guidance scale of 1.0. Furthermore, both the image-to-image and transmodal methods demonstrate improved alignment with real data. Especially, when utilizing Textual Inversion with a higher guidance scale, we observe the generation of extremely realistic images across various datasets. While lower guidance scale produce relatively noisier outputs, the notably higher top-1 accuracies obtained from lower guidance suggest the sparse existence of important samples. Although we rarely encounter samples where added noise actually results in plausible variations, we think that scaling these noisy samples, which yields more important samples over time, is the key reason for its high performance.

a bad photo of a {CLS}.	a cropped photo of a {CLS}.
a photo of many {CLS}.	a plastic {CLS}.
a sculpture of a {CLS}.	a photo of the dirty {CLS}.
a photo of the hard to see {CLS}.	a jpeg corrupted photo of a {CLS}.
a low resolution photo of the {CLS}.	a blurry photo of the {CLS}.
a rendering of a {CLS}.	a photo of the {CLS}.
graffiti of a {CLS}.	a good photo of the {CLS}.
a bad photo of the {CLS}.	a rendering of the {CLS}.
a cropped photo of the {CLS}.	a {CLS} in a video game.
a tattoo of a {CLS}.	a photo of one {CLS}.
the embroidered {CLS}.	a doodle of a {CLS}.
a photo of a hard to see {CLS}.	a close-up photo of the {CLS}.
a bright photo of a {CLS}.	a photo of a {CLS}.
a photo of a clean {CLS}.	the origami {CLS}.
a photo of a dirty {CLS}.	the {CLS} in a video game.
a dark photo of the {CLS}.	a sketch of a {CLS}.
a drawing of a {CLS}.	a doodle of the {CLS}.
a photo of my {CLS}.	a origami {CLS}.
the plastic {CLS}.	a low resolution photo of a {CLS}.
a photo of the cool {CLS}.	the toy {CLS}.
a close-up photo of a {CLS}.	a rendition of the {CLS}.
a black and white photo of the {CLS}.	a photo of the clean {CLS}.
a painting of the {CLS}.	a photo of a large {CLS}.
a painting of a {CLS}.	a rendition of a {CLS}.
a pixelated photo of the {CLS}.	a photo of a nice {CLS}.
a sculpture of the {CLS}.	a photo of a weird {CLS}.
a bright photo of the {CLS}.	a blurry photo of a {CLS}.
a cartoon {CLS}.	art of the {CLS}.
art of a {CLS}.	a drawing of the {CLS}.
a sketch of the {CLS}.	a photo of the large {CLS}.
a embroidered {CLS}.	a black and white photo of a {CLS}.
a pixelated photo of a {CLS}.	the plushie {CLS}.
itap of the {CLS}.	a dark photo of a {CLS}.
a jpeg corrupted photo of the {CLS}.	itap of a {CLS}.
a good photo of a {CLS}.	graffiti of the {CLS}.
a plushie {CLS}.	a toy {CLS}.
a photo of the nice {CLS}.	itap of my {CLS}.
a photo of the small {CLS}.	a photo of a cool {CLS}.
a photo of the weird {CLS}.	a photo of a small {CLS}.
the cartoon {CLS}.	a tattoo of the {CLS}.

Figure A9: A list of 80 templates used for CLIP templates strategy.

green iguana in a tropical rainforest	African elephants are grazing in the bush.
iguana in green on the beach	elephants in the bush in africa.
green iguana in the forest	elephants in a bush in africa.
iguana in green on the beach	an elephant in the bush in africa.
green iguana in the field	A group of elephants in a bush in Africa.
iguanas and a red and green iguana.	an elephant is grazing in a bush in africa.
iguana reversing the greens with a roar	elephants in a bush in African.
iguana in green on a sunny day	elephants and bush in the africa.
iguana in green with a shady twig	A group of elephants in the bush in Africa.
green iguana in a pond	African elephants are in a bush.
a lone iguana looks like it is a green.	African elephants roaming in a bush.
green iguana in a lagoon	An elephant in the bush in Africa.
iguana in green	An elephant in a bush in Africa.
green iguana in the wild.	elephants in the bush in africa.
iguanas in the green of the bushes	elephants and buffaloes in the bush in africa.
a lone iguana in green	An elephant is in the bush in Africa.
iguana in green on the green	An elephant in a bush in Africa.
iguana in green on a sunny day	An elephant in a bush in the country.

Figure A10: Randomly selected prompts sampled by T5 language model. We display prompts for green iguana and African bush elephant classes, decoded using top-p nucleus sampling strategy.

two iguanas are standing next to each other	a group of elephants walking around a pond
a lizard is sitting on a rock	a large elephant eating hay from a large pile
three iguanas are laying on the grass	elephants are the largest land mammals
a green iguana is sitting on a branch	a baby elephant is being held by a soldier
an iguana is sitting on top of a tree branch	an elephant walking on a dirt road
a large iguana laying on a log in a zoo	an elephant standing in the grass
a close up of an iguana with its head turned	an elephant is standing in tall grass
the lizard is brown	two elephants walking in the grass
several iguanas are sitting on the ground	an elephant standing in a field
a green iguana is standing on some steps	an elephant with tusks
a woman holding a lizard on a beach	a herd of elephants standing in the water
a lizard is sitting on top of a rock	an elephant is walking through the brush
an iguana is sitting on top of a tree stump	a woman is petting an elephant
a lizard is walking on the side of a fence	a group of people riding on an elephant
an iguana is sitting in the branches of a tree	a baby elephant walking
a lizard sitting on top of a log in the woods	two elephants walking in the grass
an iguana on a rock with its long tail	an elephant with a big nose
a lizard sitting on a rock in the sun	two elephants standing in water

Figure A11: Randomly selected prompts captioned by BLIP2 from IN100-LT images. We display prompts for green iguana and African bush elephant classes.

A green iguana rests on a tree branch in the forest.
 A green iguana curled up in the sun, its fur gleaming in the sunlight.
 A green iguana scurries through the lush forest in search of food.
 A green iguana is pacing across a lush tropical forest, its green fur gleaming in the sun.
 Imagine a green iguana sitting in a tree, sunbathing and enjoying the warm summer air.
 Imagine a green iguana standing in the middle of the tropical jungle.
 A green iguana scurries through the jungle, looking for a new home.
 The green iguana stood on a small island, its fur green and its eyes glowing in the sunlight.
 A green iguana is swimming in the ocean.

An African bush elephant grazing in the grasslands with its trunk up in the air.
 A herd of African bush elephants slowly strolling through the vast plains.
 The African bush elephant trudged through the tall grass, its trunk in the air.
 A majestic African bush elephant strolls through a lush, green forest.
 The African bush elephants walked peacefully across the lush green grasslands.
 A majestic African bush elephant gracefully walking through the dense foliage.
 Imagine a majestic African bush elephant grazing in the lush green jungle.
 An African bush elephant saunters through the forest, its large ears floppy in the heat.
 Imagine an African bush elephant galloping through the forest.

Figure A12: Randomly selected prompts sampled by Flan-T5 XXL language model. We display prompts for green iguana and African bush elephant classes, decoded using top-p nucleus sampling strategy.

green iguana with a red mohawk, green mohawk, wearing sunglasses, synthwave style,
 green iguana with a white nose and a green frog face, photograph, national geographic,
 green iguana, beautiful landscape, environment, colors, film, dramatic, cinematic,
 green iguana, intricate, elegant, highly detailed, digital painting, artstation,
 green iguana, wildlife photography on a road, high definition, shot with sigma f/ 4.2,
 green iguana with a flower face on its head, natural light, sharp, detailed face,
 green iguana in jungle, Beksinski painting, part by Adrian Ghenie and Gerhard Richter.
 green iguana with a gun in its mouth, realistic, digital art, painted by seb mckinnon,
 green iguana hybrid in the style of pixar, cinematic composition, cinematic lighting,

African bush elephant portrait, face, wearing a crown, futuristic,
 African bush elephant as a fantasy D&D character, portrait art by Donato Giancola and,
 African bush elephant with long hairs in fur and fur and fluffy fur with pastel colors,
 African bush elephant robot, 3d model, photorealistic, 3d, rendered, hd, 4k, HD, sharp,
 African bush elephant hybrid with a peacock tail and rainbow mane, rainbow mane, pastel,
 African bush elephant with a big trunk and a trunk. A frog is sitting on top of a leaf.
 African bush elephant chimera, 4k, high detail, high-resolution photograph, professional,
 African bush elephant, award winning nature photography, National Geographic,
 African bush elephant chimera, 4k, DSLR, highly detailed, professional photography,

Figure A13: Randomly selected prompts sampled by Prompt-Extend language model. We display prompts for green iguana and African bush elephant classes, decoded using top-p nucleus sampling strategy. We truncate lengthy sentences that consisted of a maximum of 60 tokens in order to fit the figure.




Figure A14: Randomly selected images from IN100 using different prompt-to-image based methods (scale=7.5). Class names are as follows: {A: green iguana, B: sunglasses, C: gazelle, D: paintbrush, E: mongoose, F: lifeboat, G: beaker, H: African bush elephant, I: tractor, J: bee, K: Border Terrier, L: bulbul, M: firetruck, N: guinea pig, O: Irish Setter, P: Sealyham Terrier, Q: bald eagle, R: toucan, S: leopard, T: ptarmigan}




Figure A15: Randomly selected images from IN100 using different prompt-to-image based methods (scale=1.0). Class names are as follows: {A: green iguana, B: sunglasses, C: gazelle, D: paintbrush, E: mongoose, F: lifeboat, G: beaker, H: African bush elephant, I: tractor, J: bee, K: Border Terrier, L: bulbul, M: firetruck, N: guinea pig, O: Irish Setter, P: Sealyham Terrier, Q: bald eagle, R: toucan, S: leopard, T: ptarmigan}




Figure A16: Randomly selected images from IN100 using image-to-image and transmodal based methods (scale=7.5). IN100-LT samples were used as source images. Class names and the number of real samples are as follows: {A: (green iguana, 1000), B: (sunglasses, 765), C: (gazelle, 651), D: (paintbrush, 472), E: (mongoose, 342), F: (lifeboat, 223), G: (beaker, 180), H: (African bush elephant, 153), I: (tractor, 117), J: (bee, 80), K: (Border Terrier, 61), L: (bulbul, 55), M: (firetruck, 40), N: (guinea pig, 32), O: (Irish Setter, 26), P: (Sealyham Terrier, 21), Q: (bald eagle, 17), R: (toucan, 12), S: (leopard, 8), T: (ptarmigan, 5)}



Figure A17: Randomly selected images from IN100 using image-to-image and transmodal based methods (scale=1.0). IN100-LT samples were used as source images. Class names and the number of real samples are as follows: {A: (green iguana, 1000), B: (sunglasses, 765), C: (gazelle, 651), D: (paintbrush, 472), E: (mongoose, 342), F: (lifeboat, 223), G: (beaker, 180), H: (African bush elephant, 153), I: (tractor, 117), J: (bee, 80), K: (Border Terrier, 61), L: (bulbul, 55), M: (firetruck, 40), N: (guinea pig, 32), O: (Irish Setter, 26), P: (Sealyham Terrier, 21), Q: (bald eagle, 17), R: (toucan, 12), S: (leopard, 8), T: (ptarmigan, 5)}



Figure A18: Qualitative results for iNaturalist2018. Due to the lack of class names, we list class indices. {A: 5533, B: 6334, C: 1211, D: 2457, E: 2616, F: 3508, G: 3976, H: 4543, I: 4094, J: 4071 }

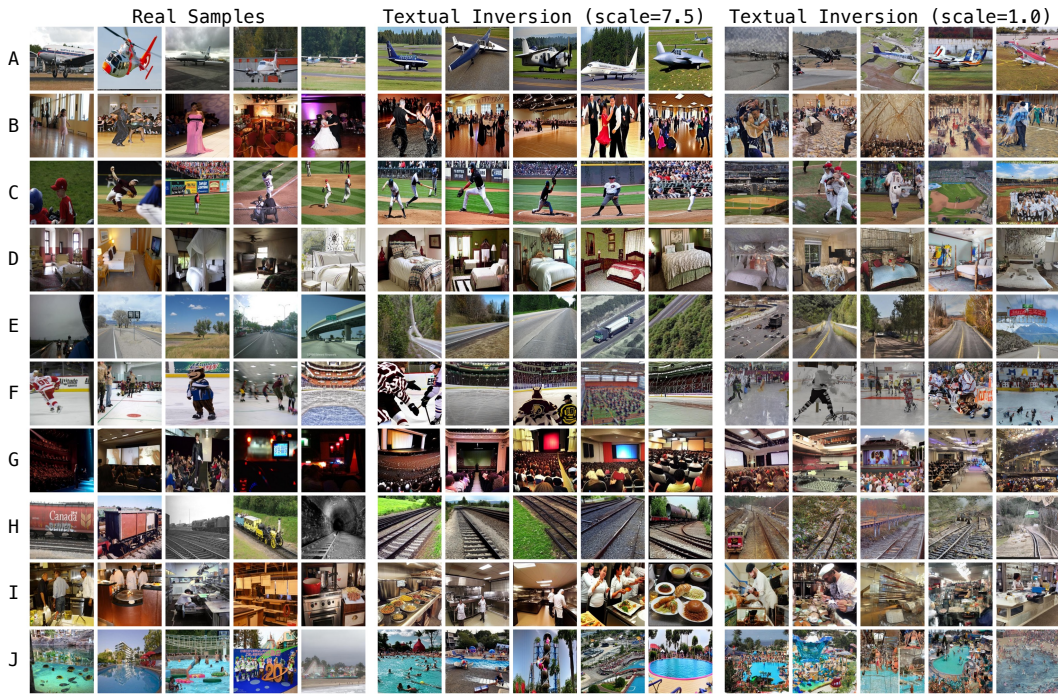


Figure A19: Qualitative results for Places-LT. Class names are as follows: {A: airfield, B: ballroom, C: baseball field, D: bedroom, E: highway, F: ice skating rink indoor, G: movie theater indoor, H: railroad track, I: restaurant kitchen, J: water park }

E Broader impacts and compute resources

Fill-Up aims to mitigate the challenge of imbalanced real-world distributions by leveraging large-scale generative models. Given the brittle nature of training generative models under imbalanced data and the heavy computational resources required for full fine-tuning of state-of-the-art models, our approach, which focuses on lightweight optimization of a single text token per class, shows a promising direction for handling long-tailed distributions. In the domain of long-tailed recognition, where synthetic data generation approaches have relatively lagged behind state-of-the-art methods, our approach provides strong intuitions and promising results by leveraging a pre-trained generative model. With the rapid development of cutting-edge generative models, we expect more approaches based on synthetic data generation to be actively discussed and explored in the field.

Recent explosive developments in generative modeling brought forth numerous potential dangers, encompassing ethical concerns, privacy issues, legal implications, and much more. Built on top of recent advancements, our method also inherently shares potential risks and challenges associated with these models. Two specific issues that we have observed pertain to privacy and nudity, particularly in classes that involve human subjects. In specific classes such as sunscreen, where individuals wearing exposed clothing are prevalent, the optimization process of Textual Inversion may yield unintended images that could potentially be deemed NSFW (Not Suitable For Work). This issue can be further exacerbated by the tendency of generative models to memorize instances, such as real human beings, that they encounter during the training process. Safeguarding privacy and maintaining high ethical standards are our top priorities throughout this work. We will take additional precautions in the future application and release of any related assets. Our code and selected data will be made public upon paper publication after careful review.

In terms of computing resources, we employ a combination of RTX 3090 and A100 GPUs. With a single RTX 3090, Textual Inversion takes approximately 30 minutes per class, following our heuristic, on ImageNet-LT dataset. Additionally, sampling a single image with a scale of 1.0 takes about 3 seconds. For classification, training ResNet50 on ImageNet-LT with 2.6M filled-up images requires one day of computation with eight A100 GPUs. Computational requirements for IN100 ablation, Places-LT, and iNaturalist2018 scale with mentioned timings.