Deliberate Practice with Synthetic Data

Reyhane Askari-Hemmat[1*] **Mohammad Pezeshki**[1*]
Pietro Astolfi[1] **Melissa Hall**[1] **Florian Bordes**[1]
Jakob Verbeek[1] **Michal Drozdzal**[1] **Adriana Romero-Soriano**[1,2,3,4]

[1]FAIR at Meta - Montreal, Paris, and New York City labs
[2]Mila [3] McGill University [4]Canada CIFAR AI chair

Abstract

Deliberate practice for humans is the process of improving one's skills by leveraging external feedback while actively seeking out and correcting mistakes. The current status quo in machine learning is to use static datasets, composed of real or generated data, to train models. While state-of-the-art generative models can serve as an infinite source of synthetic data to train downstream models, prior work has shown that simply increasing the dataset size results in diminishing improvements in model accuracy. In this work, we design a framework that generates synthetic data to improve the performance of a downstream machine learning model. The framework incorporates feedback from the downstream model to refine the generated data used to train the model throughout the training process. In particular, we employ deliberate practice for neural network training to generate challenging synthetic examples tailored to the model's weaknesses at any stage of training, replacing easier, less informative examples in the dataset. With a fixed-size synthetic dataset throughout training, this approach yields over 14% and 8% accuracy improvement on ImageNet-100 and ImageNet-1000, respectively.

1 Introduction

The availability of large amounts of labeled data is a crucial factor in the success of many machine learning models. However, collecting and labeling data can be time-consuming and expensive, and in some cases it is impossible to obtain enough labeled data to train a high-performing model. This has led to a growing interest in synthetic data generation using large-scale pre-trained generative models [4, 8, 12]. In this paper, we focus on using diffusion models [5] as the data source for downstream applications. While synthetic data generation holds potential, a naive approach of generating large quantities of labeled synthetic data and using it to train a neural network quickly reaches a point of diminishing returns. As more synthetic examples are generated, the added value of each new example declines because synthetic data often lacks the diversity of real-world data [3]. In many cases, achieving a similar level of performance to real data would require generating orders of magnitude more synthetic data [3, 11], leading to inefficiencies in computation and storage. Therefore, there is a necessity for more advanced techniques of synthetic data generation that can close the gap between synthetic and real data performance.

Drawing inspiration from the concept of *deliberate practice* in human learning, where learners improve their skills by focusing on tasks that lie at the edge of their current abilities, we introduce *deliberate practice for synthetic data generation*. Our framework leverages a generative model, such as a diffusion model, to generate synthetic examples conditioned on known labels for a downstream model, such as a classifier. However, unlike previous approaches that generates data in a static fashion [3, 8], our framework incorporates a *feedback loop* from the classifier being trained, back to

*Equal contribution

NeurIPS 2024 Workshop on Adaptive Foundation Models.

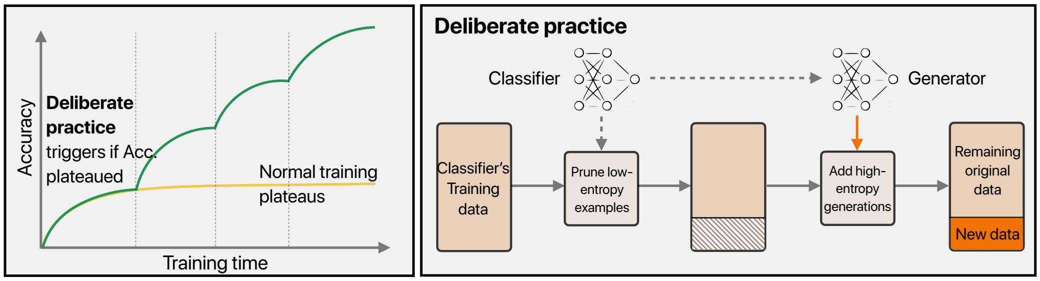

Figure 1: Schematic illustration of the Deliberate Practice framework.

the generative model. This feedback loop ensures that the generative model is continuously informed about the current weaknesses and strengths of the classifier. Consequently, the generative model focuses on generating new examples that are deliberately challenging for the classifier, thereby promoting more effective and targeted learning.

In this paper, we apply *Deliberate Practice* (DP) to image classification tasks, using an off-the-shelf large-scale pre-trained diffusion model as the data source. DP operates iteratively: starting with an initial set of synthetic data, we train a classifier until its performance on a real validation set plateaus. At this point, we leverage the classifier's current state to guide the diffusion model in generating examples that are specifically challenging, based on a measure of difficulty. These harder examples are added to the training set, replacing easier examples that are no longer informative for the classifier. This controlled pruning and replacement strategy maintains a constant dataset size, preventing the explosion of synthetic data while ensuring the dataset remains challenging and useful for the classifier.

To validate the effectiveness of our approach, we conduct experiments on ImageNet-100 [10] and ImageNet-1000 [1], demonstrating significant improvements in classifier performance compared to naive synthetic data generation methods. We also present a series of ablation studies to investigate the contribution of each component of our framework. Our findings suggest that the deliberate practice framework can provide a more sustainable path forward for training large-scale models when real data is limited or unavailable.

2 Preliminaries: steering generations of diffusion kodels

In this work, we focus on Latent Diffusion Models (LDMs) and specifically Denoising Diffusion Implicit Models (DDIMs) [9], a variant of diffusion models that has a deterministic sampling process, allowing for faster inference while maintaining generation quality.

During inference, DDIMs perform a reverse diffusion process, where the model sequentially denoises a noisy sample. This reverse process at timestep t is defined as:

$$x_{t-1} = \sqrt{\xi_{t-1}}\hat{x}_{0,t} + \underbrace{\sqrt{1-\xi_{t-1}-\sigma_t^2}\cdot\epsilon_\theta^{(t)}(x_t)}_{\text{direction pointing to } x_t} + \underbrace{\sigma_t\epsilon_t}_{\text{random noise}}, \tag{1}$$

where x_t is the current noisy sample, $\epsilon_t \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ represents Gaussian noise, and ξ_t and σ_t are time-dependent coefficients. The term $\hat{x}_{0,t}$ is the estimated denoised sample, computed using the following approximation:

$$\hat{x}_{0,t} = \text{DDIMApprox}(x_t) := \frac{x_t - \sqrt{1-\xi_t}\epsilon_\theta^{(t)}(x_t)}{\sqrt{\xi_t}}, \tag{2}$$

where $\epsilon_\theta^{(t)}$ represents the score function, approximated by a pretrained denoising network parameterized by θ, which predicts the noise at time step t. Formally, the score function is defined as:

$$\epsilon_\theta^{(t)} \approx \nabla_x \log p(x_t). \tag{3}$$

It is often desirable to generate data samples conditioned on specific inputs, such as class labels or text prompts. Latent Diffusion Models (LDMs) typically incorporate these inputs directly into the score function during training, enabling conditional generation. However, even when a diffusion model lacks explicit conditioning, classifier-guidance [2] can be used to steer the generative process by leveraging a pretrained classifier to estimate $\log p(y|x_t)$. In this case, the score function is modified as:

$$\epsilon_\theta^{(t)} \approx \nabla_x \log p_\gamma(x_t|y) = \nabla_x \log p(x_t) + \gamma \nabla_x \log p(y|x_t), \tag{4}$$

where γ is a scaling factor that controls the strength of the class-conditional generation.

A similar approach has been proposed to generate 'useful' examples for a classification task [4]. By defining the 'usefulness' of an example based on the entropy of a classifier (f_ϕ) on that example, the sampling process can be adjusted to generate high-entropy examples through the so-called *feedback-guided sampling*:

$$\nabla_{x_t} \log \hat{p}_\omega(x_t|h, y) = \nabla_{x_t} \log \hat{p}_\theta(x_t) + \gamma \nabla_{x_t} \log \hat{p}_\theta(y|x_t) + \omega \nabla_{x_t} H(f_\phi(\hat{x}_{0,t})), \tag{5}$$

where $h \in \{0, 1\}$ is a binary variable that describes whether a sample is useful for the classifier or not and $H(f_\phi(\hat{x}_{0,t}))$ is the entropy of the classifier.

3 Deliberate practice for neural networks

Our proposed framework revolves around the dynamic interaction between model training, feedback-guided synthetic data generation, and strategic pruning of the training set. This process ensures that the model continuously receives challenging and informative training examples, resulting in a more efficient learning process. The framework consists of three key components: patience, feedback, and pruning, which are applied iteratively.

Patience. The patience mechanism acts as a monitoring tool to measure the classifier's performance and determine when adjustments are needed in the training data. After each epoch of training, we evaluate the model's validation accuracy using a real, held-out validation set. If the validation accuracy improves, training proceeds as usual. However, if there is no improvement, a patience counter increments. The patience counter tracks how many consecutive epochs pass without any improvement in the validation performance. If this counter exceeds a pre-defined patience threshold, we assume that the current training data may no longer be sufficiently challenging for the model. At this point, a refinement mechanism for the training set is triggered to intervene by generating more informative examples that are tailored to the model's weaknesses and swapping them with easy examples.

Feedback. Instead of naively generating additional synthetic data, we apply feedback guidance [4] using the *entropy of the classifier* at the current epoch as the feedback criterion. Entropy serves as a measure of the model's uncertainty about its predictions. Examples with high entropy are those for which the model is uncertain, making them harder for the model to classify correctly. By actively generating high-entropy examples, we generate synthetic examples that target the model's weaknesses. This ensures that training is focused on overcoming the model's current limitations, rather than sub-optimally training on examples the model has already learned well.

Pruning. Once the feedback mechanism generates new, challenging examples, a pruning mechanism is applied on the training set. The goal of this step is to replace least informative examples with the newly generated synthetic ones, maintaining a fixed dataset size. During the pruning process, we evaluate the difficulty of examples in the current training set, again using the entropy of the classifier as the metric. By continuously pruning easy examples and replacing them with harder ones, we ensure that the dataset remains challenging.

This iterative process continues throughout the model's training, ensuring the model focuses on learning from the most informative examples, see Figure 1 for a schematic overview.

4 Experiments

We use a large-scale pretrained latent diffusion model (LDM-1.5) [7] to generate synthetic examples. For optimization, we use the Warmup-Stable-Decay (WSD) learning rate scheduler [6], which stabilizes training when samples are swapped and maintains a constant learning rate throughout most of the process, allowing the model to effectively learn from newly generated examples.

Our training begins with a static synthetic dataset. When the patience threshold is reached, a refinement mechanism is triggered. At this point, the classifier evaluates the dataset to identify low-entropy examples, and feedback-guided generation is used to generate more challenging examples.

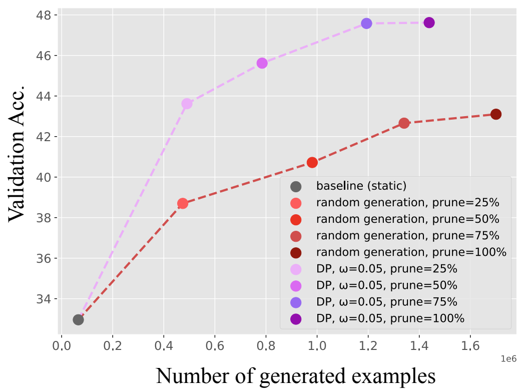

	Val. Acc.	**Test Acc.**
IN-100 static	32.96	35.06
IN-100 random gen.	43.10	45.35
IN-100, DP	47.62	49.84
IN-1000	32.94	33.65
IN-1000, DP	41.13	42.04

Figure 2: (**Left**) Scaling ability of DP compared to the baseline (static) and random generation on ImageNet-100. We observe that DP allows for much better scaling. (**Right**) ImageNet-100 and ImageNet-1000 classification results on a ViT-B model with **synthetic** data.

We then swap the easier examples with these harder ones, allowing the model to continue training with more informative data.

We tune the hyperparameters ω and λ, finding that the optimal λ without feedback is 2, while with feedback, values of 3 or 4 yield the best results. An ω value of 0.05 proves optimal in most cases. The training set size remains fixed throughout at 65k examples for ImageNet-100 and 650k examples for ImageNet-1000. We use the real images in the validation set to control the patience counter. To evaluate our method, we use the full real training sets of ImageNet-100 and ImageNet-1000 as test sets.

Figure 2 **(left)** shows how the validation accuracy scales as the model sees more examples. Although the training set size remains fixed, we track the total number of generated examples to observe scaling. On ImageNet-100, deliberate practice (DP) significantly improves scaling. For instance, without DP, the best accuracy reaches 43.1% after seeing more than 1.7 million examples, while with DP, we achieve 43.62% with only 490k examples. Figure 2 **(right)** also compares the validation and test accuracy with and without DP for both ImageNet-100 and ImageNet-1000.

We also perform an ablation study on the size of the pruning set, choice of scheduler and patience hyper-parameter. We observe that WSD reaches better performance than cosine scheduler in this setup. Furthermore, we study the pruning size as a function of the size of the training set. We observe when a larger data is pruned and replaced, we reach better accuracy but that comes at the cost of generating more examples.

Table 1: Ablation study on Imagenet-100

	Scheduler	**Patience**	**Pruning**	**Val. Acc.**	**Test Acc.**
IN-100 static	Cosine	-	0%	29.78	31.74
IN-100 static	WSD	-	0%	32.96	35.06
IN-100, DP	WSD	10	25%	42.90	45.08
IN-100, DP	WSD	8	25%	44.14	45.89
IN-100, DP	WSD	8	50%	45.62	47.81
IN-100, DP	WSD	8	75%	47.58	49.19
IN-100, DP	WSD	8	100%	47.62	49.84
IN-100, random gen.	WSD	8	100%	43.1	45.35

5 Conclusion

In this work, we introduced a novel framework for improving downstream model performance by applying the concept of deliberate practice to neural networks. Rather than relying on static datasets, our approach incorporates feedback from the model during training to generate challenging synthetic examples that target the model's current weaknesses. This dynamic feedback loop ensures that the training dataset remains informative and difficult, without requiring an ever-expanding dataset size. Our results demonstrate significant improvements on ImageNet-100 and ImageNet-1000.

References

[1] Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hierarchical image database. In *CVPR*, 2009.

[2] Prafulla Dhariwal and Alexander Nichol. Diffusion models beat GANs on image synthesis. In *NeurIPS*, 2021.

[3] Lijie Fan, Kaifeng Chen, Dilip Krishnan, Dina Katabi, Phillip Isola, and Yonglong Tian. Scaling laws of synthetic images for model training... for now. In *CVPR*, 2024.

[4] Reyhane Askari Hemmat, Mohammad Pezeshki, Florian Bordes, Michal Drozdzal, and Adriana Romero-Soriano. Feedback-guided data synthesis for imbalanced classification. *arXiv preprint*, 2310.00158, 2023.

[5] Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *NeurIPS*, 2020.

[6] Shengding Hu, Yuge Tu, Xu Han, Chaoqun He, Ganqu Cui, Xiang Long, Zhi Zheng, Yewei Fang, Yuxiang Huang, Weilin Zhao, et al. Minicpm: Unveiling the potential of small language models with scalable training strategies. *arXiv preprint*, 2404.06395, 2024.

[7] Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Björn Ommer. High-resolution image synthesis with latent diffusion models. In *CVPR*, 2022.

[8] Joonghyuk Shin, Minguk Kang, and Jaesik Park. Fill-up: Balancing long-tailed data with generative models. *arXiv preprint*, 2306.07200, 2023.

[9] Jiaming Song, Chenlin Meng, and Stefano Ermon. Denoising diffusion implicit models. *arXiv preprint*, 2010.02502, 2020.

[10] Yonglong Tian, Dilip Krishnan, and Phillip Isola. Contrastive multiview coding. In *Computer Vision–ECCV 2020: 16th European Conference, Glasgow, UK, August 23–28, 2020, Proceedings, Part XI 16*, pages 776–794. Springer, 2020.

[11] Yonglong Tian, Lijie Fan, Kaifeng Chen, Dina Katabi, Dilip Krishnan, and Phillip Isola. Learning vision from models rivals learning vision from data. In *CVPR*, 2024.

[12] Yifan Zhang, Daquan Zhou, Bryan Hooi, Kai Wang, and Jiashi Feng. Expanding small-scale datasets with guided imagination. In *NeurIPS*, 2024.