

From Reconstruction to Compression: Reinforcement Learning in Diffusion-Based Dataset Distillation

Anonymous submission

Abstract

Dataset distillation synthesizes compact datasets that retain the training utility of much larger ones. While diffusion models are natural candidates for this task due to their generative capabilities, there are few methods that adopt them in dataset distillation compared to the matching-based approaches and label-relaxation approaches. A key reason is the fundamental mismatch between diffusion objectives and distillation goals: diffusion models are trained to reconstruct high-fidelity data, whereas distillation requires compressed, task-relevant representations. We address this gap by proposing a reinforcement learning (RL)-guided framework that steers diffusion models from reconstruction toward compression. By formulating sampling as a decision process, we optimize the generative trajectory using rewards derived from student model performance. This enables the generation of synthetic samples that maximize learning utility under strict compression budgets. Unlike prior static modifications of the diffusion process, our method dynamically adapts generation based on downstream outcomes. Experiments on standard benchmarks show that our RL-guided diffusion approach consistently improves both performance and efficiency, advancing the frontier of generative dataset distillation. Our code will be available at <https://github.com/Galleriess/DiT\GRPO>

Introduction

Dataset distillation (Zhao, Mopuri, and Bilen 2021; Wang et al. 2018; Cazenavette et al. 2022; Chen et al. 2025; Kim et al. 2022; Cui et al. 2023; Gu et al. 2024) emerges as a scalable alternative to coreset selection, with a critical shift in paradigm: instead of selecting a subset from the original dataset, it aims to synthesize a small number of synthetic samples that can train models to comparable performance. This synthesis-oriented nature makes generative models, particularly diffusion models (Peebles and Xie 2023; Chen et al. 2025), natural candidates for distillation backbones. Given their ability to model complex data distributions and generate diverse samples, diffusion models appear well-suited to construct informative, compact datasets. As visualized in Figure 1, this shift in objective—from faithfully reconstructing individual samples to selectively synthesizing task-informative ones—calls for a fundamental rethinking of how diffusion models are employed in this context.

However, despite the popularity of dataset distillation (Sachdeva and McAuley 2023; Lei and Tao 2024) and the great success of diffusion models (Ho, Jain, and Abbeel 2020; Peebles and Xie 2023), their integration remains rare. We identify a key reason behind this gap: a fundamental mismatch in objectives. Diffusion models are trained to reconstruct individual data samples with high fidelity by reversing a noise corruption process. In contrast, dataset distillation is inherently a compression task, aiming to concentrate task-relevant information into a minimal number of synthetic instances. As a result, the stronger a diffusion model becomes at reconstructing original data, the less effective it is for generating compressed data optimized for downstream learning.

Existing dataset distillation methods primarily fall into two families: (i) matching-based approaches that directly optimize synthetic samples to approximate gradients or training trajectories (Nguyen, Chen, and Lee 2021; Cazenavette et al. 2022; Du et al. 2023; Du, Shi, and Zhou 2023); and (ii) label-relaxation approaches such as SRe2L (Yin, Xing, and Shen 2023), which guide learning through softened targets. While both have achieved considerable progress, they suffer from scalability and generalization bottlenecks—either due to reliance on differentiable supervision or overly rigid label semantics. These limitations further motivate a flexible, model-driven distillation framework, one that can generate rather than optimize, and adapt based on downstream training outcomes.

To this end, we introduce *Compression-Oriented Distillation* (COD), a novel framework that reformulates diffusion sampling as a reinforcement learning (RL) (Guo et al. 2025; Ouyang et al. 2022; Fan et al. 2023; Black et al. 2024; Zhou et al. 2024) problem aimed at utility-aware data compression. Instead of statically following the reverse denoising process, we learn a policy that dynamically controls the generative trajectory to favor samples that are compact yet highly effective for downstream training. By directly optimizing this policy with task-driven reward signals, our method moves beyond heuristic guidance and enables principled generation of high-utility synthetic data under strict budget constraints.

We instantiate this framework COD using Group Relative Policy Optimization (GRPO) (Guo et al. 2025), a lightweight yet stable policy optimization method that avoids explicit value estimation. To guide sample genera-




Figure 1: **Illustration of reconstruction-oriented vs. compression-oriented distillation.** Top: Existing diffusion-based distillation reconstructs high-fidelity samples by directly mapping a latent code z to individual data points, optimizing for pixel-level realism. Bottom: Our proposed Compression-Oriented Distillation introduces reward-guided sampling, enabling z to dynamically steer generation toward task-informative and compact representations, thereby capturing dense supervision with fewer samples.

tion, we design a reward function that combines two complementary components: (1) an entropy-based signal (R_{Ent}) that promotes informative samples by maximizing predictive uncertainty (Paul, Ganguli, and Dziugaite 2021), and (2) a diversity-aware penalty (R_{Div}) that discourages redundancy by comparing with a memory bank of previously generated outputs. This reward-driven feedback loop steers the diffusion model beyond pixel-level fidelity, enabling it to explore and exploit regions of the data space that are optimized for learning efficiency (Black et al. 2024). Compared to prior methods like Minimax Diffusion that statically reshape sampling behavior, our approach offers dynamic, goal-aware control over generative processes.

In summary, this paper presents the first comprehensive study of reinforcement learning for controlling diffusion-based dataset distillation. By reinterpreting generative modeling as a compression-driven decision process, we bridge

the gap between reconstruction-centric generation and training-centric distillation, setting the stage for a new class of adaptive, goal-aware synthetic data pipelines. Our contributions are summarized as follows:

- We identify a fundamental mismatch between diffusion models and dataset distillation: diffusion prioritizes reconstruction, while distillation demands compression. This insight explains the limited integration of the two paradigms.
- We propose *Compression-Oriented Distillation (COD)*, a novel framework that formulates diffusion sampling as a reinforcement learning process guided by downstream utility.
- We instantiate COD using Group Relative Policy Optimization (GRPO) with a reward function combining entropy-based informativeness and diversity-aware regularization, enabling principled and adaptive sample generation.

Preliminaries

Problem Formulation: Dataset Distillation

Given a large-scale dataset $\mathcal{T} = \{(x_i, y_i)\}_{i=1}^N$, where $x_i \in \mathbb{R}^d$ are data samples drawn i.i.d. from a natural distribution $\mathcal{D}$. We denote $y_i \in \mathcal{Y} = \{1, \dots, C\}$ to represent class labels. Dataset distillation aims to construct a compact synthetic dataset $\mathbb{S} = \{(s_j, \tilde{y}_j)\}_{j=1}^M$ with $M \ll N$ such that a model trained solely on $\mathbb{S}$ performs comparably to one trained on $\mathcal{T}$ (Wang et al. 2018; Cazenavette et al. 2022):

$$\mathbb{S}^* = \arg \min_{\mathbb{S} \subset \mathbb{R}^d \times \mathcal{Y}} \mathbb{E}_{(x,y) \sim \mathcal{D}} [\ell(f_{\theta_{\mathbb{S}}}(x), y)], \quad (1)$$

where $\theta_{\mathbb{S}}$ denotes the model parameters obtained by training on $\mathbb{S}$, and $\ell(\cdot)$ is a task-specific loss function (e.g., cross-entropy). Most existing dataset distillation methods fall into two broad categories: matching-based approaches and label-relaxation approaches.

Matching-based approaches optimize synthetic data by aligning gradients or training trajectories between real and synthetic datasets (Cazenavette et al. 2022; Du et al. 2023; Cui et al. 2023; Zhao, Mopuri, and Bilen 2021). A common objective is gradient matching (Zhao, Mopuri, and Bilen 2021):

$$\min_{\mathbb{S}} \sum_{(s_j, \tilde{y}_j) \in \mathbb{S}} \|\nabla_{\theta} \ell(f_{\theta}(s_j), \tilde{y}_j) - \nabla_{\theta} \ell(f_{\theta}(x_i), y_i)\|^2, \quad (2)$$

where the gradient computed on synthetic samples is forced to approximate that from real data. Trajectory-based variants extend this idea across multiple steps of optimization. While effective, such methods often require differentiability, second-order gradients, and suffer from limited scalability on larger datasets.

Label-relaxation approaches assign soft labels to synthetic samples to improve generalization (Yin, Xing, and Shen 2023; Shao et al. 2024; Sun et al. 2024). Instead of using hard one-hot labels y_j , each synthetic sample is paired with a learnable probability vector $\tilde{y}_j \in \Delta^{C-1}$:

$$\min_{\mathbb{S}} \sum_{(s_j, \tilde{y}_j) \in \mathbb{S}} \ell(f_{\theta}(s_j), \tilde{y}_j), \quad (3)$$

where $\tilde{y}_j$ encodes label uncertainty or class similarity. While this approach improves performances significantly, it requires a pretrained model to serve as the teacher model to generate soft labels. Xiao and He (2024) reveals that removing the soft labels will cause dramatic performance drop in Label-relaxation approaches.

Both paradigms rely on direct supervision over synthetic instances. In contrast, our framework shifts the problem toward a reward-driven generative formulation, using reinforcement learning to synthesize utility-optimized training data. Unlike coreset selection (Bachem, Lucic, and Krause 2017; Chen, Welling, and Smola 2012; Har-Peled and Kushal 2005; Sener and Savarese 2017; Tsang et al. 2005), dataset distillation synthesizes new data instances rather than selecting from $\mathcal{T}$. This makes generative models, in particular the diffusion models, a promising approach for dataset distillation.

Diffusion Models for Generative Synthesis

Diffusion models (Ho, Jain, and Abbeel 2020; Peebles and Xie 2023; Sun et al. 2025) generate data via a two-stage process: a forward noising process and a reverse denoising process. Let $x_0 \sim \mathcal{D}$ denote a real data sample; f_θ denote the denoising network f parameterized with θ . The forward process gradually corrupts x_0 with Gaussian noise, yielding a sequence $\{x_t\}_{t=0}^T$. The reverse process then aims to iteratively reconstruct x_0 from pure noise $x_T \sim \mathcal{N}(0, I)$, by learning a parameterized denoising network f_θ .

Given a discretized time schedule $\{t_i\}_{i=0}^N$, the sampling trajectory starts from $x_0 \sim \mathcal{N}(0, b(t_{\max})^2 I)$ and proceeds via the following iterative update:

$$\mathbf{x}_{i+1} := \kappa_i \mathbf{x}_i + \eta_i f_\theta(\mathbf{x}_i | t_i) + \zeta_i \tilde{\epsilon}_i, \quad (4)$$

where $\tilde{\epsilon}_i \sim \mathcal{N}(0, I)$ is an optional sampling noise term (present only in SDE-based solvers), and κ_i , η_i , and ζ_i are time-dependent coefficients derived from the training-time noise schedule.

This reverse process is fundamentally designed to reconstruct a high-fidelity individual instance from Gaussian noise. The denoising network f_θ is explicitly trained to reverse the corruption applied in the forward process, which encourages the generation of samples that closely match the data distribution in pixel space or feature space. As a result, the learned generative trajectory is inherently biased toward reproducing realistic and data-faithful samples—making it highly suitable for reconstruction tasks, but potentially sub-optimal for generating compressed or task-optimized representations such as those needed in dataset distillation.

Method

The Reverse Process Performs Denoising, Not Compression

The reverse diffusion process is inherently designed as a denoising mechanism (Ho, Jain, and Abbeel 2020), not a compression pipeline. At each time step, the denoising network f_θ estimates either the original clean sample x_0 or the noise ϵ added during the forward process, conditioned on a noisy input x_t . This iterative reconstruction from Gaussian noise




Figure 2: **Impact of Diffusion Steps on Distillation Accuracy and FID.** We vary the number of sampling steps in the reverse diffusion process and observe a trade-off between image fidelity and distillation performance. As the number of steps increases, the Fréchet Inception Distance (FID) consistently improves, indicating better reconstruction quality. However, the distillation accuracy declines, confirming that high-fidelity samples are not necessarily more informative for downstream learning. This highlights a fundamental mismatch between reconstruction and compression objectives in diffusion-based distillation.

is shown in Equation 4, where the objective is to minimize reconstruction error between predicted and true clean samples.

This training formulation encourages the model to reproduce high-fidelity instances that resemble the original dataset distribution. However, in dataset distillation, the objective shifts: instead of reproducing all modes of the data, we seek to selectively generate samples that are maximally informative for training under tight data budgets. As shown in Figure 2, while the Fréchet Inception Distance (FID) improves with more steps the classification performance of distilled datasets degrades. This confirms that better reconstruction does not equate to better compression, and high-fidelity images are not necessarily information-dense for learning.

To address this mismatch, we aim to systematically bias the reverse process away from denoising and toward compression. The most direct solution appears to be modifying the reverse process itself, which replaces the reconstruction-oriented dynamics with a utility-driven sampling process that prioritizes the generation of high-information-content samples.

However, this direction faces several fundamental limitations. First, shifting the diffusion objective toward compression would require defining or learning a new target distribution $\mathcal{S}$ that represents the ideal distilled dataset distribution. We denote by $\mathcal{S}$ the empirical distribution formed by all synthetic datasets produced by existing dataset distillation methods. Each method $\mathcal{A}_i$ generates a synthetic dataset

$S_i = \mathcal{A}_i(\mathcal{T})$, where $\mathcal{T}$ is the original training set. The distribution $\mathcal{S}$ can thus be conceptualized as:

$$\mathcal{S} = \mathbb{P}_{\mathcal{A} \sim \mathcal{M}} [\mathcal{A}(\mathcal{T})], \quad (5)$$

where $\mathcal{M}$ is the space of known distillation algorithms (e.g., DM, MTT, SRe2L, IGD).

Attempting to train a diffusion model to approximate $\mathcal{S}$ poses three major challenges. First, $\mathcal{S}$ is implicitly defined and lacks a closed-form representation—sampling from it requires exhaustively running and storing outputs from many distillation pipelines. Second, synthetic datasets from different methods are structurally inconsistent, often varying in label granularity, resolution, or supervision format, making them hard to unify under a coherent distribution. Third, even if $\mathcal{S}$ were learnable, any generative model trained to replicate it would be fundamentally limited by the diversity and quality of the existing methods. That is, the best it can do is imitate prior solutions, but never surpass them.

This motivates us to abandon the idea of statically fitting a proxy to $\mathcal{S}$ and instead adopt a reward-driven sampling mechanism that actively explores beyond it.

Reinforcement Learning for Compression-Oriented Diffusion Sampling

To explore beyond the empirical limits of $\mathcal{S}$, we propose to steer the generative trajectory using reinforcement learning (RL). Rather than statically mimicking prior synthetic datasets, our goal is to actively discover high-utility samples by assigning rewards to diffusion outcomes based on their downstream training performance. This framing naturally casts the sampling procedure as a sequential decision process, where the reverse steps of the diffusion model form a Markov chain governed by a policy π_ϕ .

The idea of prioritizing high-information-content samples is inspired by the success of dataset pruning, which shows that even within natural datasets, only a subset of examples contributes meaningfully to generalization. This observation implies that data samples are inherently unequal in the perspective of information.

Concretely, at each reverse timestep t , the policy $\pi_\phi(a_t | x_t)$ selects an action a_t —such as modifying the noise prediction or controlling the denoising step size—based on the current sample state x_t . The final output x_0 is evaluated via a reward function $R(x_0)$ that reflects its utility for distillation. This reward can be instantiated using downstream classification accuracy, teacher-student agreement, or information-theoretic proxies such as entropy or mutual information. The policy is then optimized to maximize expected reward:

$$\max_{\phi} \mathbb{E}_{x_0 \sim \pi_\phi} [R(x_0)]. \quad (6)$$

This formulation transforms the role of the diffusion model from a passive denoiser into an active sampler that learns to navigate toward information-rich regions of the data space. Unlike traditional guidance strategies, which rely on heuristics or task-agnostic priors (e.g., classifier gradients or class embeddings), our RL-based controller can be trained end-to-end to align sample generation directly with dataset distillation objectives.

In doing so, we depart from the conventional reconstruction pipeline and reframe dataset distillation as a compression-driven search problem over the generative trajectory space. This dynamic mechanism allows us not only to circumvent the ill-posed nature of $\mathcal{S}$ but also to transcend the limitations of existing synthetic datasets by continuously refining the sampling policy in response to feedback from training performance.

Instantiating RL-Based Sampling with GRPO and Entropy Rewards

To realize the RL formulation described above, we adopt Group Relative Policy Optimization (GRPO) as our policy learning algorithm. While standard reinforcement learning methods such as PPO (Schulman et al. 2017) offer stable policy improvement, they rely heavily on value function estimation and surrogate clipping objectives, which are costly and unstable in our context. Diffusion-based sampling is inherently high-dimensional and slow, and accurate value estimation across diverse generative trajectories is impractical.

GRPO circumvents these issues by discarding value estimation altogether. Instead of modeling long-term returns, GRPO computes *relative advantages* within a group of sampled actions. Specifically, for each reverse step state x_t , we sample a group of candidate actions $\{a_t^{(i)}\}_{i=1}^G$, generate corresponding samples $\{x_0^{(i)}\}$, and compute their rewards $\{r_i = R(x_0^{(i)})\}$. We then normalize the rewards using z-score normalization to obtain relative advantages:

$$\bar{r}_i = \frac{r_i - \text{mean}(\mathbf{r})}{\text{std}(\mathbf{r})}, \quad (7)$$

where $\mathbf{r} = \{r_1, \dots, r_G\}$ is the reward vector within the group. The policy π_ϕ is updated to increase the probability of actions with higher $\bar{r}_i$, encouraging exploration of trajectories that outperform their peers without needing explicit value estimation.

$$R_{\text{Ent}}(x_0) = \mathcal{H}(f_\theta(x_0)) = - \sum_{c=1}^C p_c \log p_c, \quad (8)$$

where p_c denotes the softmax probability assigned to class c by the model pretrained on original dataset. High entropy indicates model uncertainty and implies that the sample x_0 lies near the decision boundary—thus being more informative for training. Unlike label-matching losses, entropy rewards are task-agnostic, differentiable, and directly aligned with the goal of generating useful training signals.

While entropy-based rewards encourage the generation of uncertain and potentially informative samples, optimizing solely for entropy may lead to *mode collapse*—the repeated synthesis of ambiguous yet similar examples. To mitigate this, we introduce a diversity reward that explicitly penalizes redundancy among generated samples.

We maintain a memory bank $\mathcal{B}$ that stores the embeddings or output logits of previously generated synthetic samples. The details to implement the memory bank is clearly stated in the experiment section. For each new candidate x_0 , we compute its similarity to the most similar entry in the bank

Algorithm 1: Compression-Oriented Diffusion for Dataset Distillation

Require: Pretrained diffusion model f_θ , pretrained evaluation network h_{pt} , policy π_φ , memory bank $\mathcal{B}$, group size G , reward weight λ

- 1: **for** each RL iteration **do**
 - 2: Sample G initial noise vectors $\{x_T^{(i)}\}_{i=1}^G \sim \mathcal{N}(0, I)$
 - 3: **for** each $x_T^{(i)}$ **do**
 - 4: Sample actions $\{a_t^{(i)}\}$ from $\pi_\varphi(a_t \mid x_t^{(i)})$ at each timestep
 - 5: Generate sample $x_0^{(i)}$ via controlled reverse diffusion trajectory
 - 6: Compute entropy reward:

$$R_{\text{Ent}}(x_0^{(i)}) = - \sum_{c=1}^C p_c \log p_c \quad \text{where } \mathbf{p} = h_{\text{pt}}(x_0^{(i)})$$
 - 7: Compute diversity penalty:

$$R_{\text{Div}}(x_0^{(i)}) = -\lambda \cdot \max_{x' \in \mathcal{B}} \text{sim}(x_0^{(i)}, x')$$
 - 8: Total reward: $R^{(i)} = R_{\text{Ent}}(x_0^{(i)}) + R_{\text{Div}}(x_0^{(i)})$
 - 9: **end for**
 - 10: Normalize rewards: $\bar{R}^{(i)} = \frac{R^{(i)} - \text{mean}(\mathbf{R})}{\text{std}(\mathbf{R})}$
 - 11: Update policy π_φ using GRPO with $\bar{R}^{(i)}$
 - 12: Update memory bank: $\mathcal{B} \leftarrow \mathcal{B} \cup \{x_0^{(i)}\}_{i=1}^G$
 - 13: **end for**
-

and apply a penalty accordingly. The final reward function becomes:

$$R(x_0) = R_{\text{Ent}}(x_0) + R_{\text{Div}}(x_0), \quad \text{where} \quad (9)$$

$$R_{\text{Div}}(x_0) = -\lambda \cdot \max_{x' \in \mathcal{B}} \text{sim}(x_0, x'). \quad (10)$$

where $\text{sim}(x_0, x')$ measures the similarity between x_0 and a stored sample x' using cosine similarity. $R_{\text{Ent}}(x_0)$ is defined in Equation 8. The hyperparameter $\lambda > 0$ controls the trade-off between uncertainty and novelty.

This diversity-aware reward encourages the sampling policy to explore broader, less redundant regions of the generative space—promoting sample diversity without sacrificing informativeness. Empirically, we find that combining entropy and diversity signals leads to synthetic datasets that are both challenging and complementary, resulting in stronger downstream performance.

Altogether, our instantiation combines (1) a policy optimization algorithm (GRPO), (2) an information-theoretic reward signal (entropy), and (3) a diversity-aware constraint (memory bank filtering). These design choices strike a balance between sample informativeness and diversity—two pillars of effective dataset distillation.

Experiments

Experimental Setup

Datasets and baselines. We evaluate our method on three benchmark datasets with increasing resolution and complex-

ity: ImageNet-1K (224×224) and two wellknown subsets of ImageNet (Russakovsky et al. 2015): ImageNette, ImageWoof. For large-scale evaluation, we follow common practice and report top-1 classification accuracy under varying image-per-class (IPC) settings (e.g., 10, 50, 100). We compare with representative baselines including pixel-level methods (DM (Zhao and Bilen 2023), IDC-1 (Kim et al. 2022)), generative methods (DiT (Peebles and Xie 2023)), and fine-tuned diffusion (Minimax (Gu et al. 2024)). Random and Full serve as lower and upper bounds respectively. We also compare with label-relaxation methods including Sre2L (Yin, Xing, and Shen 2023) and G-VBSM (Shao et al. 2024).

Evaluation protocol. Following prior work, we train standard ConvNet or ResNet architectures on the synthetic datasets for 50 to 200 epochs, depending on resolution, using SGD or Adam optimizers. We adopt consistent training schedules across baselines for fair comparison. Unless otherwise stated, evaluation is performed on the same test sets as the original datasets. For ImageNet-1K, pretrained classifiers are also used for reward calculation but not for final evaluation. While SRe2L (Yin, Xing, and Shen 2023) adopts an evaluation protocol using soft labels to have better performance, we only adopt this protocol for ImageNet-1K experiments; all subset results are reported under standard hard-label evaluation for comparability.

Diffusion backbone. We adopt latent DiT (Peebles and Xie 2023) as our diffusion backbone, using a pretrained VAE encoder-decoder to map between image and latent space. All experiments use DDIM (Song, Meng, and Ermon 2022) sampling with 50 steps. The policy network π_φ operates over the noise prediction module of the reverse process and is trained using GRPO. For reward computation, we use a pretrained ImageNet-1K classifier f_{pt} to evaluate entropy.

Memory Bank Implementation. To support the diversity-aware reward R_{div} , we maintain a dynamic memory bank that stores previously generated synthetic samples. At the beginning of policy training, the memory bank is cold-started by populating it with a fixed number of synthetic samples generated unconditionally from the pretrained diffusion model. The total number of stored samples is set equal to the target dataset size to avoid memory growth.

During training, each newly generated sample is compared against existing entries in the memory bank. If a sample is found to be highly similar to any stored instance (based on cosine similarity in a pretrained feature space), it is discarded from reward calculation and excluded from memory bank updates. Otherwise, the sample is appended to the memory bank, and the most similar existing item is removed to maintain a fixed memory size. This design ensures continual refresh of diverse representations without allowing the memory bank to grow, enabling efficient and scalable diversity estimation.

Training details. All experiments are conducted on a single NVIDIA RTX 4090 GPU. Each GRPO update uses group size $G = 4$, and the memory bank retains up to 512

embeddings per class. We adopt cosine annealing for the policy learning rate and freeze the pretrained diffusion and classifier networks throughout the process. Further details (e.g., entropy temperature, policy depth) are detailed in Appendix A.

Experimental Results

Main Results on ImageNet. We report the top-1 classification accuracy on Nette and Woof subsets under varying architectures (ConvNet-6, ResNetAP-10, ResNet-18) (He et al. 2016) and image-per-class (IPC) budgets (10, 50, 100). The results are demonstrated in Table 1. COD achieves consistent improvements over prior diffusion-based methods (DM (Zhao and Bilen 2023), DiT (Peebles and Xie 2023)) and optimization-based methods (IDC-1 (Kim et al. 2022)), especially under low-data regimes such as IPC=10. Compared to the strongest baseline, Minimax, our method exhibits competitive performance across nearly all settings. However, the performance gap between COD and Minimax remains small. This is expected, as both approaches share a similar underlying philosophy: Minimax explicitly modifies the denoising network during training to favor discriminative gradients, while COD fine-tunes the sampling trajectory via reinforcement learning. Despite differing in implementation (training vs. inference), both methods achieve comparable expressivity in guiding generation away from pixel-level fidelity and toward task-relevant content.

We report top-1 accuracy on ImageNet-1K with IPC = 10 and 50 in Table 2. COD achieves the highest accuracy at IPC = 50, surpassing both optimization-based (SRe²L (Yin, Xing, and Shen 2023), G-VBSM (Shao et al. 2024), RDED (Sun et al. 2024)) and generative (DiT (Peebles and Xie 2023), Minimax (Gu et al. 2024)) methods. The consistent improvement demonstrates the effectiveness of our reward-driven policy in scaling to large-scale distillation.

Trade-off Between Accuracy and Fidelity. We investigate how the number of reverse diffusion steps affects the trade-off between sample fidelity and distillation performance. As shown in Figure 2, increasing steps leads to lower Fréchet Inception Distance (FID), indicating improved visual quality. However, distillation accuracy peaks at 75 steps and declines thereafter. This confirms a key insight: higher-fidelity samples are not necessarily more informative for training, and optimizing for visual realism can hurt task-specific compression.

Reward Dynamics Analysis. To understand how our reward function evolves during training, we track the total reward and its two components (R_{ent} , R_{div}) across policy updates. As shown in Figure 3, the total reward increases consistently, indicating effective policy learning. The informativeness term R_{ent} rises rapidly in early stages and then saturates, reflecting that informative sample selection is quickly optimized. In contrast, the diversity term R_{div} grows more gradually, highlighting a shift in focus from informativeness to diversity as training progresses. This dynamic illustrates the complementary nature of the reward design, encouraging both discriminative and varied sample generation over time.



Figure 3: **Reward components over training steps.** The total reward (blue) increases steadily throughout training, driven by the entropy-based component R_{ent} (red) and the diversity-based penalty R_{div} (yellow). Notably, R_{ent} saturates early, while R_{div} continues to rise, indicating a shift in policy focus from informativeness to sample diversity as training progresses.

Additional Results in Supplementary. Due to space constraints, we include several extended experiments in the supplementary material. These include (1) cross-architecture evaluation on ImageNet-1K, which demonstrates the robustness of our method across different backbone networks; (2) an ablation study isolating the effects of the entropy-based reward (R_{ent}) and the diversity-based reward (R_{div}), showing that both components contribute positively to performance, though their combination yields diminishing returns due to partial redundancy; and (3) visualizations of generated samples that qualitatively reflect the trade-off between fidelity and informativeness. All code and implementation details are also provided in the supplement for reproducibility.

Discussion

This work takes a first step toward bridging generative modeling and dataset distillation by introducing a reward-driven formulation over the diffusion sampling process. While prior approaches often rely on handcrafted objectives or direct optimization of synthetic data, our method shows that reinforcement learning can provide a principled and flexible mechanism for exploring informative regions of the sample space.

Limitations. However, our framework also presents several limitations. First, incorporating reinforcement learning—though conceptually appealing—introduces training instability. Although GRPO offers a lightweight and gradient-regularized alternative to value-based methods, it still requires careful tuning of sampling frequency, reward scaling, and update schedules to achieve consistent convergence. Second, while our policy successfully shifts the generative behavior from reconstruction toward compression, it does so by modifying the sampling trajectory rather than the underlying diffusion model itself. The denoising backbone

Table 1: **Comparison of distillation performance across multiple methods, architectures, and datasets.** We report top-1 classification accuracy (%) on Nette and Woof subsets of the ImageNet-1K dataset under varying architectures (ConvNet-6, ResNetAP-10, ResNet-18) and image-per-class (IPC) budgets (10, 50, 100). COD (Ours) consistently achieves competitive or superior performance across settings, particularly under low IPC (e.g., 10), demonstrating its advantage in generating informative and compressed synthetic datasets. *Full* denotes training on the complete original dataset and serves as an upper bound.

Subset Architecture IPC	Nette									Woof								
	ConvNet-6			ResNetAP-10			ResNet-18			ConvNet-6			ResNetAP-10			ResNet-18		
	10	50	100	10	50	100	10	50	100	10	50	100	10	50	100	10	50	100
Random	46.0 ± 0.5	71.8 ± 1.2	79.9 ± 0.8	54.2 ± 1.2	77.3 ± 1.0	81.1 ± 0.6	55.8 ± 1.0	75.8 ± 1.1	82.0 ± 0.4	24.3 ± 1.1	41.3 ± 0.6	52.2 ± 0.4	29.4 ± 0.8	47.2 ± 1.3	59.2 ± 0.9	27.7 ± 0.9	47.9 ± 1.8	61.5 ± 1.3
DM	49.8 ± 1.1	70.3 ± 0.8	78.5 ± 0.8	60.2 ± 0.7	76.7 ± 1.0	80.9 ± 0.7	60.9 ± 0.7	75.0 ± 1.0	81.5 ± 0.4	26.9 ± 1.2	43.8 ± 1.1	50.1 ± 0.9	29.8 ± 1.0	47.1 ± 1.1	56.4 ± 0.8	30.2 ± 0.6	46.2 ± 0.6	60.2 ± 1.0
IDC-1	48.2 ± 1.2	72.4 ± 0.7	80.6 ± 1.1	60.4 ± 0.6	77.4 ± 0.7	81.5 ± 1.2	61.0 ± 0.8	77.8 ± 0.7	81.7 ± 0.8	33.3 ± 1.1	42.6 ± 0.9	51.0 ± 1.1	38.5 ± 0.7	48.3 ± 1.0	56.1 ± 0.9	36.7 ± 0.8	48.3 ± 0.8	57.7 ± 0.8
DiT	56.2 ± 1.3	73.3 ± 0.9	78.2 ± 0.3	62.8 ± 0.8	76.9 ± 0.5	80.1 ± 1.1	62.5 ± 0.9	75.2 ± 0.9	77.8 ± 0.6	32.3 ± 0.8	46.5 ± 0.8	53.4 ± 0.3	34.7 ± 0.5	49.3 ± 0.2	58.3 ± 0.8	34.7 ± 0.4	50.1 ± 0.5	58.9 ± 1.3
Minimax	58.2 ± 0.9	76.6 ± 0.2	81.1 ± 0.3	63.2 ± 1.0	78.2 ± 0.7	81.3 ± 0.9	64.9 ± 0.6	78.1 ± 0.6	81.3 ± 0.7	33.5 ± 1.4	50.7 ± 1.8	57.1 ± 1.9	39.2 ± 1.3	56.3 ± 1.0	64.5 ± 0.2	37.6 ± 0.9	57.1 ± 0.6	65.7 ± 0.4
COD (Ours)	59.2 ± 0.9	74.8 ± 0.5	78.8 ± 1.6	62.6 ± 0.3	78.2 ± 0.8	81.0 ± 1.0	63.8 ± 1.2	78.0 ± 0.9	81.6 ± 1.3	36.0 ± 0.9	47.8 ± 0.8	55.2 ± 1.5	41.6 ± 1.3	58.2 ± 0.8	65.6 ± 0.6	44.0 ± 1.8	59.2 ± 1.0	65.4 ± 0.9
Full	94.3 ± 0.5	94.3 ± 0.5	94.3 ± 0.5	94.6 ± 0.5	94.6 ± 0.5	94.6 ± 0.5	95.3 ± 0.6	95.3 ± 0.6	95.3 ± 0.6	85.9 ± 0.4	85.9 ± 0.4	85.9 ± 0.4	87.2 ± 0.6	87.2 ± 0.6	87.2 ± 0.6	89.0 ± 0.6	89.0 ± 0.6	89.0 ± 0.6

Table 2: Top-1 accuracy (%) on ImageNet-1K under different distillation methods with IPC = 10 and 50. COD achieves the highest accuracy when IPC=50, outperforming optimization-based (SRe²L, G-VBSM, RDED) and generative (DiT, Minimax) baselines. This demonstrates the effectiveness of reward-driven sampling in scaling dataset distillation to challenging large-scale benchmarks.

Dataset	IPC	SRe ² L	G-VBSM	RDED	DiT	Minimax	COD
ImageNet-1K	10	21.3 ± 0.6	31.4 ± 0.5	42.0 ± 0.1	39.6 ± 0.4	44.3 ± 0.5	44.0 ± 0.6
	50	46.8 ± 0.2	51.8 ± 0.4	56.5 ± 0.1	52.9 ± 0.6	58.6 ± 0.3	59.4 ± 0.4

remains trained to match the natural data distribution, and therefore retains an inherent bias toward data fidelity. As a result, the full potential of compression-oriented generation is still constrained by the original training objective of the generative model.

These limitations point to promising future directions, such as integrating downstream utility signals into the training of the generative model itself, or developing more stable and expressive learning frameworks beyond policy optimization to further improve the quality and utility of distilled samples.

Future Work. More broadly, we view *Compression-Oriented Distillation* as a paradigm shift for the dataset distillation community. Rather than treating generative models as static decoders of the original dataset, we advocate for a dynamic, policy-guided generation process in which synthetic data is optimized for task-specific utility. Our framework—based on reinforcement learning and built upon a transformer-based diffusion backbone—demonstrates that modern generative architectures can be harnessed not just for realism, but for strategic, goal-aware data construction.

We believe this direction opens up a rich avenue for fu-

ture research: leveraging increasingly powerful generative models, especially diffusion and transformer-based architectures, not merely as sample generators, but as active agents in data compression, selection, and synthesis. As foundation models continue to scale in capacity and generality, coupling them with task-aware decision-making mechanisms may fundamentally redefine how we construct and optimize training datasets across domains.

Conclusion

We introduced *Compression-Oriented Distillation* (COD), a reinforcement learning framework that guides diffusion models to generate informative and compact synthetic data for dataset distillation. By shifting the objective from reconstruction to compression, our method departs from static denoising and instead learns a dynamic sampling policy optimized for downstream utility. Through entropy-driven and diversity-driven rewards, our approach enables principled control over generative trajectories without modifying the diffusion training objective. This work bridges generative modeling and data distillation, paving the way for future research that further integrates task-aware objectives with advanced generative architectures.

References

- Bachem, O.; Lucic, M.; and Krause, A. 2017. Practical coreset constructions for machine learning. *arXiv preprint arXiv:1703.06476*.
- Black, K.; Janner, M.; Du, Y.; Kostrikov, I.; and Levine, S. 2024. Training Diffusion Models with Reinforcement Learning. *arXiv:2305.13301*.
- Cazenavette, G.; Wang, T.; Torralba, A.; Efros, A. A.; and Zhu, J. 2022. Dataset Distillation by Matching Training Trajectories. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 10708–10717.
- Chen, M.; Du, J.; Huang, B.; Wang, Y.; Zhang, X.; and Wang, W. 2025. Influence-guided diffusion for dataset distillation. In *The Thirteenth International Conference on Learning Representations*.
- Chen, Y.; Welling, M.; and Smola, A. 2012. Super-samples from kernel herding. *arXiv preprint arXiv:1203.3472*.
- Cui, J.; Wang, R.; Si, S.; and Hsieh, C. 2023. Scaling Up Dataset Distillation to ImageNet-1K with Constant Memory. In *Proc. Int. Conf. Mach. Learn. (ICML)*, 6565–6590.
- Du, J.; Jiang, Y.; Tan, V. Y. F.; Zhou, J. T.; and Li, H. 2023. Minimizing the Accumulated Trajectory Error to Improve Dataset Distillation. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 3749–3758.
- Du, J.; Shi, Q.; and Zhou, J. T. 2023. Sequential Subset Matching for Dataset Distillation. In *Adv. Neural Inf. Process. Syst. (NeurIPS)*.
- Fan, Y.; Watkins, O.; Du, Y.; Liu, H.; Ryu, M.; Boutilier, C.; Abbeel, P.; Ghavamzadeh, M.; Lee, K.; and Lee, K. 2023. Dpok: Reinforcement learning for fine-tuning text-to-image diffusion models. *Advances in Neural Information Processing Systems*, 36: 79858–79885.
- Gu, J.; Vahidian, S.; Kungurtsev, V.; Wang, H.; Jiang, W.; You, Y.; and Chen, Y. 2024. Efficient dataset distillation via minimax diffusion. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 15793–15803.
- Guo, D.; Yang, D.; Zhang, H.; Song, J.; Zhang, R.; Xu, R.; Zhu, Q.; Ma, S.; Wang, P.; Bi, X.; et al. 2025. Deepseek-r1: Incentivizing reasoning capability in llms via reinforcement learning. *arXiv preprint arXiv:2501.12948*.
- Har-Peled, S.; and Kushal, A. 2005. Smaller coresets for k-median and k-means clustering. In *Proceedings of the twenty-first annual symposium on Computational geometry*, 126–134.
- He, K.; Zhang, X.; Ren, S.; and Sun, J. 2016. Deep Residual Learning for Image Recognition. In *Proc. IEEE/CVF Conf. Comput. Vis. Pattern Recognit. (CVPR)*, 770–778.
- Ho, J.; Jain, A.; and Abbeel, P. 2020. Denoising diffusion probabilistic models. *Advances in neural information processing systems*, 33: 6840–6851.
- Kim, J.-H.; Kim, J.; Oh, S. J.; Yun, S.; Song, H.; Jeong, J.; Ha, J.-W.; and Song, H. O. 2022. Dataset condensation via efficient synthetic-data parameterization. In *International Conference on Machine Learning*, 11102–11118. PMLR.
- Lei, S.; and Tao, D. 2024. A Comprehensive Survey of Dataset Distillation. *IEEE Trans. Pattern Anal. Mach. Intell.*, 46(1): 17–32.
- Nguyen, T.; Chen, Z.; and Lee, J. 2021. Dataset Meta-Learning from Kernel Ridge-Regression. In *Proc. Int. Conf. Learn. Represent. (ICLR)*.
- Ouyang, L.; Wu, J.; Jiang, X.; Almeida, D.; Wainwright, C.; Mishkin, P.; Zhang, C.; Agarwal, S.; Slama, K.; Ray, A.; et al. 2022. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35: 27730–27744.
- Paul, M.; Ganguli, S.; and Dziugaite, G. K. 2021. Deep Learning on a Data Diet: Finding Important Examples Early in Training. In *Adv. Neural Inf. Process. Syst. (NeurIPS)*, 20596–20607.
- Peebles, W.; and Xie, S. 2023. Scalable diffusion models with transformers. In *Proceedings of the IEEE/CVF international conference on computer vision*, 4195–4205.
- Russakovsky, O.; Deng, J.; Su, H.; Krause, J.; Satheesh, S.; Ma, S.; Huang, Z.; Karpathy, A.; Khosla, A.; Bernstein, M.; Berg, A. C.; and Fei-Fei, L. 2015. ImageNet Large Scale Visual Recognition Challenge. *arXiv:1409.0575*.
- Sachdeva, N.; and McAuley, J. J. 2023. Data Distillation: A Survey. *Trans. Mach. Learn. Res.*
- Schulman, J.; Wolski, F.; Dhariwal, P.; Radford, A.; and Klimov, O. 2017. Proximal policy optimization algorithms. *arXiv preprint arXiv:1707.06347*.
- Sener, O.; and Savarese, S. 2017. Active learning for convolutional neural networks: A core-set approach. *arXiv preprint arXiv:1708.00489*.
- Shao, S.; Yin, Z.; Zhou, M.; Zhang, X.; and Shen, Z. 2024. Generalized large-scale data condensation via various backbone and statistical matching. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 16709–16718.
- Song, J.; Meng, C.; and Ermon, S. 2022. Denoising Diffusion Implicit Models. *arXiv:2010.02502*.
- Sun, P.; Shi, B.; Yu, D.; and Lin, T. 2024. On the diversity and realism of distilled dataset: An efficient dataset distillation paradigm. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*, 9390–9399.
- Sun, Q.; Jiang, Z.; Zhao, H.; and He, K. 2025. Is Noise Conditioning Necessary for Denoising Generative Models? *arXiv preprint arXiv:2502.13129*.
- Tsang, I. W.; Kwok, J. T.; Cheung, P.-M.; and Cristianini, N. 2005. Core vector machines: Fast SVM training on very large data sets. *Journal of Machine Learning Research*, 6(4).
- Wang, T.; Zhu, J.; Torralba, A.; and Efros, A. A. 2018. Dataset Distillation. *arXiv preprint arXiv:1811.10959*.
- Xiao, L.; and He, Y. 2024. Are large-scale soft labels necessary for large-scale dataset distillation? *arXiv preprint arXiv:2410.15919*.
- Yin, Z.; Xing, E. P.; and Shen, Z. 2023. Squeeze, Recover and Relabel: Dataset Condensation at ImageNet Scale From A New Perspective. In *Adv. Neural Inf. Process. Syst. (NeurIPS)*.

Zhao, B.; and Bilen, H. 2023. Dataset Condensation with Distribution Matching. In *Proc. IEEE/CVF Winter Conf. Appl. Comput. Vis. (WACV)*, 6503–6512.

Zhao, B.; Mopuri, K. R.; and Bilen, H. 2021. Dataset Condensation with Gradient Matching. In *Proc. Int. Conf. Learn. Represent. (ICLR)*.

Zhou, M.; Yin, Z.; Shao, S.; and Shen, Z. 2024. Self-supervised Dataset Distillation: A Good Compression Is All You Need. *arXiv preprint arXiv:2404.07976*.

Reproducibility Checklist

1. General Paper Structure

- 1.1. Includes a conceptual outline and/or pseudocode description of AI methods introduced (yes/partial/no/NA) [yes](#)
- 1.2. Clearly delineates statements that are opinions, hypothesis, and speculation from objective facts and results (yes/no) [yes](#)
- 1.3. Provides well-marked pedagogical references for less-familiar readers to gain background necessary to replicate the paper (yes/no) [yes](#)

2. Theoretical Contributions

- 2.1. Does this paper make theoretical contributions? (yes/no) [yes](#)

If yes, please address the following points:

- 2.2. All assumptions and restrictions are stated clearly and formally (yes/partial/no) [yes](#)
- 2.3. All novel claims are stated formally (e.g., in theorem statements) (yes/partial/no) [yes](#)
- 2.4. Proofs of all novel claims are included (yes/partial/no) [yes](#)
- 2.5. Proof sketches or intuitions are given for complex and/or novel results (yes/partial/no) [yes](#)
- 2.6. Appropriate citations to theoretical tools used are given (yes/partial/no) [yes](#)
- 2.7. All theoretical claims are demonstrated empirically to hold (yes/partial/no/NA) [yes](#)
- 2.8. All experimental code used to eliminate or disprove claims is included (yes/no/NA) [yes](#)

3. Dataset Usage

- 3.1. Does this paper rely on one or more datasets? (yes/no) [yes](#)

If yes, please address the following points:

- 3.2. A motivation is given for why the experiments are conducted on the selected datasets (yes/partial/no/NA) [yes](#)
- 3.3. All novel datasets introduced in this paper are included in a data appendix (yes/partial/no/NA) [NA](#)
- 3.4. All novel datasets introduced in this paper will be made publicly available upon publication of the paper with a license that allows free usage for research purposes (yes/partial/no/NA) [NA](#)
- 3.5. All datasets drawn from the existing literature (potentially including authors' own previously published work) are accompanied by appropriate citations (yes/no/NA) [yes](#)
- 3.6. All datasets drawn from the existing literature (potentially including authors' own previously published work) are publicly available (yes/partial/no/NA) [yes](#)
- 3.7. All datasets that are not publicly available are described in detail, with explanation why publicly available alternatives are not scientifically satisfying (yes/partial/no/NA) [yes](#)

4. Computational Experiments

- 4.1. Does this paper include computational experiments? (yes/no) [yes](#)

If yes, please address the following points:

- 4.2. This paper states the number and range of values tried per (hyper-) parameter during development of the paper, along with the criterion used for selecting the final parameter setting (yes/partial/no/NA) [yes](#)
- 4.3. Any code required for pre-processing data is included in the appendix (yes/partial/no) [yes](#)
- 4.4. All source code required for conducting and analyzing the experiments is included in a code appendix (yes/partial/no) [yes](#)
- 4.5. All source code required for conducting and analyzing the experiments will be made publicly available upon publication of the paper with a license that allows free usage for research purposes (yes/partial/no) [yes](#)
- 4.6. All source code implementing new methods have comments detailing the implementation, with references to the paper where each step comes from (yes/partial/no) [yes](#)
- 4.7. If an algorithm depends on randomness, then the method used for setting seeds is described in a way sufficient to allow replication of results (yes/partial/no/NA) [yes](#)
- 4.8. This paper specifies the computing infrastructure

used for running experiments (hardware and software), including GPU/CPU models; amount of memory; operating system; names and versions of relevant software libraries and frameworks (yes/partial/no) [yes](#)

- 4.9. This paper formally describes evaluation metrics used and explains the motivation for choosing these metrics (yes/partial/no) [yes](#)
- 4.10. This paper states the number of algorithm runs used to compute each reported result (yes/no) [yes](#)
- 4.11. Analysis of experiments goes beyond single-dimensional summaries of performance (e.g., average; median) to include measures of variation, confidence, or other distributional information (yes/no) [yes](#)
- 4.12. The significance of any improvement or decrease in performance is judged using appropriate statistical tests (e.g., Wilcoxon signed-rank) (yes/partial/no) [yes](#)
- 4.13. This paper lists all final (hyper-)parameters used for each model/algorithm in the paper's experiments (yes/partial/no/NA) [yes](#)