arXiv:2603.22364v4 [cs.LG] 20 Oct 2026

MCLR: Improving Conditional Modeling via Inter-Class Likelihood-Ratio Maximization and Unifying Classifier-Free Guidance with Alignment Objectives

Xiang Li[1,‡], **Yixuan Jia**[1], **Xiao Li**[1], **Jeffrey A. Fessler**[1], **Rongrong Wang**[2], **Qing Qu**[1]

University of Michigan[1] · Michigan State University[2]

[‡] Corresponding author

Abstract

Diffusion models achieve strong generative performance but often rely on classifier-free guidance (CFG), an inference-time intervention that modifies the sampling trajectory. Since standard denoising score matching (DSM) should recover the target distribution when the score is learned accurately, the practical necessity of CFG suggests deficiencies in learned conditional scores. We identify weak inter-class separation as one such limitation: samples generated from the same noise under different labels can retain similar class-agnostic structure. To address this deficiency, we introduce **MCLR**, a principled fine-tuning objective that explicitly **m**aximizes inter-**c**lass **l**ikelihood-**r**atios, thereby strengthening class-specific structures under standard sampling. Across diverse models and datasets, MCLR substantially improves guidance-free conditional generation and narrows the gap to CFG. Theoretically, we prove that the CFG-guided score is the unique minimizer of a sample-adaptively weighted, ELBO-approximated MCLR objective, interpreting CFG as an implicit inference-time contrastive alignment procedure. We further compare MCLR with related contrastive objectives, including DPO, CCA, and CFM, clarifying their connections and key differences in an otherwise fragmented literature.

Keywords: Diffusion Models, Classifier-Free Guidance, Contrastive Fine-Tuning, Alignment Algorithms

Date: Originally released on arXiv in May 2026; revised in October 2026
Correspondence: forkobe@umich.edu
Resources: Code Repository

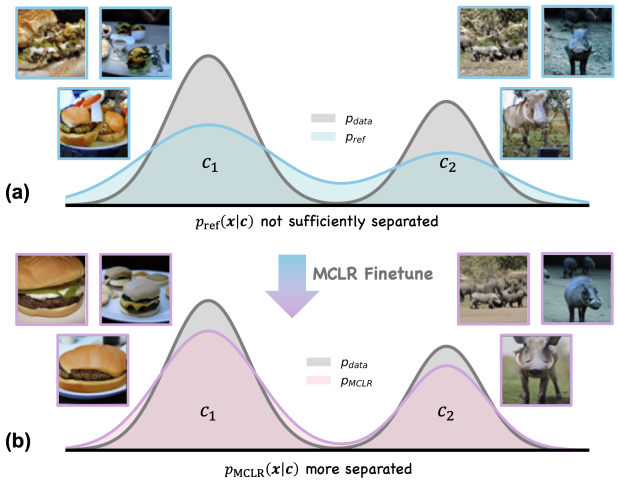


Figure 1: **Conceptual Illustration of MCLR.** (a) Samples generated from two classes using the same initial noises exhibit high visual similarity despite different conditioning labels, indicating insufficient separation of the learned conditional distributions. (b) MCLR mitigates this issue by encouraging inter-class separation, resulting in generations with more distinct class-specific features.

Contents

1 Introduction **3**

2 Preliminaries **4**
- 2.1 Basics of Diffusion Models . . . 4
- 2.2 Evidence Lower Bound for Diffusion Models . . . 5
- 2.3 Classifier-Free Guidance . . . 5
- 2.4 Direct Preference Optimization . . . 6

3 Motivation and Method **6**
- 3.1 A Practical Limitation of DSM: Weak Class Specificity . . . 6
- 3.2 Maximum Inter-Class Likelihood-Ratio Training . . . 6
- 3.3 Theoretical Analysis of MCLR . . . 8
- 3.4 Diffusion-Based MCLR: Approximating Log-Likelihood with ELBO . . . 9
- 3.5 CFG as a Contrastive Alignment Algorithm: A Mechanistic Interpretation . . . 11

4 Alternative Inter-Class Contrastive Objectives **12**
- 4.1 Adapting DPO for Improved Conditional Modeling . . . 12
- 4.2 Contrastive Flow Matching . . . 13

5 Related Work **13**

6 Experimental Results **14**

7 Limitations and Conclusions **16**

A Weak Inter-Class Separation in Diffusion Models **21**

B Theoretical Analysis of MCLR **22**
- B.1 Equivalent Forms of MCLR Objective . . . 22
- B.2 Distributional Analysis of MCLR . . . 22
- B.3 Distributional Analysis of Fine-tuning with MCLR . . . 26
- B.4 MCLR for Diffusion Models . . . 27
  - B.4.1 Parameterization-Free Analysis under a Pointwise Norm Bound . . . 29
  - B.4.2 Extension to Fine-Tuning from a Pretrained Score Model . . . 35
  - B.4.3 Connection to Classifier-Free Guidance . . . 36
  - B.4.4 Population-Level Simulation under the Mixture-Error Model . . . 39
- B.5 Comparison with Contrastive Flow Matching . . . 40

C Theoretical Analysis of CC-DPO **43**
- C.1 Basics of DPO . . . 43
- C.2 Improving Conditional Modeling with CC-DPO . . . 45

D Theoretical Analysis of CCA **47**

E Practical Implementation Details **49**
E.1 Approximating Log-Likelihood with ELBO . . . 49
E.2 Building Training Data from a Minibatch . . . 52
E.3 Overall Algorithm . . . 52
E.4 Hyperparameters . . . 53
E.5 Computational Efficiency . . . 53

F Additional Experimental Results **54**
F.1 Progressive Class Separation and the Fidelity-Diversity Trade-off. . . . 54
F.2 MCLR Outperforms Training-time Baselines. . . . 54
F.3 MCLR Achieves Comparable Performance as CFG. . . . 54

G Ablation Study **56**
G.1 Effects of DSM and KL Regularization . . . 56
G.2 Empirical Comparison with Contrastive Flow Matching . . . 57
G.3 Hyperparameter Sensitivity of CC-DPO . . . 58
G.4 Hyperparameter Sensitivity of CCA . . . 58

H Discussion on Related Works **58**
H.1 Visual Generation without Guidance . . . 58
H.2 Inference-Time Alignment via Guidance . . . 59

I Broader Impact **59**

1 Introduction

Diffusion models [HJA20; Son+21b; Kar+22; Lip+23] have become the dominant paradigm for high-fidelity generative modeling. These models generate samples by reversing a forward noising process using learned score functions, typically trained via denoising score matching [Vin11]. Although the reverse sampling process is theoretically guaranteed to recover the target distribution [And82] when the score functions are learned accurately, in practice, it often yields samples of noticeably inferior quality: conditional generation frequently appears visually incoherent or insufficiently faithful to the intended class or prompt [BN24]. In fact, nearly all high-quality generation by diffusion models relies heavily on classifier-free guidance (CFG) [HS22], an inference-time modification of the reverse sampling process that injects an additional guidance term. While CFG substantially improves sample quality, reducing FID scores by up to 75% as shown in seminal works [PX23; Yu+25], its empirical necessity exposes a gap between the theoretical optimality of DSM and its practical behavior. This raises a central question:

> ***What deficiency in DSM-trained conditional score functions does CFG compensate for, and can this deficiency be addressed through a principled training objective to enable high-quality conditional generation under standard sampling?***

Recent study [LWQ26] suggests that standard conditional models suffer from insufficient inter-class separation: generated samples are less distinguishable across classes than real data, indicating that class-dependent structures are not fully captured by diffusion models. Motivated by this insight, we propose **MCLR**, a principled alignment objective that explicitly prompts inter-class

separability by **M**aximizing the inter-**C**lass log-**L**ikelihood **R**atio. By encouraging the model to amplify density differences between a target class and other classes, MCLR strengthens class-specific structures in the learned score function. Empirically, models fine-tuned with MCLR exhibit CFG-like improvements under standard sampling, substantially improving guidance-free conditional generation and narrowing the gap to inference-time CFG.

Beyond the empirical benefits, we provide a theoretical result showing that the CFG-guided score coincides exactly with the optimizer of a weighted MCLR objective. This establishes a formal equivalence between CFG and alignment-based training objectives, revealing CFG as an implicit inference-time contrastive alignment algorithm.

The core principle of MCLR is to leverage inter-class contrastive signals to improve conditional modeling, which can also be realized with broader contrastive-based approaches. To assess the uniqueness of MCLR and for completeness, we analyze two alternative objectives: a conditional adaptation of DPO [Raf+23], termed Conditional Contrastive DPO (CC-DPO), and Contrastive Flow Matching (CFM) [Sto+25]. We show that CC-DPO shares the gamma-powered population optimum of CCA [Che+25b], whereas CFM induces a global mean-centered correction rather than MCLR's locally density-adaptive correction. Experiments show that MCLR consistently outperforms these training-time contrastive alternatives across diverse settings.

Summary of Contributions. Our main contributions are as follows:

- **A Principled Alignment Objective for Conditional Modeling.** We propose **MCLR**, a theoretically grounded fine-tuning objective that explicitly maximizes inter-class log-likelihood ratios to improve conditional generative modeling. Across diverse models and datasets, MCLR achieves substantial improvements in sample fidelity under standard reverse sampling, consistently outperforming training-time contrastive alternatives such as CC-DPO, CCA and CFM.
- **Theoretical Equivalence between CFG and Alignment Objectives.** We prove that the classifier-free guidance (CFG)–induced score coincides exactly with the optimizer of a weighted MCLR objective. This establishes a formal connection between CFG and alignment-based training, providing a mechanistic interpretation of CFG as an implicit inference-time alignment algorithm.
- **Understanding Contrastive Alternatives.** We show that CC-DPO and CCA share the same gamma-powered population optimum, whereas CFM rescales the conditional score and introduces a global mean-centered correction rather than MCLR's locally density-adaptive correction.

2 Preliminaries

2.1 Basics of Diffusion Models

Let $p_{\text{data}}(\boldsymbol{x})$ denote the ground-truth data distribution. Diffusion models construct a forward noising process that gradually perturbs p_{data} into a simple prior distribution using a stochastic differential equation (SDE):

$$\mathrm{d}\boldsymbol{x} = \boldsymbol{f}(\boldsymbol{x}, t)\,\mathrm{d}t + g(t)\,\mathrm{d}\boldsymbol{w}, \tag{1}$$

where $\boldsymbol{f}(\cdot, t)$ is the drift coefficient, $g(t)$ is the diffusion coefficient, and $\boldsymbol{w}$ denotes the standard Brownian motion. Let $p_t(\boldsymbol{x})$ be the marginal distribution of $\boldsymbol{x}(t)$, and $p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})$ the transition density

from $\boldsymbol{x}(0)$ to $\boldsymbol{x}(t)$. For sufficiently large T, $p_T(\boldsymbol{x})$ becomes indistinguishable from a tractable prior $\pi(\boldsymbol{x})$, e.g., an isotropic Gaussian. The SDE (1) admits a reverse-time probability-flow ODE [Son+21b]:

$$\mathrm{d}\boldsymbol{x} = \left[\boldsymbol{f}(\boldsymbol{x}, t) - \frac{1}{2} g^2(t) \nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x})\right] \mathrm{d}t. \tag{2}$$

Sampling from the reverse ODE requires the score function $\nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x})$, which can be approximated using a deep network $s_{\boldsymbol{\theta}}(\boldsymbol{x}, t)$ trained via the denoising score matching (DSM) objective:

$$\mathcal{J}_{\mathrm{DSM}}(\boldsymbol{\theta}; w(\cdot)) := \frac{1}{2} \int_0^T \mathbb{E}_{p(\boldsymbol{x}),\, p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})} \left[w(t) \left\| \nabla_{\boldsymbol{x}_t} \log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t) \right\|_2^2 \right] \mathrm{d}t, \tag{3}$$

where $w(t)$ is a positive weighting function. For conditional diffusion models, the score network takes a conditional embedding $\boldsymbol{c}$ as input, and the DSM objective naturally extends to the conditional setting by taking the expectation over class labels and class-conditional data distributions:

$$\mathcal{J}_{\mathrm{DSM}}(\boldsymbol{\theta}, \boldsymbol{c}; w(\cdot)) := \frac{1}{2} \int_0^T \mathbb{E}_{\boldsymbol{c}, p(\boldsymbol{x}|\boldsymbol{c}), p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})} [w(t) \| \nabla_{\boldsymbol{x}_t} \log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c}) \|_2^2] \mathrm{d}t. \tag{4}$$

Accordingly, the conditional reverse ODE is:

$$\mathrm{d}\boldsymbol{x} = \left[f(\boldsymbol{x}, t) - \frac{1}{2} g^2(t) \nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x}|\boldsymbol{c})\right] \mathrm{d}t. \tag{5}$$

2.2 Evidence Lower Bound for Diffusion Models

Let $p_{\boldsymbol{\theta}}^{\mathrm{ode}}(\boldsymbol{x})$ denote the distributions induced by the reverse ODE (2). Theorem 2 of [Son+21a] shows that, under certain regularity conditions, the log-likelihood satisfies:

$$\underbrace{\mathbb{E}_{p(\boldsymbol{x})}[\log p_{\boldsymbol{\theta}}^{\mathrm{ode}}(\boldsymbol{x})]}_{\text{Maximum Likelihood Estimation}} = -\mathcal{J}_{\mathrm{DSM}}(\boldsymbol{\theta}; g^2(\cdot)) + C, \tag{6}$$

where C is a constant independent of $\boldsymbol{\theta}$. Equation (6) is analogous to the evidence lower bound (ELBO) in variational autoencoders, revealing a fundamental connection between DSM and maximum likelihood estimation (MLE), which enables likelihood-based training in diffusion models [Mar+23; Wal+24; Zhe+25].

2.3 Classifier-Free Guidance

Although the reverse ODE in (5) is theoretically guaranteed to sample from the target conditional distribution, its practical generation quality is often unsatisfactory. In practice, high-quality conditional generation requires modifying the standard reverse process with an additional guidance term, known as classifier-free guidance (CFG) [HS22], which leads to the perturbed reverse ODE:

$$\mathrm{d}\boldsymbol{x} = \left[\boldsymbol{f}(\boldsymbol{x}, t) - \frac{1}{2} g^2(t) (\nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x}|\boldsymbol{c}) + \gamma \underbrace{(\nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x}|\boldsymbol{c}) - \nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x}))}_{\text{CFG guidance}}\right] \mathrm{d}t, \tag{7}$$

where γ controls the guidance strength. Intuitively, CFG sharpens the class- or conditional-specific structure by amplifying the difference between conditional and unconditional scores. However, the CFG-perturbed reverse sampling process (7) does not, in general, correspond to any known

forward process [BN24]. Despite recent progress towards understanding CFG [Wu+24; Chi+24; BN24; Pav+25; LWQ26; LJ25; JSG25; Yan+26; Ven+26], its underlying mechanisms remain only partially understood. This post-hoc modification of the sampling procedure motivates the search for principled alternatives that reproduce CFG-like effects while preserving theoretical consistency, which may, in turn, shed light on the mechanisms of CFG itself.

2.4 Direct Preference Optimization

Direct Preference Optimization (DPO) aligns a reference (base) model p_{ref} by increasing the relative likelihood of preferred samples over non-preferred samples [Raf+23; Wal+24]. Under the Bradley–Terry model [BT52],

$$p(\boldsymbol{x}_w \succ \boldsymbol{x}_l|\boldsymbol{c}) = \text{Sigmoid}\big(r(\boldsymbol{x}_w|\boldsymbol{c}) - r(\boldsymbol{x}_l|\boldsymbol{c})\big), \tag{8}$$

where $\boldsymbol{x}_w$ and $\boldsymbol{x}_l$ denote preferred and non-preferred samples, and $\text{Sigmoid}(x) := 1/(1+\exp(-x))$. DPO parameterizes the reward as $r_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}) = \beta \log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}) - \beta \log p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})$, which leads to the objective

$$\max_{\boldsymbol{\theta}} \ \mathbb{E}_{(\boldsymbol{c},\boldsymbol{x}_w,\boldsymbol{x}_l)\sim S}\left[\log \text{Sigmoid}\left(\beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}_w|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}_w|\boldsymbol{c})} - \beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}_l|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}_l|\boldsymbol{c})}\right)\right], \tag{9}$$

where p_{ref} denotes the pretrained base (reference) model, β controls the strength of the KL regularization between $p_{\boldsymbol{\theta}}$ and p_{ref}, and S represents the preference dataset (see a more detailed introduction in Section C.1). We later adapt this preference-style objective to conditional generation by treating target-class samples as preferred and mismatched-class samples as non-preferred.

3 Motivation and Method

This section first identifies weak class specificity as a practical limitation of standard DSM training (section 3.1). To remedy this issue, we introduce MCLR (Sections 3.2 to 3.4), and establish its connection to CFG (Section 3.5).

3.1 A Practical Limitation of DSM: Weak Class Specificity

In theory, if the score functions are learned accurately, standard reverse diffusion sampling should produce samples from the target conditional distribution. In practice, however, conditional sampling often exhibits weak class-specific structure [LWQ26]. A common failure mode is that samples generated from the same initial noise but different class conditions share similar global layouts, while class-discriminative features are attenuated or missing, as shown in Figure 1(a) and Figure 2. This suggests that the learned conditional distributions are insufficiently distinguishable from one another. We provide additional quantitative evidence based on pairwise inter-class FID, together with controlled qualitative comparisons, in Appendix A. Figure 1 illustrates this conceptually: the learned class-conditional distributions of a base model p_{ref} exhibit substantially less separation than the ground-truth data distribution p_{data}. This motivates modifying the standard DSM to more explicitly exploit the information carried by the conditioning variable.

3.2 Maximum Inter-Class Likelihood-Ratio Training

Motivated by the preceding discussion, we propose improving conditional modeling by explicitly encouraging inter-class separation in conditional models, as shown conceptually in Figure 1(b).

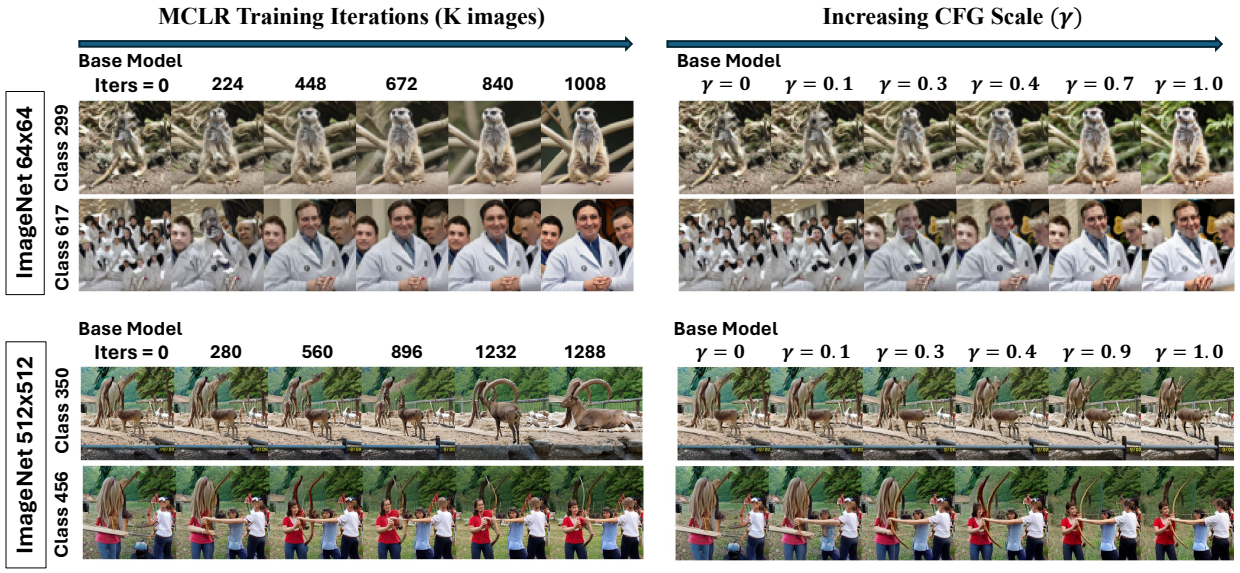


Figure 2: **Effects of MCLR vs. CFG.** We visualize the progressive emergence of class-specific structure under MCLR training and CFG. MCLR training induces effects analogous to increasing guidance strength in CFG, both substantially improving class-specific patterns. For each image block, samples from different classes are generated from the same initial noises.

Let $p_\theta(\boldsymbol{x}|\boldsymbol{c})$ denote the model's conditional distribution for class condition $\boldsymbol{c}$, where $\boldsymbol{\theta}$ is the model parameter. We consider the following objective:

$$\max_{\theta} \mathbb{E}_{\boldsymbol{c},p(\boldsymbol{x}|\boldsymbol{c})} \log p_\theta(\boldsymbol{x}|\boldsymbol{c}) + \frac{\eta}{2} \mathbb{E}_{\boldsymbol{c},\tilde{\boldsymbol{c}},\boldsymbol{x}\sim p(\cdot|\boldsymbol{c}),\boldsymbol{y}\sim p(\cdot|\tilde{\boldsymbol{c}})} \left[\log \frac{p_\theta(\boldsymbol{x}|\boldsymbol{c})}{p_\theta(\boldsymbol{x}|\tilde{\boldsymbol{c}})} + \log \frac{p_\theta(\boldsymbol{y}|\tilde{\boldsymbol{c}})}{p_\theta(\boldsymbol{y}|\boldsymbol{c})} \right], \tag{10}$$

where $\boldsymbol{c}$ and $\tilde{\boldsymbol{c}}$ are two randomly sampled classes, and η controls the regularization strength. Compared to standard MLE (equivalently, DSM under appropriate weighting) in (6), eq. (10) introduces an additional regularizer, which we refer to as **M**aximum Inter-**C**lass **L**ikelihood-**R**atio (MCLR).

Specifically, given a sample $\boldsymbol{x}$ (or $\boldsymbol{y}$) drawn from class $\boldsymbol{c}$ (or $\tilde{\boldsymbol{c}}$), MCLR explicitly encourages its corresponding log-likelihood under the true class $\log p_\theta(\boldsymbol{x}|\boldsymbol{c})$ (or $\log p_\theta(\boldsymbol{y}|\tilde{\boldsymbol{c}})$) to be higher relative to its log-likelihood $\log p_\theta(\boldsymbol{x}|\tilde{\boldsymbol{c}})$ (or $\log p_\theta(\boldsymbol{y}|\boldsymbol{c})$) under a mismatched class $\tilde{\boldsymbol{c}}$ (or $\boldsymbol{c}$). In doing so, MCLR drives the model to increase the inter-class likelihood ratio, thereby pushing $p_\theta(\boldsymbol{x}|\boldsymbol{c})$ to concentrate more probability mass in regions where the true class is favored over competing classes. Such regions typically correspond to samples with more pronounced class-specific features.

Intuitively, MCLR encourages the model to fully exploit the label-conditioned information. Standard conditional training simply feeds the conditional label $\boldsymbol{c}$ alongside the data to the model, optimizes the MLE (or DSM) objective, and relies on deep networks to automatically discover the conditional structure. If labels are under-exploited, the conditional distributions can collapse and become weakly distinguishable, yielding $p_\theta(\boldsymbol{x}|\boldsymbol{c}) \approx p_\theta(\boldsymbol{x}|\tilde{\boldsymbol{c}})$ regardless of the true class of $\boldsymbol{x}$. Maximizing log-likelihood ratio $\log \frac{p_\theta(\boldsymbol{x}|\boldsymbol{c})}{p_\theta(\boldsymbol{x}|\tilde{\boldsymbol{c}})}$ penalizes this collapse and encourages the model to discriminate between different conditions (classes) by leveraging the information encoded in the labels.

Applying the linearity of expectation to (10) and simplifying leads to the following equivalent

form that we use throughout, unless otherwise stated:

$$\max_{\theta} \; \mathbb{E}_{c,\, p(x|c)}\big[\log p_\theta(x|c)\big] + \eta \underbrace{\mathbb{E}_{c,\, \tilde{c},\, x\sim p(\cdot|c),\, y\sim p(\cdot|\tilde{c})}\left[\log \frac{p_\theta(x|c)}{p_\theta(y|c)}\right]}_{\text{MCLR Regularization}}. \tag{11}$$

Fine-tuning with MCLR. When a pretrained (suboptimal) model $p_{\text{ref}}(x)$ that lacks class specificity is available, we can fine-tune it using MCLR in combination with KL regularization:

$$\max_{\theta} -\mathbb{E}_c\left[D_{\text{KL}}\big(p_{\text{ref}}(x|c)\|p_\theta(x|c)\big)\right] + \eta\, \mathbb{E}_{c,\, \tilde{c}, x\sim p(\cdot|c), y\sim p(\cdot|\tilde{c})}\left[\log \frac{p_\theta(x|c)}{p_\theta(y|c)}\right]. \tag{12}$$

We next analyze the optimal solution of (11) and (12), providing insight into MCLR's effects.

3.3 Theoretical Analysis of MCLR

We now analyze the optimal distribution induced by MCLR. Define:

$$h(x|c) := p(x|c) + \eta\,\big(p(x|c) - p(x)\big). \tag{13}$$

Theorem 1. *Suppose $h(x|c)$ is supported on a bounded set K, and restrict the model class to densities satisfying $p_\theta(x|c) \geq \delta > 0$ on K. These conditions are mainly technical: bounded support is natural for image data with bounded pixel values, and the positive density floor ensures that the log-likelihood objective is well-defined. Then the optimal solution to (11) is:*

$$p_{\theta^*}(x|c) \;=\; \begin{cases} \max\left\{\frac{h(x|c)}{Z(c)}, \delta\right\}, & x \in K, \\ 0, & x \notin K, \end{cases} \tag{14}$$

where $Z(c)$ is the normalizing constant.

The formal assumptions and proof are provided in Section B.2. Intuitively, the optimal conditional distribution $p_{\theta^*}(x|c)$ induced by MCLR is proportional to $h(x|c)$ whenever $h(x|c) > Z(c)\delta$, while being clipped to the floor δ elsewhere. In the limit $\delta \to 0$, the floor disappears and the optimal distribution approaches:

$$p_{\theta^*}(x|c) = \frac{h^+(x|c)}{\int_x h^+(x|c)dx}, \tag{15}$$

where $h^+(x|c) := \max\{h(x|c), 0\}$, such that the negative part of $h(x|c)$ is truncated to zero and then normalized to form a valid distribution. One can interpret (14) and (15) as the "sum-of-difference" distribution: MCLR reshapes $p(x|c)$ according to the difference between $p(x|c)$ and $p(x)$, such that the regions where $p(x|c) > p(x)$ (i.e., samples with strong class-specific features) are amplified, while regions where $p(x|c) < p(x)$ (i.e., ambiguous samples lying near class boundaries) are suppressed.

For the fine-tuning objective in (12), it is easy to show that the optimal solution admits the same form as (14), but in this case $h(x|c) = p_{\text{ref}}(x|c) + \eta\,(p(x|c) - p(x))$. This directly leads to the following corollary (Section B.3 provides the proof).

Corollary 1. *If the base model p_{ref} satisfies the mixture error model:*

$$p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c}) = (1-\eta)\, p(\boldsymbol{x}|\boldsymbol{c}) + \eta\, p(\boldsymbol{x}), \tag{16}$$

where $\eta \in [0,1]$, then fine-tuning $p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})$ with MCLR regularization (12) recovers the ground-truth conditional distribution $p(\boldsymbol{x}|\boldsymbol{c})$.

The mixture error model (16) posits that the base model suffers from cross-class "leakage"; specifically, the learned conditional density $p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})$ is a convex combination of the ground-truth conditional distribution $p(\boldsymbol{x}|\boldsymbol{c})$ and the unconditional counterpart $p(\boldsymbol{x})$, which corresponds to a weighted average over all class-conditional distributions. MCLR counteracts the leakage by adding the contrastive density difference $p(\boldsymbol{x}|\boldsymbol{c}) - p(\boldsymbol{x})$ to the base model, thereby suppressing regions dominated by competing classes and strengthening class-specific regions. Although the error structure of practical diffusion models is likely more complex than this simple mixture model, this model provides a useful intuition for why MCLR can mitigate insufficient class separation.

3.4 Diffusion-Based MCLR: Approximating Log-Likelihood with ELBO

The objective above is written in terms of log-likelihoods. This is directly applicable to autoregressive models [Tia+24], but exact likelihood computation for diffusion models is too expensive for training. We therefore approximate log-likelihood differences using the diffusion ELBO (6). This yields an ELBO-approximated version of MCLR (69), whose full derivation is given in Section E.1.

Denoising Interpretation of MCLR. According to the equivalence between score function and optimal MMSE denoiser (see Section E.1 for details), the score function can be parameterized as: $\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}, t, \boldsymbol{c}) = \frac{\mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x};\sigma(t),\boldsymbol{c})-\boldsymbol{x}}{\sigma^2(t)}$, where $\sigma(t)$ is the standard deviation of additive noise at time t. With a customized noise-time distribution $p(t)$ and weighting function $w(t)$, the MCLR regularization becomes

$$\mathbb{E}_{\substack{\boldsymbol{c},\tilde{\boldsymbol{c}},\, t\sim p(t) \\ \boldsymbol{x}\sim p(\cdot|\boldsymbol{c}),\, \boldsymbol{x}_t\sim p_{0t}(\cdot|\boldsymbol{x}) \\ \boldsymbol{y}\sim p(\cdot|\tilde{\boldsymbol{c}}),\, \boldsymbol{y}_t\sim p_{0t}(\cdot|\boldsymbol{y})}} \left[w(t) \Big(\|\boldsymbol{x} - \mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x}_t; \sigma(t), \boldsymbol{c})\|_2^2 - \|\boldsymbol{y} - \mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{y}_t; \sigma(t), \boldsymbol{c})\|_2^2 \Big) \right]. \tag{17}$$

Eq. (17) provides a denoising-perspective interpretation of MCLR: it encourages the target-condition denoiser $\mathcal{D}(\cdot; \sigma(t), \boldsymbol{c})$ to achieve lower reconstruction error on samples of the target condition $\boldsymbol{c}$ than on samples from a mismatched condition $\tilde{\boldsymbol{c}}$. We use this formulation in practical implementation.

Population-Level Analysis of Diffusion-Based MCLR. Although likelihood-based objectives are commonly adapted to score-based diffusion models through ELBO approximations [Wal+24; Zhe+25; Xue+25], the resulting diffusion objective need not preserve the likelihood-level optimum. In particular, the population-optimal score of diffusion-based MCLR does not generally coincide with the score associated with the forward-noised version of the likelihood-based solution in (14). We therefore directly characterize the population-optimal score field induced by the diffusion-based MCLR objective (69) by treating the score function $\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c})$ as the optimization variable, without imposing any additional assumptions on the model parameterization. However, we do impose the following regularity constraint on the score function, which bounds its maximum norm:

$$\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c})\|_2 \le M, \tag{18}$$

where M is a positive constant. This constraint is required for the diffusion-based MCLR objective (69) and (70) to be well-posed so that a pointwise minimizer exists. Consider a $\boldsymbol{x}_t$ such that:

$$p_t(\boldsymbol{x}_t) \gg p_t(\boldsymbol{x}_t|\boldsymbol{c}), \tag{19}$$

i.e., $\boldsymbol{x}_t$ has a much larger density under the unconditional distribution compared to that under the conditional distribution. At such $\boldsymbol{x}_t$, the negative term in (69) and (70) dominates the training objective, and the objective value can be pushed arbitrarily small by simply increasing the norm of $\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c})$ or $\mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c})$ to be arbitrarily large. Intuitively, MCLR relies on the positive term to prevent the degenerate behavior caused by the negative term, but at $\boldsymbol{x}_t$ where $p_t(\boldsymbol{x}_t) \gg p_t(\boldsymbol{x}_t|\boldsymbol{c})$, the regularization effect of the positive term becomes too weak and the MCLR loss can be pushed arbitrarily small, while the score function or denoiser fails to learn any useful class-specific structure. Therefore, we impose a pointwise bound on the score magnitude. This constraint also reflects the implicit magnitude control induced in practice by finite network capacity and regularization. Now we are ready to present the main theorem.

Define:

$$h_t(\boldsymbol{x}_t|\boldsymbol{c}) := p_t(\boldsymbol{x}_t|\boldsymbol{c}) + \eta \left(p_t(\boldsymbol{x}_t|\boldsymbol{c}) - p_t(\boldsymbol{x}_t)\right), \tag{20}$$

which can be interpreted as a noisy version of the "sum-of-difference" contrastive function (13). Then we have the following theorem.

Theorem 2. *Consider the population diffusion-based MCLR objective over score fields satisfying* $\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c})\|_2 \leq M$. *For any* $(\boldsymbol{x}_t, t, \boldsymbol{c})$ *such that* $\nabla_{\boldsymbol{x}_t} h_t(\boldsymbol{x}_t|\boldsymbol{c}) \neq \boldsymbol{0}$, *the pointwise optimal score is*

$$\boldsymbol{s}_{\boldsymbol{\theta}^*}(\boldsymbol{x}_t, t, \boldsymbol{c}) = \begin{cases} \nabla_{\boldsymbol{x}_t} \log h_t(\boldsymbol{x}_t|\boldsymbol{c}), & h_t(\boldsymbol{x}_t \mid \boldsymbol{c}) > 0 \text{ and } \|\nabla_{\boldsymbol{x}_t} \log h_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2 \leq M, \\ M \dfrac{\nabla_{\boldsymbol{x}_t} h_t(\boldsymbol{x}_t|\boldsymbol{c})}{\|\nabla_{\boldsymbol{x}_t} h_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2}, & \text{otherwise}. \end{cases} \tag{21}$$

Thus, the diffusion-based MCLR score is locally governed by the noise-mollified sum-of-difference function h_t. In regions where $h_t > 0$ (i.e., $\frac{p_t(\boldsymbol{x}_t|\boldsymbol{c})}{p_t(\boldsymbol{x}_t)} > \frac{\eta}{1+\eta}$), it follows $\nabla \log h_t$, up to magnitude clipping determined by M. Note that:

$$\nabla_{\boldsymbol{x}_t} \log h_t(\boldsymbol{x}_t|\boldsymbol{c}) = \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}) + \alpha_t(\eta, \boldsymbol{x}_t)\big(\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}) - \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t)\big), \tag{22}$$

where $\alpha_t(\eta, \boldsymbol{x}_t) = \frac{1}{(1+\frac{1}{\eta})\frac{p_t(\boldsymbol{x}_t|\boldsymbol{c})}{p_t(\boldsymbol{x}_t)} - 1}$. Hence, in this region, diffusion-based MCLR takes the form of CFG with a sample-adaptive guidance strength $\alpha_t(\eta, \boldsymbol{x}_t)$, which is a monotonically decreasing function of the conditional-unconditional likelihood ratio $\frac{p_t(\boldsymbol{x}_t|\boldsymbol{c})}{p_t(\boldsymbol{x}_t)}$. Therefore, the guidance induced by MCLR becomes weaker when the ratio is large, because the current sample already provides strong evidence for the target condition, and becomes stronger as this relative evidence decreases.

When $h_t \leq 0$, the optimal field saturates at the norm bound and points along the steepest-ascent direction of h_t, thereby potentially steering samples toward regions with stronger relative evidence for the target condition.

The above discussion focuses on non-critical points. At critical points where $\nabla_{\boldsymbol{x}_t} h_t(\boldsymbol{x}_t|\boldsymbol{c}) = \boldsymbol{0}$, h_t provides no local directional information, so the MCLR objective may not identify a unique optimal direction. This degeneracy reflects the absence of local first-order gradient information rather than a failure of MCLR. In practice, the learned score values at the critical points are implicitly determined by the shared network parameterization and neighboring training samples. The complete

characterization at critical points, the proof and synthetic simulation are deferred to Section B.4. Importantly, as shown in Corollary 2, diffusion-based MCLR recovers the ground-truth conditional distribution under the mixture-error model (16).

3.5 CFG as a Contrastive Alignment Algorithm: A Mechanistic Interpretation

We now connect classifier-free guidance (CFG) to diffusion-based MCLR developed above. In particular, we show that the CFG-guided score in (7) arises as the unique minimizer of a sample-adaptive weighted, ELBO-approximated MCLR objective. This result provides a formal interpretation of CFG as an inference-time procedure that implicitly performs contrastive alignment between the conditional distribution and the unconditional mixture (the proof is provided in Section B.4.3).

Theorem 3. *For any time sampling distribution $p(t)$ and weighting function $w(t) > 0$, the CFG-guided score*

$$s_{cfg}(x_t, t, c) := \nabla_{x_t} \log p_t(x_t|c) + \eta \left(\nabla_{x_t} \log p_t(x_t|c) - \nabla_{x_t} \log p_t(x_t)\right)$$

is the unique minimizer of a sample-adaptive weighted ELBO-approximated MCLR objective:

$$\begin{aligned} s_{cfg}(\cdot) = \underset{s_\theta(\cdot)}{arg\,min}\ & \mathbb{E}_{c,t\sim p(t),\, x\sim p(x|c),\, x_t\sim p_{0t}(x_t|x)}\Big[w(t)\|\nabla_{x_t} \log p_{0t}(x_t|x) - s_\theta(x_t, t, c)\|_2^2\Big] \\ & + \eta\, \mathbb{E}_{\substack{c,\tilde{c},t\sim p(t),x\sim p(\cdot|c),y\sim p(\cdot|\tilde{c}) \\ p_{0t}(x_t|x),p_{0t}(y_t|y)}}\Bigg[w(t)\Big(\|\nabla_{x_t} \log p_{0t}(x_t|x) - s_\theta(x_t, t, c)\|_2^2 \\ & - \frac{p_t(y_t|c)}{p_t(y_t)}\|\nabla_{y_t} \log p_{0t}(y_t|y) - s_\theta(y_t, t, c)\|_2^2\Big)\Bigg], \end{aligned} \tag{23}$$

where $p_t(y_t) := \mathbb{E}_c[p_t(y_t|c)]$.

The regularization term in (23) has the same contrastive structure as the standard MCLR objective in (69) and its denoising form (17): it encourages the denoiser associated with the target condition c to achieve lower denoising error on samples from the target condition than on samples drawn from mismatched conditions. The key difference from standard MCLR is the sample-adaptive weight $\frac{p_t(y_t|c)}{p_t(y_t)}$. Since this weight is nonnegative, the inter-class contrastive nature of MCLR is preserved.

Intuitively, a larger weight indicates that a noisy sample y_t from a mismatched condition $\tilde{c}$ is more likely to be confused with samples from the target condition c. Hence, the adaptive weight makes the CFG-equivalent objective (23) more selective over negative samples: samples y_t from mismatched conditions that have high relative density under the target condition c receive larger weights, while samples that are already easily distinguishable from the target condition (those with small $\frac{p_t(y_t|c)}{p_t(y_t)}$) receive smaller weights. Thus, the adaptive weight can be interpreted as a form of hard-negative mining, emphasizing ambiguous or class-confusable samples that provide stronger contrastive signal.

This perspective clarifies the relationship between MCLR and CFG. Standard MCLR uses a uniform contrastive penalty over mismatched samples, whereas CFG corresponds to a sample-adaptive weighted version of MCLR. Therefore, CFG can be interpreted as an **inference-time** contrastive alignment procedure, while MCLR provides a **training-time** mechanism for partially internalizing this effect into the model. In practice, we implement the standard, uniformly weighted MCLR objective rather than the sample-adaptive version. Nevertheless, as shown empirically, standard MCLR induces similar qualitative effects and substantially narrows the gap between unguided sampling and CFG.

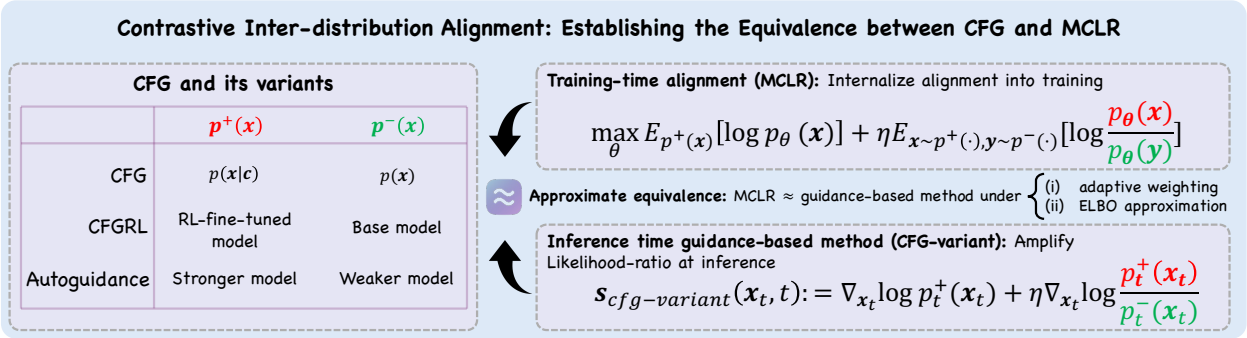


Figure 3: **A Unified Framework Connecting CFG Variants with Contrastive Alignment.** CFG-based methods can be interpreted as implicitly optimizing a likelihood-ratio–based contrastive alignment objective between two distributions $p^+(x)$ and $p^-(x)$ at inference time.

Understanding CFG-Variants through the Alignment Lens. Finally, Theorem 3 provides a unified perspective for interpreting CFG variants. By modifying the source distributions from which positive and negative samples are drawn in the contrastive objective, one recovers several existing guidance mechanisms as special cases.

For example in (23), if x is sampled from a stronger model and y is sampled from a weaker model, the resulting optimal solution recovers the score of Autoguidance [Kar+24a]. Similarly, if x is sampled from a reinforcement-learning fine-tuned model and y is sampled from the base model, the induced solution corresponds to the score used in CFGRL [Fra+25]. Thus, MCLR provides a unified alignment interpretation for a broad class of guidance algorithms. Figure 3 summarizes this unified framework and Section B.4.3 provides detailed generalization and discussion.

4 Alternative Inter-Class Contrastive Objectives

4.1 Adapting DPO for Improved Conditional Modeling

The core principle behind MCLR is to improve conditional modeling by encouraging the model to exploit class-dependent structures through inter-class contrast, which can also be instantiated via other contrastive objectives such as DPO. To study this connection, we adapt DPO to conditional generation by treating samples from the target class c as preferred (x_w) and samples from other randomly selected classes as non-preferred (x_l). This leads to the following objective:

$$\min_\theta \; -\mathbb{E}_{c,\tilde{c},\, x_w\sim p(\cdot|c),\, x_l\sim p(\cdot|\tilde{c})} \log \text{Sigmoid}\left(\beta \log \frac{p_\theta(x_w|c)}{p_{\text{ref}}(x_w|c)} - \beta \log \frac{p_\theta(x_l|c)}{p_{\text{ref}}(x_l|c)}\right). \tag{24}$$

We refer to (24) as Conditional Contrastive DPO (CC-DPO). Similar to MCLR, CC-DPO objective also admits a closed-form solution, as stated in the following theorem (see proof in Section C.2).

Theorem 4. *Under standard regularity conditions, the optimal solution to (24) is:*

$$p_{\theta^*}(x|c) = \frac{1}{\tilde{Z}(c)} p_{ref}(x|c) \left(\frac{p(x|c)}{p(x)}\right)^{\frac{1}{\beta}}, \tag{25}$$

where $\tilde{Z}(c) = \int_x p_{ref}(x|c) \left(\frac{p(x|c)}{p(x)}\right)^{\frac{1}{\beta}} dx$ is the normalizing constant. Moreover, this optimal solution coincides with that induced by Conditional Contrastive Alignment (CCA) [Che+25b].

Comparison between MCLR and CC-DPO. The optimal solution (25) is known as the *gamma-powered* distribution, which was initially conjectured to characterize the effect of classifier-free guidance (CFG) [HS22], but was later shown not to correspond to the true CFG dynamics [Kar+24a; BN24]. Although both methods amplify regions where $p(\boldsymbol{x}|\boldsymbol{c}) > p(\boldsymbol{x})$ and suppress regions where $p(\boldsymbol{x}|\boldsymbol{c}) < p(\boldsymbol{x})$, they do so through fundamentally different transformations: MCLR modifies the distribution additively through the density difference $p(\boldsymbol{x}|\boldsymbol{c}) - p(\boldsymbol{x})$, whereas CC-DPO reweights it multiplicatively through the density ratio $\left(\frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})}\right)^{\frac{1}{\beta}}$. The multiplicative form of CC-DPO can be overly aggressive. Consider a point $\tilde{\boldsymbol{x}}$ such that $p(\tilde{\boldsymbol{x}}) \approx 0$ while $p(\tilde{\boldsymbol{x}}|\boldsymbol{c}) > 0$, a situation that naturally arises when $\boldsymbol{c}$ are minority classes. In this case, the ratio $\left(\frac{p(\tilde{\boldsymbol{x}}|\boldsymbol{c})}{p(\tilde{\boldsymbol{x}})}\right)^{\frac{1}{\beta}}$ becomes ill-conditioned, strongly amplifying the density at $\tilde{\boldsymbol{x}}$, therefore potentially driving the learned conditional distribution toward degenerate or unstable solutions. By contrast, the optimal solution (14) induced by MCLR remains well-behaved. We include CC-DPO and CCA as contrastive baselines in our experiments.

4.2 Contrastive Flow Matching

Contrastive Flow Matching (CFM) [Sto+25] provides another way to introduce conditional contrast into diffusion training, but it induces a fundamentally different optimal score field compared to MCLR. Specifically, we show that its optimal score is:

$$\boldsymbol{s}_{\theta^*}(\boldsymbol{x}_t, t, \boldsymbol{c}) = (1+\eta)\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}) + \eta \frac{b_t}{a_t}(\boldsymbol{x}_t - \boldsymbol{\mu}_t), \tag{26}$$

where $\boldsymbol{\mu}_t := \mathbb{E}_{\boldsymbol{x}_t \sim p_t(\boldsymbol{x}_t)}[\boldsymbol{x}_t]$ is the unconditional mean, η is the CFM regularization strength and a_t, b_t are determined by the noise schedule. Thus, CFM rescales the conditional score and introduces a global mean-centered affine correction determined by the displacement between the current noisy sample and the unconditional mean. Unlike CFG and MCLR, the CFM correction is not explicitly adaptive to the local conditional–unconditional density geometry. Consequently, CFM can promote global class separation, but its correction need not vanish even in regions that are already strongly associated with the target condition. We provide the complete theoretical analysis in Section B.5. Under matched fine-tuning settings, MCLR consistently outperforms CFM (see Section G.2).

5 Related Work

Our work contributes to recent efforts [Che+25b; Che+25a; Tan+25] that seek to improve guidance-free conditional generation by modifying the training objective rather than applying classifier-free guidance (CFG) at inference time. The closest related method is Conditional Contrastive Alignment (CCA) [Che+25b], which learns the gamma-powered distribution via Noise Contrastive Estimation [GH10]. We show that MCLR outperforms CCA in diffusion models while achieving comparable performance in autoregressive settings, and further reveal a previously unexplored equivalence between CC-DPO and CCA.

Another related direction is Guidance-Free Training (GFT) [Che+25a; Tan+25], which aims to reproduce CFG-induced score functions through modified DSM objectives. While GFT mimics the functional form of CFG during training, MCLR approaches the problem from a likelihood-ratio perspective and reveals the contrastive structure implicit in CFG, offering a complementary mechanistic interpretation.

Several additional works [Yan+24; Lee+25; KMS24; YAA25] employ class-wise contrastive objectives to improve conditional generation, but lack a formal characterization of the underlying theoretical properties.

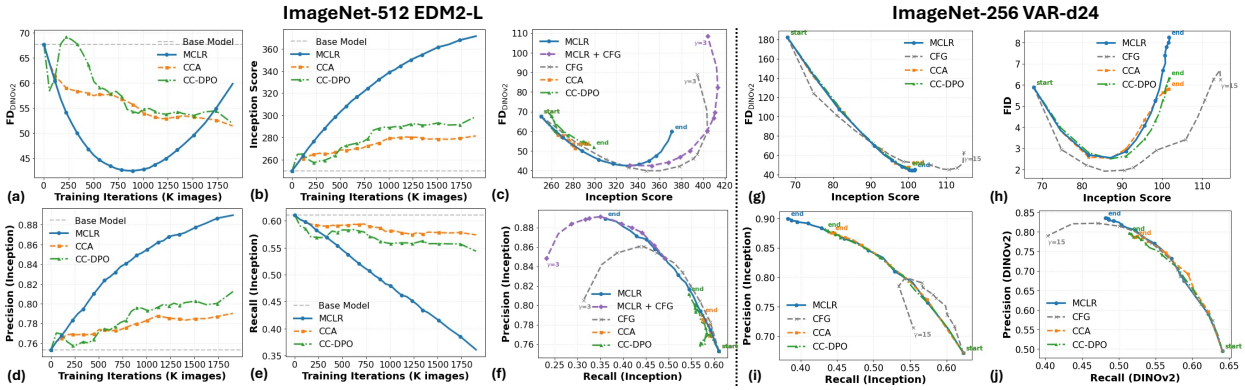


Figure 4: **Quantitative Results for EDM2-L and VAR-d24.** (a), (b), (d), and (e) show the evolution of FD, Inception Score, Precision, and Recall, respectively, as functions of training iterations. (c), (g), and (h) illustrate the FD–IS trade-offs, while (f), (i), and (j) depict the Precision–Recall trade-offs. For EDM2 models, we evaluate classifier-free guidance (CFG) scales $\gamma \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.7, 0.9, 1, 1.5, 2.0, 3.0\}$. For Var-d24 model, $\gamma \in \{0.5, 0.8, 1.1, 1.5, 1.7, 2.0, 2.5, 3.0, 4.0, 5.0, 7.0, 10.0, 15.0\}$. **Start** denotes the performance of the base model, while **End** denotes the model obtained after a fixed finetuning duration.

Finally, Direct Discriminative Optimization (DDO) [Zhe+25] contrasts real and synthetic samples rather than class-conditional distributions, and is included as a baseline in our experiments. Additional discussion of related work is deferred to Section H.

6 Experimental Results

Experiment Setup. In our experiments, we focus on fine-tuning pretrained models. For diffusion models, we fine-tune pretrained EDM2 models [Kar+24b] on ImageNet-64×64 and ImageNet-512×512, and SiT (with REPA) model [Yu+25] on ImageNet-256×256. For visual autoregressive models, we fine-tune VAR-d24 [Tia+24] on ImageNet-256×256. We compare against CFG (without interval) and several training-time baselines, including CC-DPO, CCA, CFM and DDO.

The generative performance is evaluated with Fréchet Distance (FD) [Heu+17], Precision and Recall [Kyn+19], and Inception score (IS) [Sal+16]. For diffusion models, we primarily report FD computed with DINOv2 features ($\text{FD}_{\text{DINOv2}}$). While FID (Inception-based FD) is widely used in the diffusion literature, we find that for strong pretrained models such as EDM2 which already achieves strong FID (1.5-2), it can be insensitive to perceptually meaningful improvement: although both CFG and MCLR lead to visually pronounced quality improvements, the FID score often does not improve and may even degrade. In contrast, $\text{FD}_{\text{DINOv2}}$ consistently captures these improvements, aligning with prior findings that it correlates more strongly with human evaluations [Ste+23]. For base models that are less strong including VAR and SiT, both FID and $\text{FD}_{\text{DINOv2}}$ improve consistently. Additional implementation details, ablations, and qualitative results deferred to Sections F and G.

Progressive Class Separation and Fidelity–Diversity Trade-off. We first analyze the training dynamics induced by MCLR and their effect on conditional generation. Qualitatively, as shown in Figures 2 and 10 to 21, at early stages of training, images generated from the same initial noises share similar global structures across different class conditions, indicating weak class-conditional modeling by the base model. As training proceeds, generated images gradually develop distinct

class-specific structures, reflecting increasing inter-class separation and improved conditional modeling.

Quantitatively, as shown in Figure 4(b,d), both Inception Score (IS) and Precision increase steadily during training, indicating that the generated samples become more class-discriminative and visually faithful. However, excessive training reduces within-class diversity, as reflected by a decrease in Recall (Figure 4(e)). As a consequence, $\text{FD}_{\text{DINOv2}}$ initially improves as conditional modeling strengthens, but later deteriorates when diversity decreases, as shown in Figure 4(a), revealing a characteristic fidelity–diversity trade-off.

This behavior closely mirrors the effect of increasing the guidance scale in CFG, where stronger guidance improves class fidelity at the cost of diversity. Similar trade-offs are observed for other contrastive alignment objectives such as CC-DPO and CCA. Following prior works, we therefore report results at the checkpoint achieving the best $\text{FD}_{\text{DINOv2}}$ score in Table 1 for each algorithm.

MCLR outperforms training-time baselines. Table 1 shows that MCLR consistently achieves the best $\text{FD}_{\text{DINOv2}}$ score among training-time baselines. On EDM2-S, MCLR improves FD-DINOv2 from 95.20 to 52.69, outperforming DDO, CCA, and CC-DPO. On EDM2-L, MCLR reaches 42.50 $\text{FD}_{\text{DINOv2}}$, substantially outperforming the strongest training-time baseline. On VAR-d24 and SiT (with REPA), MCLR achieves the best $\text{FD}_{\text{DINOv2}}$ among all compared methods, including CFG.

Moreover, for EDM models, as shown in Figure 4 (c,f), MCLR traverses a significantly wider fidelity–diversity trade-off region steadily with a faster training speed. In contrast, CCA and CC-DPO converge early at suboptimal local minima and exhibit zigzag optimization trajectories, as evidenced by slowly improving or stalled learning curves. A controlled comparison with CFM under matched fine-tuning settings leads to a similar conclusion: MCLR achieves larger and more sustained improvements across different DSM regularization strengths (see Section G.2).

MCLR Substantially Narrows the Gap to CFG. CFG still achieves the best $\text{FD}_{\text{DINOv2}}$ on EDM2 models, but MCLR substantially narrows the gap under standard guidance-free sampling. For example, on EDM2-L, MCLR obtains 42.50 FD-DINOv2 compared with CFG's 39.86, while using half the NFE. When evaluated using Precision–Recall, MCLR exhibits competitive performance relative to CFG. Specifically, as shown in Figure 4(f), MCLR matches CFG in the high-recall regime corresponding to early training stages, and attains a substantially higher best-case precision in the high-precision regime at later training stages, where CFG begins to produce images with oversaturated colors. Moreover, applying CFG on top of an MCLR-fine-tuned model further narrows the performance gap between the two methods, yielding higher best-case Inception Score and Precision on both EDM2-S and EDM2-L models (see Figure 4(c,f) and Figure 7(c,f,i)).

Qualitatively, MCLR and CFG produce highly similar visual effects, both substantially enhancing class-specific structures in the generated images, as shown in Figures 2 and 10 to 21. These observations are consistent with our theoretical analysis in Section 3.5, which interprets CFG as an implicit contrastive alignment method.

Equivalence between CCA and CC-DPO. Finally, we observe that CC-DPO consistently matches or outperforms CCA (see Section G.4). This empirical observation aligns with our theoretical analysis in Section 4.1, which establishes the equivalence between these two objectives.

Table 1: Quantitative results for MCLR, CFG and training-time baselines on ImageNet. For each algorithm, the metrics are reported at the model checkpoint (or CFG scale) achieving the best-case FD_{DINOv2}. Precision and Recall are computed using Inception features. The best and second-best results are highlighted in **bold** and underline, respectively.

(a) EDM2 (ImageNet-64 × 64)

Method	NFE	FD_{DINOv2}↓	Prec.↑	Rec.↑	IS↑
EDM2-S	63	95.20	0.705	0.614	60.43
+CFG	126	**43.75**	**0.800**	0.565	**127.40**
+DDO	63	72.97	0.689	**0.642**	65.84
+CCA	63	62.36	0.762	0.557	76.13
+CC-DPO	63	60.98	0.784	0.536	86.11
+MCLR (Ours)	63	52.69	**0.800**	0.505	90.68

(b) EDM2 (ImageNet-512 × 512)

Method	NFE	FD_{DINOv2}↓	Prec.↑	Rec.↑	IS↑
EDM2-L	63	67.70	0.753	0.610	250.07
+CFG	126	**39.86**	0.844	0.512	**360.30**
+DDO	63	49.47	0.737	**0.652**	268.72
+CCA	63	51.45	0.790	0.574	281.45
+CC-DPO	63	51.92	0.812	0.544	298.89
+MCLR (Ours)	63	42.50	**0.849**	0.492	332.02

(c) VAR (ImageNet-256 × 256)

Method	NFE	FD_{DINOv2}↓	Prec.↑	Rec.↑	IS↑
VAR-d24	10	182.12	0.672	**0.623**	67.92
+CFG	20	45.08	0.798	0.542	**110.70**
+CCA	10	46.82	0.873	0.448	98.92
+CC-DPO	10	46.63	0.881	0.433	100.12
+MCLR (Ours)	10	**44.31**	**0.893**	0.404	100.84

(d) SiT-XL/2+REPA (ImageNet-256 × 256)

Method	NFE	FD_{DINOv2}↓	Prec.↑	Rec.↑	IS↑
SiT-XL/2+REPA	50	184.36	0.603	**0.683**	142.24
+CFG	100	50.23	0.833	0.506	**385.87**
+CCA	50	60.20	0.769	0.551	254.56
+CC-DPO	50	47.20	**0.839**	0.451	329.83
+MCLR (Ours)	50	**45.96**	0.825	0.488	311.00

7 Limitations and Conclusions

This work introduced MCLR, a training-time objective that substantially improves the unguided base model and narrows the gap to inference-time CFG. Beyond empirical improvements, our analysis reveals a close connection between MCLR and CFG, identifying the implicit inter-class contrastive alignment mechanism of CFG.

Limitations and Future Directions. As a training-time method, MCLR produces a fixed model and cannot dynamically adjust the fidelity–diversity trade-off at inference time as CFG does. CFG also typically achieves stronger best-case FD and Inception Score. Bridging training-time objectives and inference-time guidance remains an important future direction. More fundamentally, these observations raise a broader question:

Should alignment be performed at training time, or is inference-time guidance a more effective paradigm?

Recent studies [Fra+25; Jin+25; Che+25c; JWL26] suggest that inference-time guidance can, in certain settings, outperform training-based alignment algorithms. These observations point toward a promising research direction: developing alignment algorithms that operate directly at inference time. Exploring this possibility may lead to a broader shift from traditional training-time alignment toward inference-time alignment paradigms.

Lastly, the lack of inter-class separation is only one practical limitation of the standard DSM; systematically identifying and addressing additional limitations remains an important future direction.

Data Availability Statement

The code and instructions for reproducing the experimental results will be made available in our code repository.

Acknowledgment

We acknowledge funding support from NSF CCF-2212066, NSF CCF- 2212326, NSF IIS 2402950, and ONR N000142512339. This research used the Delta advanced computing and data resource which is supported by the National Science Foundation (award OAC 2005572) and the State of Illinois. Delta is a joint effort of the University of Illinois Urbana-Champaign and its National Center for Supercomputing Applications[Boe+23].

References

[And82] Brian DO Anderson. "Reverse-time diffusion equation models". In: *Stochastic Processes and their Applications* 12.3 (1982), pp. 313–326.

[BN24] Arwen Bradley and Preetum Nakkiran. "Classifier-free guidance is a predictor-corrector". In: *arXiv preprint arXiv:2408.09000* (2024).

[Boe+23] Timothy J Boerner, Stephen Deems, Thomas R Furlani, Shelley L Knuth, and John Towns. "Access: Advancing innovation: Nsf's advanced cyberinfrastructure coordination ecosystem: Services & support". In: *Practice and Experience in Advanced Research Computing 2023: Computing for the Common Good*. 2023, pp. 173–176.

[BT52] Ralph Allan Bradley and Milton E Terry. "Rank analysis of incomplete block designs: I. The method of paired comparisons". In: *Biometrika* 39.3/4 (1952), pp. 324–345.

[BV04] Stephen P Boyd and Lieven Vandenberghe. *Convex optimization*. Cambridge university press, 2004.

[Che+25a] Huayu Chen, Kai Jiang, Kaiwen Zheng, Jianfei Chen, Hang Su, and Jun Zhu. "Visual Generation Without Guidance". In: *Forty-second International Conference on Machine Learning*. 2025.

[Che+25b] Huayu Chen, Hang Su, Peize Sun, and Jun Zhu. "Toward Guidance-Free AR Visual Generation via Condition Contrastive Alignment". In: *The Thirteenth International Conference on Learning Representations*. 2025.

[Che+25c] Min Cheng, Fatemeh Doudi, Dileep Kalathil, Mohammad Ghavamzadeh, and Panganamala R Kumar. "Diffusion Blend: Inference-Time Multi-Preference Alignment for Diffusion Models". In: *arXiv preprint arXiv:2505.18547* (2025).

[Chi+24] Muthu Chidambaram, Khashayar Gatmiry, Sitan Chen, Holden Lee, and Jianfeng Lu. "What does guidance do? A fine-grained analysis in a simple setting". In: *The Thirty-eighth Annual Conference on Neural Information Processing Systems*. 2024.

[Fra+25] Kevin Frans, Seohong Park, Pieter Abbeel, and Sergey Levine. "Diffusion guidance is a controllable policy improvement operator". In: *arXiv preprint arXiv:2505.23458* (2025).

[GH10] Michael Gutmann and Aapo Hyvärinen. "Noise-contrastive estimation: A new estimation principle for unnormalized statistical models". In: *Proceedings of the thirteenth international conference on artificial intelligence and statistics*. JMLR Workshop and Conference Proceedings. 2010, pp. 297–304.

[HE26] Peter Holderrieth and Ezra Erives. *Introduction to Flow Matching and Diffusion Models*. 2026. arXiv: 2506.02070. URL: https://diffusion.csail.mit.edu/.

[Heu+17] 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: *Advances in neural information processing systems* 30 (2017).

[HJA20] Jonathan Ho, Ajay Jain, and Pieter Abbeel. "Denoising diffusion probabilistic models". In: *Advances in neural information processing systems* 33 (2020), pp. 6840–6851.

[HS22] Jonathan Ho and Tim Salimans. "Classifier-free diffusion guidance". In: *arXiv preprint arXiv:2207.12598* (2022).

[Jin+25] Luozhijie Jin, Zijie Qiu, Jie Liu, Zijie Diao, Lifeng Qiao, Ning Ding, Alex Lamb, and Xipeng Qiu. "Inference-time alignment control for diffusion models with reinforcement learning guidance". In: *arXiv preprint arXiv:2508.21016* (2025).

[JSG25] Cheng Jin, Qitan Shi, and Yuantao Gu. "Stage-wise Dynamics of Classifier-Free Guidance in Diffusion Models". In: *arXiv preprint arXiv:2509.22007* (2025).

[JWL26] Zhou Jiang, Yandong Wen, and Zhen Liu. "Rethinking Preference Alignment for Diffusion Models with Classifier-Free Guidance". In: *arXiv preprint arXiv:2602.18799* (2026).

[Kar+22] Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. "Elucidating the design space of diffusion-based generative models". In: *Advances in neural information processing systems* 35 (2022), pp. 26565–26577.

[Kar+24a] Tero Karras, Miika Aittala, Tuomas Kynkäänniemi, Jaakko Lehtinen, Timo Aila, and Samuli Laine. "Guiding a diffusion model with a bad version of itself". In: *Advances in Neural Information Processing Systems* 37 (2024), pp. 52996–53021.

[Kar+24b] Tero Karras, Miika Aittala, Jaakko Lehtinen, Janne Hellsten, Timo Aila, and Samuli Laine. "Analyzing and improving the training dynamics of diffusion models". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2024, pp. 24174–24184.

[KMS24] Zahra Kadkhodaie, Stéphane Mallat, and Eero P Simoncelli. "Feature-guided score diffusion for sampling conditional densities". In: *arXiv preprint arXiv:2410.11646* (2024).

[Kyn+19] Tuomas Kynkäänniemi, Tero Karras, Samuli Laine, Jaakko Lehtinen, and Timo Aila. "Improved precision and recall metric for assessing generative models". In: *Advances in neural information processing systems* 32 (2019).

[Lee+25] Jaa-Yeon Lee, Byunghee Cha, Jeongsol Kim, and Jong Chul Ye. "Aligning text to image in diffusion models is easier than you think". In: *arXiv preprint arXiv:2503.08250* (2025).

[Lip+23] Yaron Lipman, Ricky TQ Chen, Heli Ben-Hamu, Maximilian Nickel, and Matthew Le. "Flow Matching for Generative Modeling". In: *The Eleventh International Conference on Learning Representations*. 2023.

[LJ25] Gen Li and Yuchen Jiao. "Provable Efficiency of Guidance in Diffusion Models for General Data Distribution". In: *Forty-second International Conference on Machine Learning*. 2025.

[LWQ26] Xiang Li, Rongrong Wang, and Qing Qu. "Towards understanding the mechanisms of classifier-free guidance". In: *Advances in Neural Information Processing Systems* 38 (2026), pp. 62835–62889.

[Mar+23] Morteza Mardani, Jiaming Song, Jan Kautz, and Arash Vahdat. "A Variational Perspective on Solving Inverse Problems with Diffusion Models". In: *The Twelfth International Conference on Learning Representations*. 2023.

[Miy+61] Koichi Miyasawa et al. "An empirical Bayes estimator of the mean of a normal population". In: *Bull. Inst. Internat. Statist* 38.181-188 (1961), pp. 1–2.

[OLV18] Aaron van den Oord, Yazhe Li, and Oriol Vinyals. "Representation learning with contrastive predictive coding". In: *arXiv preprint arXiv:1807.03748* (2018).

[Pav+25] Krunoslav Lehman Pavasovic, Jakob Verbeek, Giulio Biroli, and Marc Mezard. "Classifier-Free Guidance: From High-Dimensional Analysis to Generalized Guidance Forms". In: *arXiv preprint arXiv:2502.07849* (2025).

[PX23] William Peebles and Saining Xie. "Scalable diffusion models with transformers". In: *Proceedings of the IEEE/CVF international conference on computer vision*. 2023, pp. 4195–4205.

[Raf+23] Rafael Rafailov, Archit Sharma, Eric Mitchell, Christopher D Manning, Stefano Ermon, and Chelsea Finn. "Direct preference optimization: Your language model is secretly a reward model". In: *Advances in Neural Information Processing Systems* 36 (2023), pp. 53728–53741.

[Sal+16] Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen. "Improved techniques for training gans". In: *Advances in neural information processing systems* 29 (2016).

[Son+21a] Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. "Maximum likelihood training of score-based diffusion models". In: *Advances in neural information processing systems* 34 (2021), pp. 1415–1428.

[Son+21b] Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben Poole. "Score-Based Generative Modeling through Stochastic Differential Equations". In: *International Conference on Learning Representations*. 2021.

[SS19] Simo Särkkä and Arno Solin. "Applied stochastic differential equations". In: *Institute of mathematical statistics textbooks* (2019).

[Ste+23] George Stein, Jesse Cresswell, Rasa Hosseinzadeh, Yi Sui, Brendan Ross, Valentin Villecroze, Zhaoyan Liu, Anthony L Caterini, Eric Taylor, and Gabriel Loaiza-Ganem. "Exposing flaws of generative model evaluation metrics and their unfair treatment of diffusion models". In: *Advances in Neural Information Processing Systems* 36 (2023), pp. 3732–3784.

[Sto+25] George Stoica, Vivek Ramanujan, Xiang Fan, Ali Farhadi, Ranjay Krishna, and Judy Hoffman. *Contrastive Flow Matching*. 2025. arXiv: 2506.05350 [cs.CV]. URL: https://arxiv.org/abs/2506.05350.

[Tan+25] Zhicong Tang, Jianmin Bao, Dong Chen, and Baining Guo. "Diffusion models without classifier-free guidance". In: *arXiv preprint arXiv:2502.12154* (2025).

[Tia+24] Keyu Tian, Yi Jiang, Zehuan Yuan, Bingyue Peng, and Liwei Wang. "Visual autoregressive modeling: Scalable image generation via next-scale prediction". In: *Advances in neural information processing systems* 37 (2024), pp. 84839–84865.

[Ven+26] Enrico Ventura, Beatrice Achilli, Luca Ambrogioni, and Carlo Lucibello. "Emergence of Distortions in High-Dimensional Guided Diffusion Models". In: *arXiv preprint arXiv:2602.00716* (2026).

[Vin11] Pascal Vincent. "A Connection Between Score Matching and Denoising Autoencoders". In: *Neural Computation* 23.7 (2011), pp. 1661–1674. DOI: 10.1162/NECO_a_00142.

[Wal+24] Bram Wallace, Meihua Dang, Rafael Rafailov, Linqi Zhou, Aaron Lou, Senthil Purushwalkam, Stefano Ermon, Caiming Xiong, Shafiq Joty, and Nikhil Naik. "Diffusion model alignment using direct preference optimization". In: *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2024, pp. 8228–8238.

[Wu+24] Yuchen Wu, Minshuo Chen, Zihao Li, Mengdi Wang, and Yuting Wei. "Theoretical insights for diffusion guidance: A case study for Gaussian mixture models". In: *Forty-first International Conference on Machine Learning*. 2024.

[Xue+25] Shuchen Xue, Chongjian Ge, Shilong Zhang, Yichen Li, and Zhi-Ming Ma. "Advantage weighted matching: Aligning rl with pretraining in diffusion models". In: *arXiv preprint arXiv:2509.25050* (2025).

[YAA25] Junno Yun, Yaşar Utku Alçalar, and Mehmet Akçakaya. "No Alignment Needed for Generation: Learning Linearly Separable Representations in Diffusion Models". In: *arXiv preprint arXiv:2509.21565* (2025).

[Yan+24] Divin Yan, Lu Qi, Vincent Tao Hu, Ming-Hsuan Yang, and Meng Tang. "Training class-imbalanced diffusion model via overlap optimization". In: *arXiv preprint arXiv:2402.10821* (2024).

[Yan+26] Ruofeng Yang, Yiyu Qiu, Bo Jiang, Cheng Chen, and Shuai Li. *Elucidating Guidance in Variance Exploding Diffusion Models: Fast Convergence and Better Diversity*. OpenReview. https://openreview.net/forum?id=tWbasgJA68. 2026.

[Yu+25] Sihyun Yu, Sangkyung Kwak, Huiwon Jang, Jongheon Jeong, Jonathan Huang, Jinwoo Shin, and Saining Xie. "Representation Alignment for Generation: Training Diffusion Transformers Is Easier Than You Think". In: *The Thirteenth International Conference on Learning Representations*. 2025.

[Zhe+25] Kaiwen Zheng, Yongxin Chen, Huayu Chen, Guande He, Ming-Yu Liu, Jun Zhu, and Qinsheng Zhang. "Direct discriminative optimization: Your likelihood-based visual generative model is secretly a gan discriminator". In: *arXiv preprint arXiv:2503.01103* (2025).

Appendices

A Weak Inter-Class Separation in Diffusion Models

In Section 3.1, we argue that conditional diffusion models trained with denoising score matching (DSM) may suffer from weak inter-class separation, yielding class-conditional distributions that are less separated than their real-data counterparts. In this section, we provide quantitative evidence for this phenomenon by comparing pairwise Fréchet Inception Distances (FID) among real class-conditional distributions with the corresponding distances among generated class-conditional distributions. Specifically, for two classes c_1 and c_2, let $\mathcal{D}_{c_1}$ and $\mathcal{D}_{c_2}$ denote the corresponding real training data sets, and let $\mathcal{G}_{c_1}$ and $\mathcal{G}_{c_2}$ denote the sample sets generated by a conditional diffusion model. We compute $\mathrm{FID}(\mathcal{D}_{c_1}, \mathcal{D}_{c_2})$ and $\mathrm{FID}(\mathcal{G}_{c_1}, \mathcal{G}_{c_2})$. Note that FID is a measure of distances between two distributions in the InceptionV3 feature space, and a larger FID means the two measured distributions are more separated. Therefore, if $\mathrm{FID}(\mathcal{G}_{c_1}, \mathcal{G}_{c_2}) < \mathrm{FID}(\mathcal{D}_{c_1}, \mathcal{D}_{c_2})$, the generated conditional distributions for c_1 and c_2 are less separated than their real-data counterparts.

We perform this evaluation on eight selected ImageNet classes following [Kyn+19]: four identified as relatively easy to model—great Pyrenees (GP), Egyptian cat (EC), broccoli (BR) and lemon (LE)—and four relatively difficult classes—cornet/trumpet (CT), park bench (PB), bubble (BU), and baseball player (BP). As shown in [Kyn+19], the easy classes are largely textural or are common in the datasets while the hard classes often contain precise global structure or are underrepresented in the dataset. We generate samples using SiT-XL/2 with REPA and a 50-step Euler ODE sampler, without classifier-free guidance (CFG). For each class, we use all available training images and generate an equal number of samples for computing FID.

Table 2 reports the pairwise inter-class FID. Across all 28 evaluated class pairs, the generated pairwise FID is lower than its real-data counterpart. This suggests that the generated class-conditional distributions are systematically less separated than the corresponding real-data distributions. We next examine whether CFG and MCLR can recover this missing separation. Table 3 reports the mean pairwise FID over all 28 class pairs. Both methods increase inter-class separation. Figure 5 qualitatively demonstrate the weak class separation issue, where initial random noise is fixed across all class conditions and generation methods.

Table 2: **Pairwise inter-class FID for real and generated samples.** Each entry reports $\mathrm{FID}(\mathcal{D}_{c_i}, \mathcal{D}_{c_j})/\mathrm{FID}(\mathcal{G}_{c_i}, \mathcal{G}_{c_j})$. Only the upper triangle is shown. Across all 28 class pairs, the generated pairwise FID is lower than the corresponding real-data pairwise FID, indicating weaker inter-class separation in the generated distributions. The larger value in each entry is shown in bold.

	GP	EC	CT	PB	BR	LE	BU	BP
GP	–	**227**/186	**234**/201	**294**/213	**328**/309	**249**/232	**324**/262	**284**/248
EC		–	**233**/196	**286**/208	**315**/312	**254**/240	**307**/242	**275**/247
CT			–	**281**/190	**303**/287	**238**/217	**286**/198	**265**/212
PB				–	**380**/327	**299**/243	**347**/239	**309**/223
BR					–	**310**/264	**382**/346	**356**/328
LE						–	**321**/274	**306**/273
BU							–	**354**/276
BP								–

Table 3: **Mean pairwise inter-class FID over the 28 class pairs.** The real training data serve as the reference for inter-class separation. The absolute gap measures deviation from this reference; lower is better. Here, γ denotes the CFG scale, with $\gamma = 0$ corresponding to unguided conditional sampling.

Source / Model	Sampling	Mean Pairwise FID	Absolute Gap ↓
Real training data	–	298.1	–
SiT-XL/2 + REPA	Unguided ($\gamma = 0$)	249.8	48.3
SiT-XL/2 + REPA	CFG ($\gamma = 1$)	310.5	12.4
SiT-XL/2 + REPA + MCLR	Unguided ($\gamma = 0$)	287.7	**10.4**

Lastly, we note that similar evidence of weak inter-class separation was previously reported by [LWQ26] using EDM on ImageNet-64 × 64 [Kar+22]. Our results extend this observation to the more recent transformer-based SiT-XL/2 architecture with REPA at 256 × 256 resolution, suggesting that the phenomenon persists across model architectures and image resolutions.

B Theoretical Analysis of MCLR

B.1 Equivalent Forms of MCLR Objective

By the linearity of expectation, it is easy to show that the MCLR objective (10) admits the following two equivalent forms:

$$\max_{\theta} \ \mathbb{E}_{\boldsymbol{c},\, p(\boldsymbol{x}|\boldsymbol{c})}\Big[\log p_\theta(\boldsymbol{x}|\boldsymbol{c})\Big] + \eta \underbrace{\mathbb{E}_{\boldsymbol{c},\,\tilde{\boldsymbol{c}},\,\boldsymbol{x}\sim p(\cdot|\boldsymbol{c}),\boldsymbol{y}\sim p(\cdot|\tilde{\boldsymbol{c}})}\Big[\log \frac{p_\theta(\boldsymbol{x}|\boldsymbol{c})}{p_\theta(\boldsymbol{y}|\boldsymbol{c})}\Big]}_{\text{MCLR Regularization (Form I)}}, \tag{27}$$

and

$$\max_{\theta} \ \mathbb{E}_{\boldsymbol{c},\, p(\boldsymbol{x}|\boldsymbol{c})}\Big[\log p_\theta(\boldsymbol{x}|\boldsymbol{c})\Big] + \eta \underbrace{\mathbb{E}_{\boldsymbol{c},\,\tilde{\boldsymbol{c}},\,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c})}\Big[\log \frac{p_\theta(\boldsymbol{x}|\boldsymbol{c})}{p_\theta(\boldsymbol{x}|\tilde{\boldsymbol{c}})}\Big]}_{\text{MCLR Regularization (Form II)}}. \tag{28}$$

While the main text presents MCLR in Form I, the appendix derives most results from Form II. Because these two formulations are equivalent, we use them interchangeably throughout the analysis.

B.2 Distributional Analysis of MCLR

In this section, we provide the proof for Theorem 1. We first restate the assumptions and theorem.

Assumption 1. *The function $h(\boldsymbol{x}|\boldsymbol{c})$ has compact support; that is,*

$$\operatorname{supp} h(\boldsymbol{x}|\boldsymbol{c}) \ \subseteq \ K, \qquad |K| < \infty, \tag{29}$$

where $|K|$ denotes the volume of set K.



Figure 5: **Qualitative Comparison of Inter-Class Separation.** All panels share the same initial noise; columns vary the class condition, while rows vary the generation method. Under unguided sampling, the baseline model preserves similar class-agnostic structures across different labels. Increasing the CFG strength progressively enhances class-specific structure. MCLR produces a similar effect under unguided sampling, without applying CFG at inference time.

Assumption 2. *For* $\forall \boldsymbol{x} \in K$,

$$p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}) \geq \delta > 0, \qquad \delta < \frac{1}{|K|}. \tag{30}$$

Assumption 1 is mild in practice, as image data typically occupies a bounded pixel range. Assumption 2 is a regularity assumption that ensures the log-likelihood is well-defined, avoiding the singularity that occurs when evaluating $\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})$ at zero density. Since $p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})$ integrates to one over K, we necessarily have $\delta \leq \frac{1}{|K|}$. In the following, we further assume $\delta < \frac{1}{|K|}$.

Theorem 1. *Under the two assumptions, the optimal solution to* (11) *is:*

$$p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c}) \;=\; \begin{cases} \max\left\{\frac{h(\boldsymbol{x}|\boldsymbol{c})}{Z(\boldsymbol{c})}, \delta\right\}, & \boldsymbol{x} \in K, \\ 0, & \boldsymbol{x} \notin K, \end{cases} \tag{31}$$

where $Z(\boldsymbol{c})$ *is the normalizing constant.*

Proof. Without loss of generality, we consider discrete classes in this proof; the theorem extends straightforwardly to continuous classes. Suppose there are M classes $\{\boldsymbol{c}_i\}_{i=1}^{M}$, each class has prior probability $p(\boldsymbol{c}_i)$ and $\sum_{i=1}^{M} p(\boldsymbol{c}_i) = 1$, then the training objective (11) takes the following form:

$$\max_{\boldsymbol{\theta}} \mathcal{L}(\boldsymbol{\theta}) := \mathbb{E}_{\boldsymbol{c},p(\boldsymbol{x}|\boldsymbol{c})}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})\right] + \eta\mathbb{E}_{\boldsymbol{c},\tilde{\boldsymbol{c}},\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c})}\left[\log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})}{p_{\boldsymbol{\theta}}(\boldsymbol{x}|\tilde{\boldsymbol{c}})}\right] \tag{32}$$

$$= \max_{\boldsymbol{\theta}} \sum_{i=1}^{M} p(\boldsymbol{c}_i)\mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{c}_i)}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_i)\right] + \eta \sum_{i=1}^{M} p(\boldsymbol{c}_i) \sum_{j=1}^{M} p(\boldsymbol{c}_j)\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c}_i)}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_i)\right] \tag{33}$$

$$-\eta \sum_{i=1}^{M} p(\boldsymbol{c}_i) \sum_{j=1}^{M} p(\boldsymbol{c}_j)\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c}_i)}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_j)\right]. \tag{34}$$

Note that we can decompose overall objective $\mathcal{L}(\boldsymbol{\theta})$ as:

$$\mathcal{L}(\boldsymbol{\theta}) = \sum_{k=1}^{M} \mathcal{L}_k(\boldsymbol{\theta}), \tag{35}$$

where $\mathcal{L}_k(\boldsymbol{\theta})$ is the amount contributed by $p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)$:

$$\mathcal{L}_k(\boldsymbol{\theta}) = p(\boldsymbol{c}_k)\mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{c}_k)}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\right] + \eta p(\boldsymbol{c}_k)\mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{c}_k)}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\right] \tag{36}$$

$$-\eta p(\boldsymbol{c}_k) \sum_{i=1}^{M} p(\boldsymbol{c}_i)\mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{c}_i)}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\right] \tag{37}$$

$$= p(\boldsymbol{c}_k)\mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{c}_k)}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\right] + \eta p(\boldsymbol{c}_k)\mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{c}_k)}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\right] - \eta p(\boldsymbol{c}_k)\mathbb{E}_{p(\boldsymbol{x})}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\right]. \tag{38}$$

Note that we can optimize $L_k(\boldsymbol{\theta})$ for each $k \in \{1, \ldots, M\}$ to get the optimal conditional distribution $p(\boldsymbol{x}|\boldsymbol{c}_k)$ independently:

$$\max_{\boldsymbol{\theta}} \mathcal{L}_k(\boldsymbol{\theta}) \Leftrightarrow \max_{p_{\boldsymbol{\theta}}(\cdot|\boldsymbol{c}_k)} \mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{c}_k)}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\right] + \eta\mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{c}_k)}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\right] - \eta\mathbb{E}_{p(\boldsymbol{x})}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\right] \tag{39}$$

$$= \max_{p_{\boldsymbol{\theta}}(\cdot|\boldsymbol{c}_k)} \int_K \log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\left(p(\boldsymbol{x}|\boldsymbol{c}_k) + \eta(p(\boldsymbol{x}|\boldsymbol{c}_k) - p(\boldsymbol{x}))\right) \mathrm{d}\boldsymbol{x}. \tag{40}$$

In what follows, we drop the subscript k, so that the optimization problem becomes:

$$\max_{p_{\boldsymbol{\theta}}(\cdot|\boldsymbol{c})} \int_K \log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})\left(p(\boldsymbol{x}|\boldsymbol{c}) + \eta(p(\boldsymbol{x}|\boldsymbol{c}) - p(\boldsymbol{x}))\right) \mathrm{d}\boldsymbol{x} \tag{41}$$

$$= \max_{p_{\boldsymbol{\theta}}(\cdot|\boldsymbol{c})} \int_K \log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}) h(\boldsymbol{x}|\boldsymbol{c}) \mathrm{d}\boldsymbol{x} \tag{42}$$

Note that (42) shares the similar form as KL divergence, but $h(\boldsymbol{x}|\boldsymbol{c})$ is not a valid probability distribution, since there could exists $\boldsymbol{x}$ such that $h(\boldsymbol{x}|\boldsymbol{c}) < 0$. In this case, setting $p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}) = 0$ makes (42) approaches $+\infty$, hence the optimization problem does not have an attainable optimum. To make the optimization problem well-posed, we impose Assumption 2.

Define the sets

$$K_+(\boldsymbol{c}) = \{\boldsymbol{x} \in K : h(\boldsymbol{x}|\boldsymbol{c}) > 0\}, \qquad K_-(\boldsymbol{c}) = \{\boldsymbol{x} \in K : h(\boldsymbol{x}|\boldsymbol{c}) \leq 0\}. \tag{43}$$

Under Assumption 2, it is straightforward to show the optimal distribution $p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c})$ must attain the lower bound δ for $\forall \boldsymbol{x} \in K_-(\boldsymbol{c})$; otherwise increasing the density will decrease the objective (42). Furthermore for $\boldsymbol{x} \in K^c$, where $h(\boldsymbol{x}|\boldsymbol{c}) = 0$, the optimal density must satisfy $p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c}) = 0$ for almost

every $\boldsymbol{x} \in K^c$; otherwise, probability mass could be shifted from K^c to $K_+(\boldsymbol{c})$to further increase the objective (42). Therefore, the remaining optimization concerns the density on $K^+(\boldsymbol{c})$, which is the optimal solution to the following constrained optimization problem:

$$\min_{p_{\boldsymbol{\theta}}(\cdot|\boldsymbol{c})} \int_{K_+(\boldsymbol{c})} -\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}) h(\boldsymbol{x}|\boldsymbol{c}) \mathrm{d}\boldsymbol{x} \tag{44}$$

$$\text{s.t. } -p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}) + \delta \leq 0, \ \forall \boldsymbol{x} \in K_+(\boldsymbol{c}) \tag{45}$$

$$\int_{K_+(\boldsymbol{c})} p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}) \mathrm{d}\boldsymbol{x} - m(\boldsymbol{c}) = 0, \tag{46}$$

where

$$m(\boldsymbol{c}) = 1 - \int_{K_-(\boldsymbol{c})} \delta \mathrm{d}\boldsymbol{x} = 1 - \delta |K_-(\boldsymbol{c})|. \tag{47}$$

Since $\delta < \frac{1}{|K|}$ by assumption, we have $0 < m(\boldsymbol{c}) \leq 1$.

Note that this optimization problem is convex in $p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})$. Treating $p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})$ for each $\boldsymbol{x}$ as optimization variables, the constraints are affien and Slater's condition holds. Therefore strong duality applies, and the optimal solution can be characterized by the KKT conditions [BV04].

Define the Lagrangian as:

$$\mathcal{L}(p_{\boldsymbol{\theta}}(\cdot|\boldsymbol{c}), \lambda, u(\cdot)) = \int_{K_+(\boldsymbol{c})} -\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}) h(\boldsymbol{x}|\boldsymbol{c}) \mathrm{d}\boldsymbol{x} + \lambda \Big(\int_{K_+(\boldsymbol{c})} p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}) \mathrm{d}\boldsymbol{x} - m(\boldsymbol{c}) \Big) \tag{48}$$

$$+ \int_{K_+(\boldsymbol{c})} u(\boldsymbol{x})(-p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}) + \delta) \mathrm{d}\boldsymbol{x}, \tag{49}$$

where λ and $\mu(\cdot)$ are the dual variables.

For each $\boldsymbol{x} \in K_+(\boldsymbol{c})$, treating $p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c})$ as a pointwise optimization variable, let $p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c})$, λ^* and $u^*(\boldsymbol{x})$ be the corresponding optimal primal and dual variables. Applying stationary condition of the KKT conditions, we have:

$$\nabla_{p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c})} \mathcal{L}(p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c}), \lambda^*, u^*(\boldsymbol{x})) = 0 \tag{50}$$

$$\Rightarrow -\frac{h(\boldsymbol{x}|\boldsymbol{c})}{p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c})} + \lambda^* - u^*(\boldsymbol{x}) = 0. \tag{51}$$

Based on (51), we consider the following two cases:

- Suppose $u^*(\boldsymbol{x}) = 0$, we have $p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c}) = \frac{h(\boldsymbol{x}|\boldsymbol{c})}{\lambda^*}$.
- Suppose $u^*(\boldsymbol{x}) \neq 0$, by complementary slackness, we have $p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c}) = \delta$ and $u^*(\boldsymbol{x}) = \lambda^* - \frac{h(\boldsymbol{x}|\boldsymbol{c})}{\delta}$. Applying dual feasibility $u^*(\boldsymbol{x}) \geq 0$, we have $h(\boldsymbol{x}|\boldsymbol{c}) \leq \lambda^* \delta$.

Note that the above two cases can be combined as:

$$p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c}) = \max\Big\{\frac{h(\boldsymbol{x}|\boldsymbol{c})}{\lambda^*}, \delta\Big\}, \ \forall \boldsymbol{x} \in K_+(\boldsymbol{c}). \tag{52}$$

Next, we prove λ^* exists, i.e., (52) is normalizable. By applying the primal feasibility, we have:

$$\int_{K_+(\boldsymbol{c})} p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c}) \mathrm{d}\boldsymbol{x} = m(\boldsymbol{c}) \tag{53}$$

$$\Rightarrow \int_{K_+(\boldsymbol{c}) \cap \{h(\boldsymbol{x}|\boldsymbol{c}) \leq \lambda^* \delta\}} \delta \mathrm{d}\boldsymbol{x} + \int_{K_+(\boldsymbol{c}) \cap \{h(\boldsymbol{x}|\boldsymbol{c}) > \lambda^* \delta)\}} \frac{h(\boldsymbol{x}|\boldsymbol{c})}{\lambda^*} \mathrm{d}\boldsymbol{x} = m(\boldsymbol{c}). \tag{54}$$

Define:

$$A(\lambda) := \int_{K_+(\boldsymbol{c})\cap\{h(\boldsymbol{x}|\boldsymbol{c})\le\lambda\delta)\}} \delta \mathrm{d}\boldsymbol{x} + \int_{K_+(\boldsymbol{c})\cap\{h(\boldsymbol{x}|\boldsymbol{c})>\lambda\delta\}} \frac{h(\boldsymbol{x}|\boldsymbol{c})}{\lambda}\mathrm{d}\boldsymbol{x}. \tag{55}$$

Suppose $0 < \lambda_1 < \lambda_2$, we have:

$$A(\lambda_2) - A(\lambda_1) = \int_{K_+(\boldsymbol{c})\cap\{\lambda_1\delta\le h(\boldsymbol{x}|\boldsymbol{c})\le\lambda_2\delta\}} \delta \mathrm{d}\boldsymbol{x} - \int_{K_+(\boldsymbol{c})\cap\{\lambda_1\delta\le h(\boldsymbol{x}|\boldsymbol{c})\le\lambda_2\delta\}} \frac{h(\boldsymbol{x}|\boldsymbol{c})}{\lambda_1}\mathrm{d}\boldsymbol{x} \tag{56}$$

$$+ \int_{K_+(\boldsymbol{c})\cap\{h(\boldsymbol{x}|\boldsymbol{c})>\lambda_2\delta\}} (\frac{h(\boldsymbol{x}|\boldsymbol{c})}{\lambda_2} - \frac{h(\boldsymbol{x}|\boldsymbol{c})}{\lambda_1})\mathrm{d}\boldsymbol{x} \tag{57}$$

$$= \int_{K_+(\boldsymbol{c})\cap\{\lambda_1\delta\le h(\boldsymbol{x}|\boldsymbol{c})\le\lambda_2\delta\}} (\delta - \frac{h(\boldsymbol{x}|\boldsymbol{c})}{\lambda_1})\mathrm{d}\boldsymbol{x} + \int_{K_+(\boldsymbol{c})\cap\{h(\boldsymbol{x})>\lambda_2\delta\}} (\frac{h(\boldsymbol{x}|\boldsymbol{c})}{\lambda_2} - \frac{h(\boldsymbol{x}|\boldsymbol{c})}{\lambda_1})\mathrm{d}\boldsymbol{x} \tag{58}$$

$$< 0, \tag{59}$$

which implies $A(\lambda)$ is a monotonically decreasing function of $\lambda > 0$. By the assumption $\delta < \frac{1}{|K|}$, we have $\delta|K_-(\boldsymbol{c})| + \delta|K_+(\boldsymbol{c})| < 1$, which implies $\delta|K_+(\boldsymbol{c})| < m(\boldsymbol{c})$.

Since:

$$\lim_{\lambda\to 0} A(\lambda) = +\infty, \quad \lim_{\lambda\to+\infty} A(\lambda) = \delta|K_+(\boldsymbol{c})|, \tag{60}$$

as long as $\delta|K_+(\boldsymbol{c})| < m(\boldsymbol{c})$, by the intermediate value theorem, there exists a finite dual optimal point $\lambda^* > 0$ such that $A(\lambda^*) = m(\boldsymbol{c})$, i.e., the primal feasibility (54) holds. Hence, $p_{\boldsymbol{\theta}^*}(\boldsymbol{x})$ in (52) is a valid, normalizable probability distribution. Since the exact value of λ^* is dependent on the specific condition $\boldsymbol{c}$, we replace it with the notation $Z(\boldsymbol{c})$, which leads to the final result:

$$p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c}) = \begin{cases} \max\{\frac{h(\boldsymbol{x}|\boldsymbol{c})}{Z(\boldsymbol{c})}, \delta\}, & \boldsymbol{x} \in K_+(\boldsymbol{c}), \\ \delta, & \boldsymbol{x} \in K_-(\boldsymbol{c}), \\ 0, & \boldsymbol{x} \notin K. \end{cases} \tag{61}$$

, which can be further simplified as:

$$p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c}) = \begin{cases} \max\{\frac{h(\boldsymbol{x}|\boldsymbol{c})}{\lambda^*}, \delta\}, & \boldsymbol{x} \in K, \\ 0, & \boldsymbol{x} \notin K. \end{cases} \tag{62}$$

In the limit $\delta \to 0$, the optimal distribution approaches:

$$p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c}) = \frac{h^+(\boldsymbol{x}|\boldsymbol{c})}{\int_{\boldsymbol{x}} h^+(\boldsymbol{x}|\boldsymbol{c})\mathrm{d}\boldsymbol{x}}, \qquad h^+(\boldsymbol{x}|\boldsymbol{c}) := \max\{h(\boldsymbol{x}|\boldsymbol{c}), 0\} \tag{63}$$

which simply zeros out the negative part of $h(\boldsymbol{x}|\boldsymbol{c})$ and renormalizes it as a valid distribution. This completes the proof. □

B.3 Distributional Analysis of Fine-tuning with MCLR

Given a base model $p_{\text{ref}}(\boldsymbol{x})$ that lacks class specificity, we may fine-tune it using MCLR combined with KL regularization:

$$\max_{\boldsymbol{\theta}} -\mathbb{E}_{\boldsymbol{c}}\left[D_{\text{KL}}\big(p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})\|p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})\big)\right] + \eta\mathbb{E}_{\boldsymbol{c},\tilde{\boldsymbol{c}},\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c})}\left[\log\frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})}{p_{\boldsymbol{\theta}}(\boldsymbol{x}|\tilde{\boldsymbol{c}})}\right]. \tag{64}$$

Similar to (39), we can get the optimal conditional distribution $p(\boldsymbol{x}|\boldsymbol{c}_k)$ for each $k \in \{1, \ldots, M\}$ by solving the following optimization problem:

$$\arg\max_{p_{\boldsymbol{\theta}}(\cdot|\boldsymbol{c}_k)} -D_{\mathrm{KL}}\left(p_{\mathrm{ref}}(\boldsymbol{x}|\boldsymbol{c}_k)\|p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\right)+\eta\mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{c}_k)}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\right] - \eta\mathbb{E}_{p(\boldsymbol{x})}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\right] \tag{65}$$

$$= \arg\max_{p_{\boldsymbol{\theta}}(\cdot|\boldsymbol{c}_k)} \int_K \log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)\left(p_{\mathrm{ref}}(\boldsymbol{x}|\boldsymbol{c}_k) + \eta(p(\boldsymbol{x}|\boldsymbol{c}_k) - p(\boldsymbol{x}))\right) \mathrm{d}\boldsymbol{x}. \tag{66}$$

Note that optimization problem (66) shares the same structure as (40), hence by letting $h(\boldsymbol{x}|\boldsymbol{c}) := p_{\mathrm{ref}}(\boldsymbol{x}|\boldsymbol{c}) + \eta(p(\boldsymbol{x}|\boldsymbol{c}) - p(\boldsymbol{x}))$ and under the same compact-support Assumption 1, we get the same optimal solution as in stated in Theorem 1.

Importantly, in the fine-tuning setting, under a mixture error model, MCLR recovers the ground-truth conditional distribution, as stated in the following corollary.

Corollary 1. *If the base model satisfies the mixture error model:*

$$p_{ref}(\boldsymbol{x}|\boldsymbol{c}) = (1-\eta)p(\boldsymbol{x}|\boldsymbol{c}) + \eta p(\boldsymbol{x}), \tag{67}$$

then fine-tuning $p_{ref}(\boldsymbol{x}|\boldsymbol{c})$ with MCLR objective (12) recovers the ground-truth conditional distribution $p(\boldsymbol{x}|\boldsymbol{c})$.

Proof. Under the mixture-error model (67), the optimization problem (66) becomes:

$$\arg\max_{p_{\boldsymbol{\theta}}(\cdot|\boldsymbol{c}_k)} \int_K \log p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)p(\boldsymbol{x}|\boldsymbol{c}_k)d\boldsymbol{x} = \arg\min_{p_{\boldsymbol{\theta}}(\cdot|\boldsymbol{c}_k)} D_{\mathrm{KL}}(p(\boldsymbol{x}|\boldsymbol{c}_k)\|p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k)), \tag{68}$$

which has optimal solution $p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}_k) = p(\boldsymbol{x}|\boldsymbol{c}_k)$. Note that in this case, the proof does not depend on Assumption 1 and Assumption 2. This completes the proof. □

B.4 MCLR for Diffusion Models

Thus far, we have studied MCLR at the likelihood level. Although diffusion models induce likelihoods through their reverse processes, directly evaluating and optimizing these likelihoods is computationally prohibitive during training. We therefore replace each expected negative-log-likelihood term in MCLR with its diffusion ELBO surrogate [Son+21a], resulting in the following score-based objective:

$$\begin{aligned}
&\min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot)} \mathbb{E}_{\boldsymbol{c},t\sim p(t),\boldsymbol{x}\sim p(\cdot|\boldsymbol{c}),\boldsymbol{x}_t\sim p_{0t}(\cdot|\boldsymbol{x})}\left[w(t)\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right] \\
&+\eta\mathbb{E}_{\substack{\boldsymbol{c},\tilde{\boldsymbol{c}},t\sim p(t)\\ \boldsymbol{x}\sim p(\cdot|\boldsymbol{c}),\boldsymbol{x}_t\sim p_{0t}(\cdot|\boldsymbol{x})\\ \boldsymbol{y}\sim p(\cdot|\tilde{\boldsymbol{c}}),\boldsymbol{y}_t\sim p_{0t}(\cdot|\boldsymbol{y})}}\Big[w(t)\Big(\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2 \\
&\qquad\qquad - \|\nabla_{\boldsymbol{y}_t}\log p_{0t}(\boldsymbol{y}_t|\boldsymbol{y}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{y}_t,t,\boldsymbol{c})\|_2^2\Big)\Big],
\end{aligned} \tag{69}$$

and the equivalent denoising form:

$$\begin{aligned}
&\min_{\mathcal{D}_{\boldsymbol{\theta}}(\cdot)} \mathbb{E}_{\boldsymbol{c},t\sim p(t),\boldsymbol{x}\sim p(\cdot|\boldsymbol{c}),\boldsymbol{x}_t\sim p_{0t}(\cdot|\boldsymbol{x})}\left[w(t)\|\boldsymbol{x} - \mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x}_t;\sigma(t),\boldsymbol{c})\|_2^2\right] \\
&+\eta\mathbb{E}_{\substack{\boldsymbol{c},\tilde{\boldsymbol{c}},t\sim p(t)\\ \boldsymbol{x}\sim p(\cdot|\boldsymbol{c}),\boldsymbol{x}_t\sim p_{0t}(\cdot|\boldsymbol{x})\\ \boldsymbol{y}\sim p(\cdot|\tilde{\boldsymbol{c}}),\boldsymbol{y}_t\sim p_{0t}(\cdot|\boldsymbol{y})}}\Big[w(t)\Big(\|\boldsymbol{x} - \mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x}_t;\sigma(t),\boldsymbol{c})\|_2^2 - \|\boldsymbol{y} - \mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{y}_t;\sigma(t),\boldsymbol{c})\|_2^2\Big)\Big],
\end{aligned} \tag{70}$$

where a practical time-sampling distribution $p(t)$ and weighting function $w(t)$ are used (see Section E.1 for the detailed derivation).

It is common in the diffusion-model literature to adapt likelihood-based objectives to score-based models through ELBO-based substitutions [Wal+24; Zhe+25; Xue+25]. Importantly, however, (69) and (70) should be viewed as a diffusion training surrogate induced by likelihood-based MCLR, rather than, in general, an equivalent reformulation of the original likelihood-based MCLR objective.

First, the likelihood-weighted score-matching objective provides a one-sided upper bound on each expected negative log-likelihood. Because MCLR involves the difference between negative-log-likelihood terms, substituting their individual upper bounds does not itself yield a bound on the complete MCLR objective. Second, equivalence requires the diffusion likelihood bound to be tight. A sufficient condition given by [Son+21a] is that there exists a valid conditional distribution $q_0(\boldsymbol{x} \mid \boldsymbol{c})$ whose forward marginals $q_t(\boldsymbol{x}_t \mid \boldsymbol{c})$ satisfy

$$\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c}) \equiv \nabla_{\boldsymbol{x}_t} \log q_t(\boldsymbol{x}_t \mid \boldsymbol{c}), \qquad t \in [0, T], \tag{71}$$

together with the terminal condition $q_T(\cdot \mid \boldsymbol{c}) = \pi$. In particular, each time-indexed vector field must be conservative and induce a normalizable density, while the resulting family of distributions must be consistent with the prescribed forward process (1) through the Fokker–Planck–Kolmogorov equation [SS19]. These cross-time consistency conditions are not explicitly enforced by the ELBO-approximated training objective. Moreover, when $p(t)$ and $w(t)$ differ from the time distribution and likelihood weighting prescribed by the ELBO, the practical objective introduces an additional departure from likelihood-based MCLR. Consequently, rather than inferring its behavior from likelihood-level MCLR, we analyze the diffusion objective in (69) and (70) directly. Intuitively, the contrastive term encourages the denoiser conditioned on $\boldsymbol{c}$ to achieve a lower denoising error on samples from the target condition $\boldsymbol{c}$ than on samples from a mismatched condition $\tilde{\boldsymbol{c}}$. Features shared across conditions tend to benefit both denoising terms similarly and therefore provide limited contrastive advantage. The MCLR regularization term instead encourages the model to exploit structures specific to the target condition. We formalize this intuition below, first by optimizing over unconstrained score functions and then under constrained model parameterizations. Before presenting the main theorems, we introduce the following lemma, which will be useful in their proofs.

Lemma 1. *Let $p(\boldsymbol{x})$ be the clean data distribution, $p_t(\boldsymbol{x})$ be the marginal distribution of $\boldsymbol{x}(t)$, and $p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})$ be the transition density from $\boldsymbol{x}(0)$ to $\boldsymbol{x}(t)$ as defined in Section 2.1, then*

$$\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t) = \mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{x}_t)}\big[\nabla_{\boldsymbol{x}_t} \log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})\big]. \tag{72}$$

Proof.

$$\begin{aligned}
\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t) &= \frac{\nabla_{\boldsymbol{x}_t} p_t(\boldsymbol{x}_t)}{p_t(\boldsymbol{x}_t)} && (73)\\
&= \frac{\nabla_{\boldsymbol{x}_t} \int_{\boldsymbol{x}} p(\boldsymbol{x}) p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) \mathrm{d}\boldsymbol{x}}{p_t(\boldsymbol{x}_t)} && (74)\\
&= \int_{\boldsymbol{x}} \frac{p(\boldsymbol{x}) p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) \nabla_{\boldsymbol{x}_t} p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}{p_t(\boldsymbol{x}_t) p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})} \mathrm{d}\boldsymbol{x} && (75)\\
&= \int_{\boldsymbol{x}} p(\boldsymbol{x}|\boldsymbol{x}_t) \nabla_{\boldsymbol{x}_t} \log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) \mathrm{d}\boldsymbol{x} && (76)\\
&= \mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{x}_t)}\big[\nabla_{\boldsymbol{x}_t} \log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})\big]. && (77)
\end{aligned}$$

This completes the proof. □

B.4.1 Parameterization-Free Analysis under a Pointwise Norm Bound

We first analyze (69) by treating the score function $s_\theta(x_t, t, c)$ as the optimization variable, without imposing any additional assumptions on the model parameterization. However, we do impose the following regularity constraint on the score function, which bounds its maximum norm:

$$\|s_\theta(x_t, t, c)\|_2 \le M, \tag{78}$$

where M is a positive constant. This constraint is required for the diffusion-based MCLR objective (69) and (70) to be well-posed so that a pointwise minimizer exists. Consider a x_t such that:

$$p_t(x_t) \gg p_t(x_t|c), \tag{79}$$

i.e., x_t has a much larger density under the unconditional distribution compared to that under the conditional distribution. At such x_t, the negative term in (69) and (70) dominates the training objective, and the objective value can be pushed arbitrarily small by simply increasing the norm of $s_\theta(x_t, t, c)$ or $\mathcal{D}_\theta(x_t, t, c)$ to be arbitrarily large. Intuitively, MCLR relies on the positive term to prevent the degenerate behavior caused by the negative term, but at x_t where $p_t(x_t) \gg p_t(x_t|c)$, the regularization effect of the positive term becomes too weak and the MCLR loss can be pushed arbitrarily small, while the score function or denoiser fails to learn any useful class-specific structure. Therefore, we impose a pointwise bound on the score magnitude. This constraint also reflects the implicit magnitude control induced in practice by finite network capacity and regularization. Next, we will present the main theorem.

Theorem 2 (full statement). *Recall the signed sum-of-difference function introduced in (13):*

$$h(x|c) := p(x|c) + \eta\big(p(x|c) - p(x)\big). \tag{80}$$

Define its noise-mollified counterpart induced by the diffusion forward process (1) as

$$h_t(x_t|c) := \int_x h(x|c) p_{0t}(x_t|x) \mathrm{d}x \tag{81}$$

$$= p_t(x_t|c) + \eta\big(p_t(x_t|c) - p_t(x_t)\big). \tag{82}$$

Assume that $p_t(x_t|c)$ and $p_t(x_t)$ are continuously differentiable and strictly positive. Consider the population diffusion-based MCLR objective in (69) over the bounded-score function class

$$\mathcal{S}_M := \left\{ s : \|s(x_t, t, c)\|_2 \le M \text{ for all } (x_t, t, c) \right\}, \qquad M > 0. \tag{83}$$

For any fixed (x_t, t, c) satisfying $\nabla_{x_t} h_t(x_t|c) \neq \mathbf{0}$, the pointwise optimal score is

$$s_{\theta^*}(x_t, t, c) = \begin{cases} \nabla_{x_t} \log h_t(x_t|c), & \text{if } h_t(x_t \mid c) > 0 \text{ and } \|\nabla_{x_t} \log h_t(x_t|c)\|_2 \le M, \\ M \dfrac{\nabla_{x_t} h_t(x_t|c)}{\|\nabla_{x_t} h_t(x_t|c)\|_2}, & \text{otherwise.} \end{cases} \tag{84}$$

At a critical point satisfying $\nabla_{x_t} h_t(x_t|c) = \mathbf{0}$, let $\mathcal{A}^*(x_t, t, c)$ denote the set of pointwise minimizers. Then

$$\mathcal{A}^*(x_t, t, c) = \begin{cases} \{\mathbf{0}\}, & \text{if } h_t(x_t|c) > 0, \\ \{s : \|s\|_2 \le M\}, & \text{if } h_t(x_t|c) = 0, \\ \{s : \|s\|_2 = M\}, & \text{if } h_t(x_t|c) < 0. \end{cases} \tag{85}$$

Thus, at a nonpositive critical point, the pointwise MCLR objective does not identify a unique score direction.

Before providing the proof, we first discuss the intuition behind the optimal score field induced by diffusion-based MCLR. Recall that the optimal distribution induced by likelihood-based MCLR, shown in (15), is determined by the sum-of-difference function:

$$h(\boldsymbol{x}|\boldsymbol{c}) = p(\boldsymbol{x}|\boldsymbol{c}) + \eta\big(p(\boldsymbol{x}|\boldsymbol{c}) - p(\boldsymbol{x})\big). \tag{86}$$

Consequently, likelihood-based MCLR amplifies regions where the conditional density is large relative to the unconditional density, while suppressing regions where the unconditional density dominates.

Although diffusion-based MCLR is not equivalent to likelihood-based MCLR, Theorem 2 shows that its optimal score field is governed by the corresponding noise-smoothed function:

$$h_t(\boldsymbol{x}_t|\boldsymbol{c}) = p_t(\boldsymbol{x}_t|\boldsymbol{c}) + \eta\big(p_t(\boldsymbol{x}_t|\boldsymbol{c}) - p_t(\boldsymbol{x}_t)\big). \tag{87}$$

Therefore, the two formulations share the same underlying mechanism: both use the difference between the conditional and unconditional densities to emphasize structures that are characteristic of the target condition. We next explain this connection through the local geometry of the MCLR-induced score field.

We first consider a noncritical point $(\boldsymbol{x}_t, t, \boldsymbol{c})$ satisfying $\nabla_{\boldsymbol{x}_t} h_t(\boldsymbol{x}_t|\boldsymbol{c}) \neq \boldsymbol{0}$. According to (84), the behavior of the optimal score field is determined by the sign of $h_t(\boldsymbol{x}_t|\boldsymbol{c})$ and the norm of $\nabla_{\boldsymbol{x}_t} \log h_t(\boldsymbol{x}_t|\boldsymbol{c})$.

Regime 1: $h_t(\boldsymbol{x}_t|\boldsymbol{c}) > 0$**.** In this setting, if $\|\nabla_{\boldsymbol{x}_t} \log h_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2 \leq M$, the optimal score at $\boldsymbol{x}_t$ is:

$$\begin{aligned} \boldsymbol{s}_{\boldsymbol{\theta}^*}(\boldsymbol{x}_t, t, \boldsymbol{c}) &= \nabla_{\boldsymbol{x}_t} \log h_t(\boldsymbol{x}_t|\boldsymbol{c}) && (88)\\ &= \nabla_{\boldsymbol{x}_t} \log\big(p_t(\boldsymbol{x}_t|\boldsymbol{c}) + \eta(p_t(\boldsymbol{x}_t|\boldsymbol{c}) - p_t(\boldsymbol{x}_t))\big). && (89) \end{aligned}$$

Since $h_t(\boldsymbol{x}_t|\boldsymbol{c}) > 0$, $\nabla_{\boldsymbol{x}_t} \log h_t(\boldsymbol{x}_t|\boldsymbol{c})$ has the same direction as $\nabla_{\boldsymbol{x}_t} h_t(\boldsymbol{x}_t|\boldsymbol{c})$. Therefore, the MCLR-induced score field locally points in the steepest-ascent direction of h_t.

The difference term $\eta(p_t(\boldsymbol{x}_t|\boldsymbol{c}) - p_t(\boldsymbol{x}_t))$ increases the value of h_t relative to $p_t(\boldsymbol{x}_t|\boldsymbol{c})$ in regions where $p_t(\boldsymbol{x}_t|\boldsymbol{c}) > p_t(\boldsymbol{x}_t)$, while decreasing it in regions where $p_t(\boldsymbol{x}_t|\boldsymbol{c}) < p_t(\boldsymbol{x}_t)$. Consequently, the geometry of the MCLR-induced score field favors regions in which the target conditional density has a larger advantage over the unconditional density. Such regions are more characteristic of the target condition and are therefore more likely to contain strong class-specific structures.

Note that as $h_t(\boldsymbol{x}_t|\boldsymbol{c}) \to 0^+$, the norm

$$\|\nabla_{\boldsymbol{x}_t} \log h_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2 = \frac{\|\nabla_{\boldsymbol{x}_t} h_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2}{h_t(\boldsymbol{x}_t|\boldsymbol{c})} \tag{90}$$

can become arbitrarily large unless $\nabla_{\boldsymbol{x}_t} h_t(\boldsymbol{x}_t|\boldsymbol{c})$ approaches zero at a comparable rate, and this large-norm behavior could be an ideal feature. In particular, samples at the zero-level boundary, where $h_t(\boldsymbol{x}_t|\boldsymbol{c}) = 0$, satisfy:

$$p_t(\boldsymbol{x}_t) = (1 + \frac{1}{\eta}) p_t(\boldsymbol{x}_t|\boldsymbol{c}) > p_t(\boldsymbol{x}_t|\boldsymbol{c}). \tag{91}$$

Thus the boundary lies in a region where the unconditional density dominates the target-conditional density, which is not ideal as our goal is to generate samples with strong class-specific structures. MCLR without the maximum-norm constraint avoids this nonideal boundary by strongly pushing

samples near the boundary toward interior regions where $h_t(\boldsymbol{x}_t|\boldsymbol{c})$ has higher values, with a score norm potentially approaching $+\infty$. However, since practical score functions parameterized by deep networks are not likely to have an arbitrarily large norm, we add the maximum-norm constraint (83), with which, the optimal score is still in the same direction as $\nabla_{\boldsymbol{x}_t} \log h_t(\boldsymbol{x}_t|\boldsymbol{c})$, but being clipped to have a finite norm M if $\|\nabla_{\boldsymbol{x}_t} \log h_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2 > M$, yielding:

$$\boldsymbol{s}_{\theta^*}(\boldsymbol{x}_t, t, \boldsymbol{c}) = M \frac{\nabla_{\boldsymbol{x}_t} h_t(\boldsymbol{x}_t|\boldsymbol{c})}{\|\nabla_{\boldsymbol{x}_t} h_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2}, \tag{92}$$

The constraint therefore bounds the score magnitude while preserving its steepest-ascent direction. Consequently, near a regular zero-level boundary, the bounded MCLR score field has a local inward-pointing effect toward the region where $h_t(\boldsymbol{x}_t|\boldsymbol{c}) > 0$.

Regime 2: $h_t(\boldsymbol{x}_t|\boldsymbol{c}) \leq 0$. In this setting, we have:

$$p_t(\boldsymbol{x}_t) \geq (1 + \frac{1}{\eta}) p_t(\boldsymbol{x}_t|\boldsymbol{c}) > p_t(\boldsymbol{x}_t|\boldsymbol{c}). \tag{93}$$

Thus, the unconditional density dominates the target-conditional density by at least the factor $1 + \frac{1}{\eta}$, and such points are nonideal since they are less characteristic of the target condition. At such $\boldsymbol{x}_t$, the optimal MCLR-induced score is:

$$\boldsymbol{s}_{\theta^*}(\boldsymbol{x}_t, t, \boldsymbol{c}) = M \frac{\nabla_{\boldsymbol{x}_t} h_t(\boldsymbol{x}_t|\boldsymbol{c})}{\|\nabla_{\boldsymbol{x}_t} h_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2}, \tag{94}$$

which has a norm M and points toward the steepest-ascent direction that increases $h_t(\boldsymbol{x}_t|\boldsymbol{c})$. Thus, the MCLR-induced score locally points away from regions where the unconditional density strongly dominates the target-conditional density and toward regions where $h_t(\boldsymbol{x}_t|\boldsymbol{c})$ has larger values. In particular, near a regular zero-level boundary, this direction points toward the region where $h_t(\boldsymbol{x}_t|\boldsymbol{c}) > 0$. Therefore, Regimes 1 and 2 exhibit a consistent local geometry: on both sides of the boundary, the MCLR-induced score favors directions that increase the contrastive function h_t, while the norm constraint prevents its magnitude from becoming unbounded.

Behavior at Critical Points. Lastly, we consider critical points satisfying $\nabla_{\boldsymbol{x}_t} h_t(\boldsymbol{x}_t|\boldsymbol{c}) = \boldsymbol{0}$. At such points, the pointwise minimizers are characterized by (85) and depend on the sign of $h_t(\boldsymbol{x}_t|\boldsymbol{c})$.

When $h_t(\boldsymbol{x}_t|\boldsymbol{c}) > 0$, the unique pointwise minimizer is $\boldsymbol{s}_{\theta^*}(\boldsymbol{x}_t, t, \boldsymbol{c}) = \boldsymbol{0}$, which is consistent with the optimal score $\nabla_{\boldsymbol{x}_t} \log h_t(\boldsymbol{x}_t|\boldsymbol{c})$ in (84). When $h_t(\boldsymbol{x}_t|\boldsymbol{c}) = 0$, every vector satisfying $\|\boldsymbol{s}\|_2 \leq M$ is pointwise optimal. When $h_t(\boldsymbol{x}_t|\boldsymbol{c}) < 0$, every vector satisfying $\|\boldsymbol{s}\|_2 = M$ is pointwise optimal.

This degeneracy at nonpositive critical points is a natural consequence of the local geometry underlying MCLR. The direction of the MCLR-induced score is determined by the local gradient of h_t. When this gradient vanishes, h_t provides no first-order information for selecting a preferred direction. Thus, the nonuniqueness is not a failure of the MCLR formulation, but reflects the absence of local directional information at such critical points. In practical implementations, the score at these points is implicitly determined by the shared neural-network parameterization and the surrounding training samples.

Main Implications. Overall, Theorem 2 reveals two complementary mechanisms at noncritical points. Within the positive region of h_t, MCLR augments the standard conditional score with a positively weighted conditional–unconditional log-likelihood-ratio gradient, thereby providing

an adaptive local guidance toward stronger evidence for the target condition. Near or outside the zero-level boundary, the optimal score instead uses the maximum permitted magnitude in the direction that most rapidly increases h_t. Together, these two regimes provide a unified local mechanism through which diffusion-based MCLR favors class-specific structures and suppresses regions dominated by the unconditional distribution. We emphasize that this discussion characterizes the local pointwise geometry of the induced score field. A global characterization of the resulting reverse sampling trajectories and terminal distribution is left for future work.

Now the intuition is clear, and we are ready to provide the proof.

Proof. First, note optimization problem (69) is equivalent to:

$$\min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot)} \ (1+\eta)\,\mathbb{E}_{\boldsymbol{c},t\sim p(t),\,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c}),\,\boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[w(t)\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right] \tag{95}$$

$$-\eta\,\mathbb{E}_{\substack{\boldsymbol{c},\tilde{\boldsymbol{c}},t\sim p(t)\\ \boldsymbol{x}\sim p(\boldsymbol{x}|\tilde{\boldsymbol{c}}),\,\boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}}\left[w(t)\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right]. \tag{96}$$

Note that (96) can be further simplified as:

$$-\eta\,\mathbb{E}_{\boldsymbol{c},t\sim p(t),\,\boldsymbol{x}\sim p(\boldsymbol{x}),\,\boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[w(t)\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right]. \tag{97}$$

Hence (69) is equivalent to

$$\begin{aligned}\min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot)} \ &(1+\eta)\,\mathbb{E}_{\boldsymbol{c},t\sim p(t),\,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c}),\,\boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[w(t)\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right]\\ &-\eta\,\mathbb{E}_{\boldsymbol{c},t\sim p(t),\,\boldsymbol{x}\sim p(\boldsymbol{x}),\,\boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[w(t)\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right].\end{aligned} \tag{98}$$

Since the objective decomposes across t and $\boldsymbol{c}$, the optimization can be performed independently for each pair $(t,\boldsymbol{c})$. Hence, we can fix a specific pair of t and $\boldsymbol{c}$, and optimize the corresponding score function $\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x},t,\boldsymbol{c})$ independently:

$$\begin{aligned}\min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot,t,\boldsymbol{c})} \ &(1+\eta)\,\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c}),\,\boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right]\\ &-\eta\,\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x}),\,\boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right]\end{aligned} \tag{99}$$

,

$$\begin{aligned}\Leftrightarrow \min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot,t,\boldsymbol{c})} \ &(1+\eta)\,\mathbb{E}_{\boldsymbol{x}_t\sim p_t(\boldsymbol{x}_t|\boldsymbol{c}),\,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{x}_t,\boldsymbol{c})}\left[\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right]\\ &-\eta\,\mathbb{E}_{\boldsymbol{x}_t\sim p_t(\boldsymbol{x}_t),\,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{x}_t)}\left[\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right],\end{aligned} \tag{100}$$

where from (99) to (100) we use the Bayes rule:

$$p(\boldsymbol{x}|\boldsymbol{c})p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) = p(\boldsymbol{x},\boldsymbol{x}_t|\boldsymbol{c}) = p_t(\boldsymbol{x}_t|\boldsymbol{c})p(\boldsymbol{x}|\boldsymbol{x}_t,\boldsymbol{c}), \tag{101}$$

and

$$p(\boldsymbol{x})p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) = p(\boldsymbol{x},\boldsymbol{x}_t) = p_t(\boldsymbol{x}_t)p(\boldsymbol{x}|\boldsymbol{x}_t). \tag{102}$$

Note that:

$$\begin{aligned}&\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{x}_t,\boldsymbol{c})}\left[\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right]\\ &\quad=\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2-2\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{x}_t,\boldsymbol{c})}\left[\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})\right]+C_1\\ &\quad=\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2-2\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c})+C_1,\end{aligned} \tag{103}$$

where C_1 is a constant independent of $\boldsymbol{\theta}$, and the second equality follows by applying Lemma 1 to the conditional distribution $p(\boldsymbol{x}|\boldsymbol{c})$. Similarly, we have

$$\begin{aligned}&\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{x}_t)}\left[\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right]\\&=\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2-2\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t)+C_2,\end{aligned}\tag{104}$$

, where C_2 is a constant independent of $\boldsymbol{\theta}$.

Substituting (103) and (104) into (100), the optimization problem becomes equivalent to:

$$\begin{aligned}\min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot,t,\boldsymbol{c})}\int_{\boldsymbol{x}_t}&\big((1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c})-\eta p_t(\boldsymbol{x}_t)\big)\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\\-2\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T&\big((1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c})\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c})-\eta p_t(\boldsymbol{x}_t)\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t)\big)\mathrm{d}\boldsymbol{x}_t,\end{aligned}\tag{105}$$

Note (105) is minimized when every integrand is minimized, hence we fix a $\boldsymbol{x}_t$ and considering the following optimization problem:

$$\begin{aligned}\min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})}&\big((1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c})-\eta p_t(\boldsymbol{x}_t)\big)\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\\-2\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T&\big((1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c})\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c})-\eta p_t(\boldsymbol{x}_t)\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t)\big).\end{aligned}\tag{106}$$

Note that (106) is a quadratic function of the optimization variable $\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})$ and the optimal solution will depend on the sign of $(1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c})-\eta p_t(\boldsymbol{x}_t)$. Therefore, we have the following three different regimes.

Regime 1. Suppose:

$$(1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c})-\eta p_t(\boldsymbol{x}_t)>0\Leftrightarrow h_t(\boldsymbol{x}_t|\boldsymbol{c})>0,\tag{107}$$

then (106) is convex w.r.t. the optimization variable $\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})$, hence the optimal solution is its stationary point:

$$\begin{aligned}\boldsymbol{s}_{\boldsymbol{\theta}^*}(\boldsymbol{x}_t,t,\boldsymbol{c})&=\frac{(1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c})\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c})-\eta p_t(\boldsymbol{x}_t)\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t)}{(1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c})-\eta p_t(\boldsymbol{x}_t)}&(108)\\&=\nabla_{\boldsymbol{x}_t}\log\big((1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c})-\eta p_t(\boldsymbol{x}_t)\big)&(109)\\&=\nabla_{\boldsymbol{x}_t}\log h_t(\boldsymbol{x}_t|\boldsymbol{c})&(110)\end{aligned}$$

Note that when $h_t(\boldsymbol{x}_t|\boldsymbol{c})\approx 0$, the norm of (110) can become arbitrarily large since $h_t(\boldsymbol{x}_t|\boldsymbol{c})$ appears in the denominator of (108), which makes the score function ill-defined. Therefore, the maximum magnitude constraint (78) is required as a regularization. Under the constraint, by completing the square, (106) becomes equivalent to:

$$\begin{aligned}\min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})}\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})-\frac{(1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c})\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c})-\eta p_t(\boldsymbol{x}_t)\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t)}{(1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c})-\eta p_t(\boldsymbol{x}_t)}\|_2^2,&\\ s.t.\qquad\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2\le M,&\\\Leftrightarrow\min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})}\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})-\nabla_{\boldsymbol{x}_t}\log h_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2^2,&\\ s.t.\qquad\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2\le M.&\end{aligned}\tag{111}$$

Because of Cauchy–Schwarz inequality, it is easy to see that the optimal solution to (111) must be of the same direction as $\nabla_{\boldsymbol{x}_t}\log h_t(\boldsymbol{x}_t|\boldsymbol{c})$ but with a norm of M if $\|\nabla_{\boldsymbol{x}_t}\log h_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2>M$:

$$\begin{aligned}\boldsymbol{s}_{\boldsymbol{\theta}^*}(\boldsymbol{x}_t,t,\boldsymbol{c})&=M\frac{\nabla_{\boldsymbol{x}_t}\log h_t(\boldsymbol{x}_t|\boldsymbol{c})}{\|\nabla_{\boldsymbol{x}_t}\log h_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2}&(112)\\&=M\frac{\nabla_{\boldsymbol{x}_t}h_t(\boldsymbol{x}_t|\boldsymbol{c})}{\|\nabla_{\boldsymbol{x}_t}h_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2}.&(113)\end{aligned}$$

On the other hand, if $\|\nabla_{x_t} \log h_t(x_t|c)\|_2 \leq M$, the optimal score function induced by MCLR is:

$$s_{\theta^*}(x_t, t, c) = \nabla_{x_t} \log h_t(x_t|c). \tag{114}$$

Next, we consider the setting where $h_t(x_t|c) = 0$.

Regime 2. Suppose:

$$(1+\eta)p_t(x_t|c) - \eta p_t(x_t) = 0 \Leftrightarrow h_t(x_t|c) = 0. \tag{115}$$

In this case, together with the maximum magnitude constraint (78), (106) becomes a constrained linear optimization problem:

$$\begin{aligned} \min_{s_\theta(x_t,t,c)} \quad & -s_\theta(x_t, t, c)^T \nabla_{x_t} h_t(x_t|c) \\ s.t. \quad & \|s_\theta(x_t, t, c)\|_2 \leq M. \end{aligned} \tag{116}$$

Note that when $\nabla_{x_t} h_t(x_t|c) \neq \mathbf{0}$, the optimal solution to (116) must be of the same direction as the linear coefficient $\nabla_{x_t} h_t(x_t|c)$ and with norm M:

$$s_{\theta^*}(x_t, t, c) = M \frac{\nabla_{x_t} h_t(x_t|c)}{\|\nabla_{x_t} h_t(x_t|c)\|_2}. \tag{117}$$

On the other hand, if $\nabla_{x_t} h_t(x_t|c) = \mathbf{0}$, i.e., at critical point of $h_t(x_t|c)$, the optimal solution is not identifiable, since in this case $s_\theta(x_t, t, c)$ has no effect on the training objective, and its specific value can be any vector in the set $\{s : \|s\|_2 \leq M\}$ such that the maximum magnitude constraint is satisfied, and its specific value will depend on the implicit regularization of the score network. Next, we consider the setting where $h_t(x_t|c) < 0$.

Regime 3. Suppose:

$$(1+\eta)p_t(x_t|c) - \eta p_t(x_t) < 0 \Leftrightarrow h_t(x_t|c) < 0. \tag{118}$$

In this case, together with the maximum magnitude constraint (78), (106) becomes equivalent to the following optimization problem:

$$\begin{aligned} \max_{s_\theta(x_t,t,c)} \|s_\theta(x_t, t, c) - \frac{(1+\eta)p_t(x_t|c)\nabla_{x_t} \log p_t(x_t|c) - \eta p_t(x_t)\nabla_{x_t} \log p_t(x_t)}{(1+\eta)p_t(x_t|c) - \eta p_t(x_t)}\|_2^2 \\ s.t. \quad \|s_\theta(x_t, t, c)\|_2 \leq M, \\ \Leftrightarrow \max_{s_\theta(x_t,t,c)} \|s_\theta(x_t, t, c) - \frac{\nabla_{x_t} h_t(x_t|c)}{h_t(x_t|c)}\|_2^2 \\ s.t. \quad \|s_\theta(x_t, t, c)\|_2 \leq M. \end{aligned} \tag{119}$$

Note that without the maximum magnitude constraint, (119) will become unbounded above, as one can push the quadratic objective arbitrarily large by simply increasing the norm of $s_\theta(x_t, t, c)$. In this case MCLR becomes an ill-posed objective. The core reason behind this is that in this case the density of x_t under the conditional distribution is too small compared to that under the unconditional distribution, such that the negative term in MCLR dominates the overall objective, and the objective can be pushed arbitrarily small by simply increasing the norm of the score function without actually learning class-specific features. This issue is addressed with the maximum magnitude constraint,

and by Cauchy–Schwarz inequality, the optimal solution to (119) must be of opposite direction as $\frac{\nabla_{x_t} h_t(x_t|c)}{h_t(x_t|c)}$. Since $h_t(x_t|c) < 0$, the optimal solution must be of the same direction as $\nabla_{x_t} h_t(x_t|c)$:

$$s_{\theta^*}(x_t, t, c) = M \frac{\nabla_{x_t} h_t(x_t|c)}{\|\nabla_{x_t} h_t(x_t|c)\|_2}. \tag{120}$$

If $\nabla_{x_t} h_t(x_t|c) = \mathbf{0}$, the pointwise objective (119) becomes:

$$\max_{s_\theta(x_t,t,c)} \|s_\theta(x_t, t, c)\|_2^2 \tag{121}$$

$$s.t. \qquad \|s_\theta(x_t, t, c)\|_2 \le M. \tag{122}$$

Therefore, the set of minimizers is:

$$\{s : \|s\|_2 = M\}. \tag{123}$$

Combining the results from all three regimes completes the proof. □

B.4.2 Extension to Fine-Tuning from a Pretrained Score Model

Theorem 2 can be naturally extended to the fine-tuning setting where a pretrained base diffusion model denoted by $\nabla_{x_t} \log p_{\text{ref},t}(x_t|c)$ is available, in which case the diffusion-based MCLR objective becomes:

$$\begin{aligned} &\min_{s_\theta(\cdot)} \mathbb{E}_{c,t\sim p(t),x\sim p_{\text{ref}}(\cdot|c),x_t\sim p_{0t}(\cdot|x)} \left[w(t) \|\nabla_{x_t} \log p_{\text{ref},t}(x_t|c) - s_\theta(x_t, t, c)\|_2^2 \right] \\ &+\eta \mathbb{E}_{\substack{c,\tilde{c},t\sim p(t) \\ x\sim p(\cdot|c),x_t\sim p_{0t}(\cdot|x) \\ y\sim p(\cdot|\tilde{c}),y_t\sim p_{0t}(\cdot|y)}} \Big[w(t) \Big(\|\nabla_{x_t} \log p_{0t}(x_t|x) - s_\theta(x_t, t, c)\|_2^2 \\ &\qquad - \|\nabla_{y_t} \log p_{0t}(y_t|y) - s_\theta(y_t, t, c)\|_2^2 \Big) \Big]. \end{aligned} \tag{124}$$

In this setting, the optimal solution to (124) shares exactly the same structure as that of Theorem 2, except that the contrastive function $h_t(x_t|c)$ becomes:

$$h_t(x_t|c) = p_{\text{ref,t}}(x_t|c) + \eta(p_t(x_t|c) - p_t(x_t)). \tag{125}$$

Consequently, the optimal score field induced by MCLR in the fine-tuning setting pushes samples toward regions with higher value under $p_{\text{ref},t}(x_t|c) + \eta(p_t(x_t|c) - p_t(x_t))$. We omit the full proof since it largely follows the same structure as the proof of Theorem 2. In particular, the pointwise optimization problem in (106) becomes:

$$\begin{aligned} \min_{s_\theta(x_t,t,c)} &\big(p_{\text{ref},t}(x_t|c) + \eta p_t(x_t|c) - \eta p_t(x_t)\big) \|s_\theta(x_t, t, c)\|_2^2 \\ &-2 s_\theta(x_t, t, c)^T \big(p_{\text{ref},t}(x_t|c) \nabla_{x_t} \log p_{\text{ref,t}}(x_t|c) + \\ &\eta p_t(x_t|c) \nabla_{x_t} \log p_t(x_t|c) - \eta p_t(x_t) \nabla_{x_t} \log p_t(x_t)\big), \end{aligned} \tag{126}$$

and the same analysis applies.

Importantly, if the base model satisfies the mixture error model assumption (16), then fine-tuning with diffusion-based MCLR objective recovers the correct conditional distribution, as stated by the following corollary:

Corollary 2. *If the base model p_{ref} satisfies the mixture error model:*

$$p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c}) = (1-\eta)\, p(\boldsymbol{x}|\boldsymbol{c}) + \eta\, p(\boldsymbol{x}), \tag{127}$$

where $\eta \in [0,1]$, then fine-tuning $p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})$ with diffusion-based MCLR objective (124) recovers the ground-truth conditional score

$$\boldsymbol{s}_{\theta^*}(\boldsymbol{x}_t, t, \boldsymbol{c}) = \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}), \qquad t \in (0, T]. \tag{128}$$

Consequently, exact reverse sampling recovers $p(\boldsymbol{x}|\boldsymbol{c})$ in the limit $t \downarrow 0$. Moreover, in this case the pointwise norm bound is not required.

Proof. Due to the linearity of forward noising process, we have:

$$p_{\text{ref},t}(\boldsymbol{x}_t|\boldsymbol{c}) = (1-\eta)\, p_t(\boldsymbol{x}_t|\boldsymbol{c}) + \eta\, p_t(\boldsymbol{x}_t). \tag{129}$$

Substituting (129) into (126), the pointwise optimization problem becomes:

$$\begin{aligned}
&\min_{\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c})} p_t(\boldsymbol{x}_t|\boldsymbol{c})\|\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2 \\
&-2\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c})^T\big(\nabla_{\boldsymbol{x}_t} p_{\text{ref},t}(\boldsymbol{x}_t|\boldsymbol{c}) + \eta p_t(\boldsymbol{x}_t|\boldsymbol{c})\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c}) - \eta p_t(\boldsymbol{x}_t)\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t)\big), \\
&\Leftrightarrow \min_{\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c})} p_t(\boldsymbol{x}_t|\boldsymbol{c})\|\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2 - 2\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c})^T\nabla_{\boldsymbol{x}_t} p_t(\boldsymbol{x}_t|\boldsymbol{c}), \\
&\Leftrightarrow \min_{\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c})} \|\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2 - 2\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c})^T\frac{\nabla_{\boldsymbol{x}_t} p_t(\boldsymbol{x}_t|\boldsymbol{c})}{p_t(\boldsymbol{x}_t|\boldsymbol{c})}, \\
&\Leftrightarrow \min_{\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c})} \|\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2 - 2\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c})^T\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c}), \\
&\Leftrightarrow \min_{\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c})} \|\boldsymbol{s}_\theta(\boldsymbol{x}_t,t,\boldsymbol{c}) - \nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2^2.
\end{aligned} \tag{130}$$

Therefore the unique optimal solution is

$$\boldsymbol{s}_{\theta^*}(\boldsymbol{x}_t, t, \boldsymbol{c}) = \nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c}) \tag{131}$$

for every $t \in (0, T]$. Consequently, reverse sampling with this score recovers $p(\boldsymbol{x}|\boldsymbol{c})$ as $t \downarrow 0$. This completes the proof. □

Note that although likelihood-based MCLR and diffusion-based MCLR are not equivalent in general, this corollary shows that both of them correct the same mixture-error mode at their respective population optima.

B.4.3 Connection to Classifier-Free Guidance

We now connect classifier-free guidance (CFG) to diffusion-based MCLR from two perspectives. The first perspective is that the score-field induced by diffusion-based MCLR can be interpreted as CFG with a sample-adaptive weight, while the second perspective is that the CFG-guided score is the unique minimizer of a sample-adaptive weighted ELBO-approximated MCLR objective.

Perspective 1: MCLR as CFG with a Sample-Adaptive Weight. We first restrict attention to Regime 1, where $h_t(\boldsymbol{x}_t|\boldsymbol{c}) > 0$, and consider the unconstrained optimum, or equivalently, the constrained optimum at points satisfying $\|\nabla_{\boldsymbol{x}_t} \log h_t(\boldsymbol{x}_t|\boldsymbol{c})\|_2 \le M$ in Theorem 2. In this case, the MCLR-induced score is:

$$\begin{aligned}
\nabla_{\boldsymbol{x}_t} \log h_t(\boldsymbol{x}_t|\boldsymbol{c}) &= \nabla_{\boldsymbol{x}_t} \log\big(p_t(\boldsymbol{x}_t|\boldsymbol{c}) + \eta(p_t(\boldsymbol{x}_t|\boldsymbol{c}) - p_t(\boldsymbol{x}_t))\big) && (132)\\
&= \frac{(1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c})\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}) - \eta p_t(\boldsymbol{x}_t)\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t)}{(1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c}) - \eta p_t(\boldsymbol{x}_t)} && (133)\\
&= \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}) + \frac{\eta p_t(\boldsymbol{x}_t)}{(1+\eta)p_t(\boldsymbol{x}_t|\boldsymbol{c}) - \eta p_t(\boldsymbol{x}_t)}\big(\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}) - \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t)\big) && (134)\\
&= \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}) + \frac{1}{(1+\frac{1}{\eta})\frac{p_t(\boldsymbol{x}_t|\boldsymbol{c})}{p_t(\boldsymbol{x}_t)} - 1}\big(\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}) - \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t)\big) && (135)\\
&= \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}) + \alpha_t(\eta, \boldsymbol{x}_t)\big(\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}) - \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t)\big), && (136)
\end{aligned}$$

where we denote $\frac{1}{(1+\frac{1}{\eta})\frac{p_t(\boldsymbol{x}_t|\boldsymbol{c})}{p_t(\boldsymbol{x}_t)}-1}$ as $\alpha_t(\eta, \boldsymbol{x}_t)$ for notational simplicity. Since we restrict our discussion in Regime 1 where $h_t(\boldsymbol{x}_t|\boldsymbol{c}) > 0$, we have $\alpha_t(\eta, \boldsymbol{x}_t) > 0$, implying the MCLR-induced score fields add the conditional-unconditional score difference $\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}) - \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t)$, weighted by $\alpha_t(\eta, \boldsymbol{x}_t)$, to the standard conditional score $\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c})$, which shares a similar mechanism as classifier-free guidance (CFG). However, unlike in standard CFG where a constant guidance strength is used, here the guidance term is weighted by a sample-adaptive weight $\alpha_t(\eta, \boldsymbol{x}_t)$ that monotonically decreases with the likelihood-ratio $\frac{p_t(\boldsymbol{x}_t|\boldsymbol{c})}{p_t(\boldsymbol{x}_t)}$. This makes sense as $\boldsymbol{x}_t$ with high likelihood-ratio implies that such $\boldsymbol{x}_t$ is less likely to be confused with a mismatched condition, i.e., this particular sample is already well-aligned with the target condition, therefore there is no need to further apply the guidance. On the other hand, a sample $\boldsymbol{x}_t$ with low likelihood-ratio implies such $\boldsymbol{x}_t$ lies in a region providing weak relative evidence for the target condition, hence the MCLR-induced score adds a large guidance to steer the generation trajectory away from such $\boldsymbol{x}_t$.

Perspective 2: CFG as MCLR with Sample-Adaptive Negative Weighting. It turns out the CFG-guided score is the optimal solution to a sample-adaptive weighted MCLR as stated by the following theorem.

Theorem 3. *For any time sampling distribution $p(t)$ and weighting function $w(t) > 0$, the CFG-guided score*

$$\boldsymbol{s}_{cfg}(\boldsymbol{x}_t, t, \boldsymbol{c}) := \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}) + \eta\left(\nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t|\boldsymbol{c}) - \nabla_{\boldsymbol{x}_t} \log p_t(\boldsymbol{x}_t)\right)$$

is the unique minimizer of a sample-adaptive weighted ELBO-approximated MCLR objective:

$$\begin{aligned}
\boldsymbol{s}_{cfg}(\cdot) = \underset{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot)}{arg\ min}\ & \mathbb{E}_{\boldsymbol{c}, t\sim p(t),\, \boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c}),\, \boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\Big[w(t)\|\nabla_{\boldsymbol{x}_t} \log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c})\|_2^2\Big] \\
& + \eta\, \mathbb{E}_{\substack{\boldsymbol{c}, \tilde{\boldsymbol{c}}, t\sim p(t), \boldsymbol{x}\sim p(\cdot|\boldsymbol{c}), \boldsymbol{y}\sim p(\cdot|\tilde{\boldsymbol{c}}) \\ p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}), p_{0t}(\boldsymbol{y}_t|\boldsymbol{y})}}\Bigg[w(t)\Big(\|\nabla_{\boldsymbol{x}_t} \log p_{0t}(\boldsymbol{x}_t \mid \boldsymbol{x}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c})\|_2^2 \\
& \qquad - \frac{p_t(\boldsymbol{y}_t|\boldsymbol{c})}{p_t(\boldsymbol{y}_t)}\|\nabla_{\boldsymbol{y}_t} \log p_{0t}(\boldsymbol{y}_t \mid \boldsymbol{y}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{y}_t, t, \boldsymbol{c})\|_2^2\Big)\Bigg],
\end{aligned} \tag{137}$$

where $p_t(\boldsymbol{y}_t) := \mathbb{E}_{\boldsymbol{c}}[p_t(\boldsymbol{y}_t|\boldsymbol{c})]$.

Proof. First, note that optimization problem (137) is equivalent to

$$\begin{aligned}\min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot)} \; & (1+\eta)\,\mathbb{E}_{\boldsymbol{c},t\sim p(t),\,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c}),\,\boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[w(t)\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right] \\ & -\eta\,\mathbb{E}_{\boldsymbol{c},t\sim p(t),\,\boldsymbol{x}\sim p(\boldsymbol{x}),\,\boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[w(t)\frac{p_t(\boldsymbol{x}_t|\boldsymbol{c})}{p_t(\boldsymbol{x}_t)}\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right].\end{aligned} \tag{138}$$

In order to solve (138), since the objective decomposes across t and $\boldsymbol{c}$, the optimization can be performed independently for each pair $(t,\boldsymbol{c})$. Hence, we can fix a specific pair of t and $\boldsymbol{c}$, and optimize the corresponding score $\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x},t,\boldsymbol{c})$ independently:

$$\begin{aligned}\min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot,t,\boldsymbol{c})} \; & (1+\eta)\,\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c}),\,\boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right] \\ & -\eta\,\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x}),\,\boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[\frac{p_t(\boldsymbol{x}_t|\boldsymbol{c})}{p_t(\boldsymbol{x}_t)}\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right],\end{aligned} \tag{139}$$

$$\begin{aligned}\Leftrightarrow \min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot,t,\boldsymbol{c})} \; & (1+\eta)\,\mathbb{E}_{\boldsymbol{x}_t\sim p_t(\boldsymbol{x}_t|\boldsymbol{c}),\,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{x}_t,\boldsymbol{c})}\left[\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right] \\ & -\eta\,\mathbb{E}_{\boldsymbol{x}_t\sim p_t(\boldsymbol{x}_t),\,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{x}_t)}\left[\frac{p_t(\boldsymbol{x}_t|\boldsymbol{c})}{p_t(\boldsymbol{x}_t)}\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right]\end{aligned} \tag{140}$$

$$\begin{aligned}\Leftrightarrow \min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot,t,\boldsymbol{c})} \; & (1+\eta)\,\mathbb{E}_{\boldsymbol{x}_t\sim p_t(\boldsymbol{x}_t|\boldsymbol{c}),\,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{x}_t,\boldsymbol{c})}\left[\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right] \\ & -\eta\,\mathbb{E}_{\boldsymbol{x}_t\sim p_t(\boldsymbol{x}_t|\boldsymbol{c}),\,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{x}_t)}\left[\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right].\end{aligned} \tag{141}$$

From (139) to (140) we use the Bayes rule: $p(\boldsymbol{x}|\boldsymbol{c})p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) = p(\boldsymbol{x},\boldsymbol{x}_t|\boldsymbol{c}) = p_t(\boldsymbol{x}_t|\boldsymbol{c})p(\boldsymbol{x}|\boldsymbol{x}_t,\boldsymbol{c})$. From (140) to (141) we use the identity: $p_t(\boldsymbol{x}_t)p(\boldsymbol{x}|\boldsymbol{x}_t)\frac{p_t(\boldsymbol{x}_t|\boldsymbol{c})}{p_t(\boldsymbol{x}_t)} = p_t(\boldsymbol{x}_t|\boldsymbol{c})p(\boldsymbol{x}|\boldsymbol{x}_t)$.

Note that:

$$\begin{aligned}&\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{x}_t,\boldsymbol{c})}\left[\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right] \\ &\quad=\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2-2\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{x}_t,\boldsymbol{c})}\left[\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})\right]+C_1 \\ &\quad=\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2-2\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c})+C_1,\end{aligned} \tag{142}$$

where C_1 is a constant independent of $\boldsymbol{\theta}$, and the second equality follows from Lemma 1. Similarly, we have

$$\begin{aligned}&\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{x}_t)}\left[\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})-\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right] \\ &\quad=\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2-2\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t)+C_2,\end{aligned} \tag{143}$$

, where C_2 is a constant independent of $\boldsymbol{\theta}$.

Substituting (142) and (143) into (141), and letting $\boldsymbol{v}_1 = (1+\eta)\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c}) - \eta\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t)$, the optimization problem becomes equivalent to:

$$\min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot,t,\boldsymbol{c})} \; \mathbb{E}_{\boldsymbol{x}_t\sim p_t(\boldsymbol{x}_t|\boldsymbol{c})}\left[\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2-2\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T\boldsymbol{v}_1\right] \tag{144}$$

$$\Leftrightarrow \min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot,t,\boldsymbol{c})} \; \mathbb{E}_{\boldsymbol{x}_t\sim p_t(\boldsymbol{x}_t|\boldsymbol{c})}\left[\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})-\boldsymbol{s}_{\text{cfg}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\right]+C_3, \tag{145}$$

where C_3 is a constant independent of $\boldsymbol{\theta}$. Therefore, the optimal solution is:

$$\boldsymbol{s}_{\boldsymbol{\theta}^*}(\boldsymbol{x}_t,t,\boldsymbol{c}) = \boldsymbol{s}_{\text{cfg}}(\boldsymbol{x}_t,t,\boldsymbol{c}) \tag{146}$$

$$= \nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c}) + \eta\big(\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c}) - \nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t)\big). \tag{147}$$

This completes the proof. □

Note that $\boldsymbol{y}_t$ is sampled from the unconditional distribution, hence a larger weight $\frac{p_t(\boldsymbol{y}_t|\boldsymbol{c})}{p_t(\boldsymbol{y}_t)}$ indicates that it is more likely to be confused with samples from the target condition $\boldsymbol{c}$. Hence, the adaptive weight makes the CFG-equivalent objective (23) more selective over negative samples: samples $\boldsymbol{y}_t$ from mismatched conditions that have high relative density under the target condition $\boldsymbol{c}$ receive larger weights, while samples that are already easily distinguishable from the target condition (those with small $\frac{p_t(\boldsymbol{y}_t|\boldsymbol{c})}{p_t(\boldsymbol{y}_t)}$) receive smaller weights. Thus, the adaptive weight can be interpreted as a form of hard-negative mining, emphasizing ambiguous or class-confusable samples that provide stronger contrastive signal.

Extensions: CFG Variants under the Alignment Framework. The equivalence between CFG and weighted MCLR provides a unified perspective for interpreting CFG variants. This perspective is formalized in the following corollary.

Corollary 3. *Consider two distributions $p^+(\boldsymbol{x})$ and $p^-(\boldsymbol{x})$. For any time sampling distribution $p(t)$ and weighting function $w(t) > 0$, a generalized CFG-style score:*

$$\boldsymbol{s}_{\textit{cfg-variant}}(\boldsymbol{x}_t, t) := \nabla_{\boldsymbol{x}_t} \log p_t^+(\boldsymbol{x}_t) + \eta \left(\nabla_{\boldsymbol{x}_t} \log p_t^+(\boldsymbol{x}_t) - \nabla_{\boldsymbol{x}_t} \log p_t^-(\boldsymbol{x}_t)\right) \tag{148}$$

is the unique minimizer of the following MCLR-style optimization problem:

$$\begin{aligned} \min_{\boldsymbol{s}_{\boldsymbol{\theta}}(\cdot)} \; & (1+\eta)\, \mathbb{E}_{t\sim p(t),\, \boldsymbol{x}\sim p^+(\boldsymbol{x}),\, \boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[w(t)\|\nabla_{\boldsymbol{x}_t} \log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t)\|_2^2\right] \\ & - \eta\, \mathbb{E}_{t\sim p(t),\, \boldsymbol{x}\sim p^-(\boldsymbol{x}),\, \boldsymbol{x}_t\sim p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[w(t)\frac{p_t^+(\boldsymbol{x}_t)}{p_t^-(\boldsymbol{x}_t)}\|\nabla_{\boldsymbol{x}_t} \log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t)\|_2^2\right] \end{aligned} \tag{149}$$

The proof is omitted, as it follows directly by adapting the proof of Theorem 3. The optimization problem (149) shares the same structure as the ELBO-approximated weighted MCLR (138), with the only difference that positive samples are drawn from $p^+(\boldsymbol{x})$ and negative samples from $p^-(\boldsymbol{x})$. This formulation unifies a broad class of CFG-style methods, as illustrated below.

Standard CFG. Let $p^+(\boldsymbol{x}) = p(\boldsymbol{x}|\boldsymbol{c})$ and $p^-(\boldsymbol{x}) = p(\boldsymbol{x})$. This recovers the standard CFG and corresponds to the weighted MCLR formulation discussed in previous sections.

Autoguidance. Let $p^+(\boldsymbol{x})$ be the distribution induced by a strong diffusion model and $p^-(\boldsymbol{x})$ the distribution induced by a weaker model. We recover the Autoguidance [Kar+24a].

Inference-Time Alignment Guidance. Let $p^+(\boldsymbol{x})$ be the distribution induced by a diffusion model fine-tuned on high-reward data and $p^-(\boldsymbol{x})$ the distribution induced by the base model (or a low-reward fine-tuned model). We recover a family of inference-time alignment methods [Fra+25; Jin+25; Che+25c; JWL26].

Lastly, note that (149) without adaptive weighting is the ELBO-based approximation of the following likelihood-based objective, which is presented in Figure 3:

$$\max_{\boldsymbol{\theta}} \; \mathbb{E}_{p^+(\boldsymbol{x})}\left[\log p_{\boldsymbol{\theta}}(\boldsymbol{x})\right] + \eta \underbrace{\mathbb{E}_{\boldsymbol{x}\sim p^+(\cdot),\boldsymbol{y}\sim p^-(\cdot)}\left[\log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x})}{p_{\boldsymbol{\theta}}(\boldsymbol{y})}\right]}_{\text{MCLR Regularization}}. \tag{150}$$

B.4.4 Population-Level Simulation under the Mixture-Error Model

We provide a population-level synthetic simulation to illustrate the effects of MCLR and CFG in the diffusion-model setting. The ground-truth conditional distributions consist of four equally likely

two-dimensional isotropic Gaussians with standard deviation 0.5, whose means are evenly placed on a circle of radius 2.6. The unconditional distribution is their uniform mixture. We construct a weak base model satisfying the mixture-error assumption:

$$p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c}) = (1-\rho)p(\boldsymbol{x}|\boldsymbol{c}) + \rho p(\boldsymbol{x}), \tag{151}$$

where $\rho = 0.75$.

We evaluate both methods using their exact population score fields. For MCLR, we fix the score bound to $M = 5$ and vary only the contrastive strength η. We use the variance-exploding forward process $\boldsymbol{x}_t = \boldsymbol{x} + 12t\boldsymbol{\epsilon}$ and solve the corresponding probability-flow ODE from $t = 1$ to $t = 0$ using 400 Euler steps. We generate 1,000 samples per class and report the class-averaged sliced 2-Wasserstein distance (SWD), estimated using 256 random directions.

As shown in Figure 6, both MCLR and CFG reduce the inter-class leakage of the weak base model as their respective strengths increase. At the same time, overly strong contrastive or guidance effects can alter the within-class geometry of the generated distributions, demonstrating a shared trade-off between class specificity and distributional fidelity. Under the assumed mixture-error model, setting $\eta = \rho$ cancels the unconditional contamination and yields the closest recovery of the target conditional distributions.

B.5 Comparison with Contrastive Flow Matching

Basics of Flow Matching. In this section we study Contrastive Flow Matching [Sto+25] (CFM), a method that enhances class-specific features but with a different mechanism from our proposed MCLR. Unlike MCLR, CFM is derived in the flow matching framework. Specifically, flow matching defines the forward process as:

$$\boldsymbol{x}_t = \alpha_t \boldsymbol{x} + \beta_t \boldsymbol{\epsilon}, \tag{152}$$

where $\boldsymbol{\epsilon} \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$, and aims to learn a velocity field over a probability-flow ODE, $\mathrm{d}\boldsymbol{x}_t = \boldsymbol{v}(\boldsymbol{x}_t, t)\mathrm{d}t$, whose distribution at time t is the marginal distribution $p_t(\boldsymbol{x}_t)$ induced by (152). Such a velocity field can be approximated by minimizing the following flow matching loss:

$$\mathcal{L}_{\text{FM}} := \mathbb{E}_{\boldsymbol{c},t,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c}),\boldsymbol{\epsilon}\sim\mathcal{N}(\mathbf{0},\boldsymbol{I}),\boldsymbol{x}_t=\alpha_t\boldsymbol{x}+\beta_t\boldsymbol{\epsilon}}\left[\|\boldsymbol{v}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c}) - (\dot{\alpha}_t\boldsymbol{x} + \dot{\beta}_t\boldsymbol{\epsilon})\|_2^2\right], \tag{153}$$

where $\boldsymbol{v}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c})$ is known as the flow model. Define the transition probability corresponding to (152) as:

$$p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) = \mathcal{N}(\boldsymbol{x}_t; \alpha_t\boldsymbol{x}, \beta_t^2\boldsymbol{I}). \tag{154}$$

It can be shown that [HE26]:

$$\dot{\alpha}_t\boldsymbol{x} + \dot{\beta}_t\boldsymbol{\epsilon} = a_t\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) + b_t\boldsymbol{x}_t, \tag{155}$$

where:

$$a_t = (\beta_t^2\frac{\dot{\alpha}_t}{\alpha_t} - \dot{\beta}_t\beta_t), \qquad b_t = \frac{\dot{\alpha}_t}{\alpha_t}. \tag{156}$$

Therefore, by parameterizing the flow model as:

$$\boldsymbol{v}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c}) = a_t\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c}) + b_t\boldsymbol{x}_t, \tag{157}$$

the flow matching loss (153) becomes equivalent to the denoising score matching loss in (3). In fact, as shown in [HE26], the velocity field is equivalent to the score function up to the same linear transformation:

$$\boldsymbol{v}(\boldsymbol{x}_t, t, \boldsymbol{c}) = a_t\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c}) + b_t\boldsymbol{x}_t, \tag{158}$$

hence flow matching is equivalent to denoising score matching.

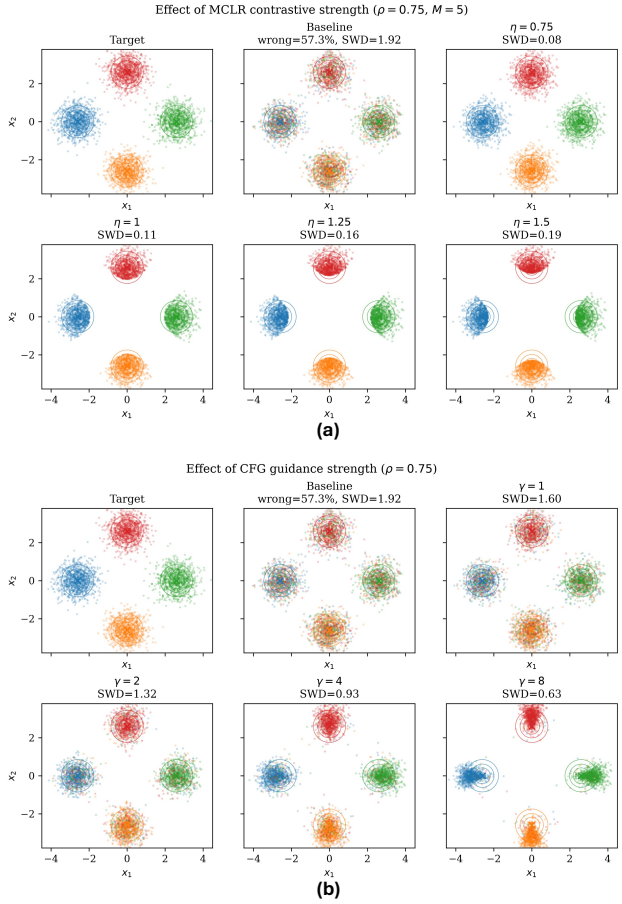


Figure 6: **Population-level comparison of MCLR and CFG.** "Target" shows samples from the ground-truth conditional distributions, whereas "Baseline" shows samples from the weak reference model under the mixture-error setting with $\rho = 0.75$. Colors denote conditioning classes. "wrong" is the fraction of samples whose most likely ground-truth class differs from the conditioning class, and SWD is the class-averaged sliced 2-Wasserstein distance to the target conditionals (lower is better). (a) MCLR with fixed $M = 5$ and varying η. (b) CFG with varying γ. Both methods reduce inter-class leakage, while large strengths can alter the within-class geometry.

Contrastive Flow Matching. Contrastive Flow Matching adds an additional contrastive term upon (153), resulting in:

$$\begin{aligned}\mathcal{L}_{\text{CFM}} = \mathbb{E}_{\boldsymbol{c},t,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c}),\boldsymbol{\epsilon}_1\sim\mathcal{N}(\boldsymbol{0},\boldsymbol{I}),\boldsymbol{x}_t=\alpha_t\boldsymbol{x}+\beta_t\boldsymbol{\epsilon}_1}\left[\|\boldsymbol{v}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})-(\dot{\alpha}_t\boldsymbol{x}+\dot{\beta}_t\boldsymbol{\epsilon}_1)\|_2^2\right]\\ +\eta\,\mathbb{E}_{\substack{\boldsymbol{c},\tilde{\boldsymbol{c}},t,\boldsymbol{\epsilon}_2\sim\mathcal{N}(\boldsymbol{0},\boldsymbol{I}),\boldsymbol{\epsilon}_3\sim\mathcal{N}(\boldsymbol{0},\boldsymbol{I}),\boldsymbol{x}\sim p(\cdot|\boldsymbol{c}),\boldsymbol{y}\sim p(\cdot|\tilde{\boldsymbol{c}})\\ \boldsymbol{x}_t=\alpha_t\boldsymbol{x}+\beta_t\boldsymbol{\epsilon}_2,\boldsymbol{y}_t=\alpha_t\boldsymbol{y}+\beta_t\boldsymbol{\epsilon}_3}}\left[\|\boldsymbol{v}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})-(\dot{\alpha}_t\boldsymbol{x}+\dot{\beta}_t\boldsymbol{\epsilon}_2)\|_2^2\right.\\ \left.-\|\boldsymbol{v}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})-(\dot{\alpha}_t\boldsymbol{y}+\dot{\beta}_t\boldsymbol{\epsilon}_3)\|_2^2\right].\end{aligned} \tag{159}$$

Since we derive MCLR in terms of the score function, for better comparison, we first convert the contrastive flow matching objective into its equivalent score function form by substituting (155)

and (157) into (159), resulting in the following optimization problem:

$$\begin{aligned}
\min_{\boldsymbol{\theta}} \mathcal{L}_{\text{CFM}} := \mathbb{E}_{\boldsymbol{c},t,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c}),p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[\|a_t \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c}) - a_t \nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})\|_2^2\right] \\
+\eta\,\mathbb{E}_{\substack{\boldsymbol{c},\tilde{\boldsymbol{c}},t,\boldsymbol{x}\sim p(\cdot|\boldsymbol{c}),\boldsymbol{y}\sim p(\cdot|\tilde{\boldsymbol{c}}) \\ p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}),p_{0t}(\boldsymbol{y}_t|\boldsymbol{y})}}\left[\|a_t \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c}) - a_t \nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})\|_2^2\right. \\
\left.-\|a_t \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c}) - a_t \nabla_{\boldsymbol{y}_t}\log p_{0t}(\boldsymbol{y}_t|\boldsymbol{y}) + b_t(\boldsymbol{x}_t - \boldsymbol{y}_t)\|_2^2\right], \\
\Leftrightarrow \min_{\boldsymbol{\theta}} \mathbb{E}_{\boldsymbol{c},t,\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c}),p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c}) - \nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})\|_2^2\right] \\
+\eta\,\mathbb{E}_{\substack{\boldsymbol{c},\tilde{\boldsymbol{c}},t,\boldsymbol{x}\sim p(\cdot|\boldsymbol{c}),\boldsymbol{y}\sim p(\cdot|\tilde{\boldsymbol{c}}) \\ p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}),p_{0t}(\boldsymbol{y}_t|\boldsymbol{y})}}\left[\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c}) - \nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})\|_2^2\right. \\
\left.-\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c}) - \nabla_{\boldsymbol{y}_t}\log p_{0t}(\boldsymbol{y}_t|\boldsymbol{y}) + \frac{b_t}{a_t}(\boldsymbol{x}_t - \boldsymbol{y}_t)\|_2^2\right].
\end{aligned} \tag{160}$$

From (160), we see that CFM is similar to MCLR in (69), as both add an additional negative term to the standard denoising score matching objective. However, unlike MCLR, the negative term in CFM does not directly correspond to a score matching target. To better understand CFM's effect, we derive the optimal solution to $\mathcal{L}_{\text{CFM}}$ in the following theorem.

Theorem 5. *The optimal solution to the CFM objective* (160) *is:*

$$\boldsymbol{s}_{\boldsymbol{\theta}^*}(\boldsymbol{x}_t,t,\boldsymbol{c}) = (1+\eta)\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c}) + \eta\frac{b_t}{a_t}(\boldsymbol{x}_t - \mathbb{E}_{p_t(\boldsymbol{x}_t)}[\boldsymbol{x}_t]). \tag{161}$$

Proof. Note that:

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{x}\sim p(\boldsymbol{x}|\boldsymbol{c}),p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c}) - \nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})\|_2^2\right] && (162)\\
&= \mathbb{E}_{\boldsymbol{x}_t\sim p_t(\boldsymbol{x}_t|\boldsymbol{c})}\left[\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2 - 2\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T\nabla_{\boldsymbol{x}_t}\log p_t(\boldsymbol{x}_t|\boldsymbol{c})\right] + C_1, && (163)
\end{aligned}$$

where C_1 is a constant independent of $\boldsymbol{\theta}$.

Moreover, we have:

$$\begin{aligned}
&\mathbb{E}_{\boldsymbol{x}\sim p(\cdot|\boldsymbol{c}),\boldsymbol{y}\sim p(\cdot),p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}),p_{0t}(\boldsymbol{y}_t|\boldsymbol{y})}\left[\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c}) - \nabla_{\boldsymbol{y}_t}\log p_{0t}(\boldsymbol{y}_t|\boldsymbol{y}) + \frac{b_t}{a_t}(\boldsymbol{x}_t - \boldsymbol{y}_t)\|_2^2\right] && (164)\\
&= \mathbb{E}_{\boldsymbol{x}\sim p(\cdot|\boldsymbol{c}),\boldsymbol{y}\sim p(\cdot),p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}),p_{0t}(\boldsymbol{y}_t|\boldsymbol{y})}\left[\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2 - 2\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T\nabla_{\boldsymbol{y}_t}\log p_{0t}(\boldsymbol{y}_t|\boldsymbol{y})\right. && (165)\\
&\qquad\left. + 2\frac{b_t}{a_t}\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T(\boldsymbol{x}_t - \boldsymbol{y}_t)\right] + C_2 && (166)\\
&= \mathbb{E}_{\boldsymbol{x}_t\sim p_t(\boldsymbol{x}_t|\boldsymbol{c}),\boldsymbol{y}_t\sim p_t(\boldsymbol{y}_t)}\left[\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2 - 2\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T\nabla_{\boldsymbol{y}_t}\log p_t(\boldsymbol{y}_t)\right. && (167)\\
&\qquad\left. + 2\frac{b_t}{a_t}\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T(\boldsymbol{x}_t - \boldsymbol{y}_t)\right] + C_2 && (168)\\
&= \mathbb{E}_{\boldsymbol{x}_t\sim p_t(\boldsymbol{x}_t|\boldsymbol{c})}\left[\|\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2 + 2\frac{b_t}{a_t}\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})^T(\boldsymbol{x}_t - \mathbb{E}_{p_t(\boldsymbol{y}_t)}[\boldsymbol{y}_t])\right] + C_2, && (169)
\end{aligned}$$

where C_2 is a constant independent of $\boldsymbol{\theta}$ and the second term in (167) vanishes since:

$$\begin{aligned}
\mathbb{E}_{\boldsymbol{y}_t\sim p_t(\boldsymbol{y}_t)}\left[\nabla_{\boldsymbol{y}_t}\log p_t(\boldsymbol{y}_t)\right] &= \int_{\boldsymbol{y}_t}\nabla_{\boldsymbol{y}_t}p_t(\boldsymbol{y}_t)\mathrm{d}\boldsymbol{y}_t && (170)\\
&= 0, && (171)
\end{aligned}$$

where the last equality follows from integration by parts applied coordinatewise, under the standard assumption that p_t decays sufficiently fast at infinity. Substituting (163) and (169) into (160), CFM becomes equivalent to the following optimization problem:

$$\min_{\theta} \mathbb{E}_{c,t,x_t\sim p_t(x_t|c)}\Big[\|s_\theta(x_t,t,c)\|_2^2 - 2\big((1+\eta)\nabla_{x_t}\log p_t(x_t|c) + \eta\frac{b_t}{a_t}(x_t - \mathbb{E}_{p_t(y_t)}[y_t])\big)^T s_\theta(x_t,t,c)\Big], \tag{172}$$

hence the optimal solution is:

$$s_{\theta^*}(x_t,t,c) = (1+\eta)\nabla_{x_t}\log p_t(x_t|c) + \eta\frac{b_t}{a_t}(x_t - \mathbb{E}_{p_t(x_t)}[x_t]). \tag{173}$$

This completes the proof. □

Theorem 5 shows that the population-optimal CFM score consists of a rescaled conditional score and a global affine correction determined by the displacement from the unconditional noisy mean. Specifically, the correction has the form $\eta(b_t/a_t)(x_t - \mu_t)$, where $\mu_t := \mathbb{E}_{p_t(x_t)}[x_t]$.

This mechanism differs from both CFG and MCLR. The CFG correction is determined by the local conditional–unconditional score difference, whereas the MCLR correction additionally contains a sample-adaptive coefficient determined by the local likelihood ratio. In contrast, the CFM correction depends only on the displacement from the global unconditional mean and a time-dependent coefficient. Hence, CFM does not explicitly adapt its correction to the local conditional–unconditional density geometry.

To illustrate this distinction, consider a well-separated Gaussian mixture and a noise schedule satisfying $b_t/a_t > 0$. In this example, the CFM correction points away from the unconditional mean whenever $x_t \neq \mu_t$. Consequently, the correction may remain nonzero even when x_t lies in a high-density, class-unambiguous region. In contrast, the CFG correction can become small when the unconditional mixture score is dominated by the target component, while MCLR additionally attenuates its adaptive coefficient as the conditional–unconditional likelihood ratio increases. Consistent with these distinct correction mechanisms, our controlled experiments in Section G.2 show that MCLR yields larger and more sustained improvements than CFM across different regularization strengths.

C Theoretical Analysis of CC-DPO

C.1 Basics of DPO

Reward Modeling. For a given prompt c and two associated outputs x_w and x_l, where the subscripts 'w' stands for 'winning' while 'l' stands for 'losing', implying that x_w is preferred over x_l, DPO models the human preference distribution with the Bradley-Terry (BT) model [BT52]:

$$p(x_w \succ x_l|c) = \frac{\exp(r^*(x_w|c))}{\exp(r^*(x_w|c)) + \exp(r^*(x_l|c))} = \text{Sigmoid}(r^*(x_w|c) - r^*(x_l|c)) \tag{174}$$

where $r^*(x|c)$ is the underlying optimal reward function for prompt c. Intuitively, the preferred samples x_w should have higher reward values compared to the non-preferred samples x_l. Assuming access to a dataset of sampled comparisons $\mathcal{S} = \{(c^i, x_w^i, x_l^i)\}_{i=1}^N$, one can learn the optimal reward

function via maximum likelihood estimation:

$$r^* = \arg\max_r \mathbb{E}_{(c,x_w,x_l)\sim\mathcal{S}}\left[\log p(x_w \succ x_l|c)\right] \tag{175}$$

$$= \arg\min_r -\mathbb{E}_{(c,x_w,x_l)\sim S}\left[\log \text{Sigmoid}(r(x_w|c) - r(x_l|c))\right]. \tag{176}$$

RL fine-tuning Phase. Assuming access to the optimal preference reward function r^*, one can fine-tune a base model $p_{\text{ref}}(x|c)$ to align with the preference dataset by optimizing the following objective:

$$\max_\theta \mathbb{E}_c\left[\mathbb{E}_{x\sim p_\theta(x|c)}\left[r^*(x|c)\right] - \beta D_{KL}(p_\theta(x|c)\|p_{\text{ref}}(x|c))\right], \tag{177}$$

where β controls the KL-regularization strength. Intuitively, objective (177) encourages the fine-tuned model to achieve high reward value in expectation, and at the same time not deviate too much from the base model.

DPO Objective. Under certain regularity conditions, the optimal solution to the fine-tuning objective (177) admits the following closed-form:

$$p_{\theta^*}(x|c) = \frac{1}{Z(c)} p_{\text{ref}}(x|c)\exp\left(\frac{1}{\beta} r^*(x|c)\right), \tag{178}$$

where $Z(c)$ is a partition function for normalizing the density. To prove this, consider optimizing the fine-tuning objective (177) for a target condition c:

$$\arg\max_\theta \mathbb{E}_{x\sim p_\theta(x|c)}\left[r^*(x|c)\right] - \beta D_{KL}(p_\theta(x|c)\|p_{\text{ref}}(x|c)) \tag{179}$$

$$= \arg\max_\theta \mathbb{E}_{x\sim p_\theta(x|c)}\left[r^*(x|c) - \beta\log\frac{p_\theta(x|c)}{p_{\text{ref}}(x|c)}\right] \tag{180}$$

$$= \arg\min_\theta \mathbb{E}_{x\sim p_\theta(x|c)}\left[\log\frac{p_\theta(x|c)}{p_{\text{ref}}(x|c)} - \frac{1}{\beta} r^*(x|c)\right] \tag{181}$$

$$= \arg\min_\theta \mathbb{E}_{x\sim p_\theta(x|c)}\left[\log\frac{p_\theta(x|c)}{\frac{1}{Z(c)} p_{\text{ref}}(x|c)\exp(\frac{1}{\beta} r^*(x|c))} - \log Z(c)\right], \tag{182}$$

where $Z(c)$ is the partition function that normalizes $p_{\text{ref}}(x|c)\exp(\frac{1}{\beta} r^*(x|c))$:

$$Z(c) = \int_x p_{\text{ref}}(x|c)\exp(\frac{1}{\beta} r^*(x|c))\mathrm{d}x, \tag{183}$$

such that $p^*(x|c) := \frac{1}{Z(c)} p_{\text{ref}}(x|c)\exp(\frac{1}{\beta} r^*(x|c))$ is a valid probability distribution. Since $Z(c)$ is independent of θ, the fine-tuning objective (182) is equivalent to:

$$\min_\theta D_{\text{KL}}(p_\theta(x|c)\|p^*(x|c)), \tag{184}$$

which achieves its minimum value 0 if and only if:

$$p_\theta(x|c) = p^*(x|c) = \frac{1}{Z(c)} p_{\text{ref}}(x|c)\exp(\frac{1}{\beta} r^*(x|c)). \tag{185}$$

With some algebra, the optimal reward function can be expressed with $p_{\theta^*}(\boldsymbol{x}|\boldsymbol{c})$:

$$r^*(\boldsymbol{x}|\boldsymbol{c}) = \beta \log \frac{p_{\theta^*}(\boldsymbol{x}|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})} + \beta \log Z(\boldsymbol{c}). \tag{186}$$

The relationship (186) between the optimal reward and the optimal fine-tuned distribution suggests a convenient parameterization of the reward model:

$$r_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}) = \beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})}. \tag{187}$$

Substituting (187) into (176) results in the DPO objective:

$$\arg\min_{\boldsymbol{\theta}} \; -\mathbb{E}_{(\boldsymbol{c},\boldsymbol{x}_w,\boldsymbol{x}_l)\sim S}\left[\log \text{Sigmoid}(\beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}_w|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}_w|\boldsymbol{c})} - \beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}_l|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}_l|\boldsymbol{c})})\right]. \tag{188}$$

In this way, one can directly fine-tune the base model without explicitly modeling the reward.

C.2 Improving Conditional Modeling with CC-DPO

To adapt DPO for improving class specificity, we may treat samples from the target class $\boldsymbol{c}$ as preferred data ($\boldsymbol{x}_w$) and samples from other randomly selected classes as non-preferred data ($\boldsymbol{x}_l$). This leads to the following objective:

$$\min_{\boldsymbol{\theta}} \; -\mathbb{E}_{\boldsymbol{c},\tilde{\boldsymbol{c}},\boldsymbol{x}_w\sim p(\boldsymbol{x}|\boldsymbol{c}),\boldsymbol{x}_l\sim p(\boldsymbol{x}|\tilde{\boldsymbol{c}})}\left[\log \text{Sigmoid}(\beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}_w|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}_w|\boldsymbol{c})} - \beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}_l|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}_l|\boldsymbol{c})})\right] \tag{189}$$

$$= \min_{\boldsymbol{\theta}} \; -\mathbb{E}_{\boldsymbol{c},\boldsymbol{x}_w\sim p(\boldsymbol{x}|\boldsymbol{c}),\boldsymbol{x}_l\sim p(\boldsymbol{x})}\left[\log \text{Sigmoid}(\beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}_w|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}_w|\boldsymbol{c})} - \beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}_l|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}_l|\boldsymbol{c})})\right], \tag{190}$$

which admits a closed-form solution as stated in the following theorem.

Theorem 4. *Under certain regularity conditions, the optimal solution to (190) is:*

$$p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c}) = \frac{1}{\tilde{Z}(\boldsymbol{c})} p_{ref}(\boldsymbol{x}|\boldsymbol{c}) \left(\frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})}\right)^{\frac{1}{\beta}}, \tag{191}$$

where $\tilde{Z}(\boldsymbol{c}) = \int_{\boldsymbol{x}} p_{ref}(\boldsymbol{x}|\boldsymbol{c}) \left(\frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})}\right)^{\frac{1}{\beta}} d\boldsymbol{x}$ *is the normalizing constant.*

Proof. From (176) to (178), it is clear that the optimal solution to CC-DPO (190) is fully determined by the base model and the optimal solution to the following reward modeling objective:

$$r^* = \arg\min_{r} \; -\mathbb{E}_{\boldsymbol{c},\boldsymbol{x}_w\sim p(\boldsymbol{x}|\boldsymbol{c}),\boldsymbol{x}_l\sim p(\boldsymbol{x})}\left[\log \text{Sigmoid}(r(\boldsymbol{x}_w|\boldsymbol{c}) - r(\boldsymbol{x}_l|\boldsymbol{c}))\right]. \tag{192}$$

Note that r^* is the collection of optimal rewards $r^*(\cdot|\boldsymbol{c}_k)$ for each $\boldsymbol{c}_k \in \{\boldsymbol{c}_i\}_{i=1}^M$. Without loss of generality, we drop the subscript and solve for the optimal reward for a single target $\boldsymbol{c}$:

$$r^*(\cdot|\boldsymbol{c}) = \arg\min_{r(\cdot|\boldsymbol{c})} \; -\mathbb{E}_{\boldsymbol{x}_w\sim p(\boldsymbol{x}|\boldsymbol{c}),\boldsymbol{x}_l\sim p(\boldsymbol{x})}\left[\log \text{Sigmoid}(r(\boldsymbol{x}_w|\boldsymbol{c}) - r(\boldsymbol{x}_l|\boldsymbol{c}))\right] \tag{193}$$

$$= \arg\max_{r(\cdot|\boldsymbol{c})} \int_{\boldsymbol{x}_w}\int_{\boldsymbol{x}_l} \log \text{Sigmoid}(r(\boldsymbol{x}_w|\boldsymbol{c}) - r(\boldsymbol{x}_l|\boldsymbol{c}))p(\boldsymbol{x}_w|\boldsymbol{c})p(\boldsymbol{x}_l)\mathrm{d}\boldsymbol{x}_w\mathrm{d}\boldsymbol{x}_l. \tag{194}$$

Note that for any arbitrary pair of points $(\boldsymbol{x}_1, \boldsymbol{x}_2)$, they contribute to the integral (194) for the following amount:

$$\log \text{Sigmoid}\left(r(\boldsymbol{x}_1|\boldsymbol{c}) - r(\boldsymbol{x}_2|\boldsymbol{c})\right) p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2) + \log \text{Sigmoid}\left(r(\boldsymbol{x}_2|\boldsymbol{c}) - r(\boldsymbol{x}_1|\boldsymbol{c})\right) p(\boldsymbol{x}_2|\boldsymbol{c})p(\boldsymbol{x}_1), \tag{195}$$

hence $r^*(\cdot|\boldsymbol{c})$ is an optimal solution to (194) if it maximizes (195) for $\forall(\boldsymbol{x}_1, \boldsymbol{x}_2)$. To find such $r^*(\cdot|\boldsymbol{c})$, let's define $w = r(\boldsymbol{x}_1|\boldsymbol{c}) - r(\boldsymbol{x}_2|\boldsymbol{c})$ and solve the following optimization problem:

$$\max_w \mathcal{L}(w) := \log \text{Sigmoid}(w)p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2) + \log \text{Sigmoid}(-w)p(\boldsymbol{x}_2|\boldsymbol{c})p(\boldsymbol{x}_1). \tag{196}$$

Note that:

$$\nabla_w \mathcal{L}(w) = \text{Sigmoid}(-w)p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2) - \text{Sigmoid}(w)p(\boldsymbol{x}_2|\boldsymbol{c})p(\boldsymbol{x}_1), \tag{197}$$

which implies the stationary point w^* must satisfy:

$$\begin{aligned}
&\text{Sigmoid}(w^*)p(\boldsymbol{x}_2|\boldsymbol{c})p(\boldsymbol{x}_1) = \text{Sigmoid}(-w^*)p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2) && (198)\\
\Rightarrow\ &\text{Sigmoid}(w^*)\big(p(\boldsymbol{x}_2|\boldsymbol{c})p(\boldsymbol{x}_1) + p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2)\big) = p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2) && (199)\\
\Rightarrow\ &\text{Sigmoid}(w^*) = \frac{p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2)}{p(\boldsymbol{x}_2|\boldsymbol{c})p(\boldsymbol{x}_1) + p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2)} && (200)\\
\Rightarrow\ &\text{Sigmoid}(w^*) = \frac{1}{1 + \frac{p(\boldsymbol{x}_2|\boldsymbol{c})p(\boldsymbol{x}_1)}{p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2)}} && (201)\\
\Rightarrow\ &\text{Sigmoid}(w^*) = \frac{1}{1 + \exp\left(-\log \frac{p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2)}{p(\boldsymbol{x}_1)p(\boldsymbol{x}_2|\boldsymbol{c})}\right)} && (202)\\
\Rightarrow\ &w^* = \log \frac{p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2)}{p(\boldsymbol{x}_1)p(\boldsymbol{x}_2|\boldsymbol{c})}, && (203)
\end{aligned}$$

which implies:

$$r^*(\boldsymbol{x}|\boldsymbol{c}) = \log \frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})} + K, \tag{204}$$

where K is any finite constant. Moreover, since:

$$\begin{aligned}
\nabla_w^2 \mathcal{L}(w) &= -\text{Sigmoid}(w)\text{Sigmoid}(-w)p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2) - \text{Sigmoid}(-w)\text{Sigmoid}(w)p(\boldsymbol{x}_2|\boldsymbol{c})p(\boldsymbol{x}_1) && (205)\\
&< 0, && (206)
\end{aligned}$$

we know $\mathcal{L}(w)$ is concave and w^* is the global maximizer and consequently $r^*(\boldsymbol{x}|\boldsymbol{c})$ is the optimal reward function. Since this optimal reward function is consistent for any arbitrary pair of points $(\boldsymbol{x}_1, \boldsymbol{x}_2)$, it is the global reward function that maximizes the full objective (193).

Substituting (204) into (178), we obtain the optimal solution to CC-DPO:

$$p_{\theta^*}(\boldsymbol{x}|\boldsymbol{c}) = \frac{1}{\tilde{Z}(\boldsymbol{c})} p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c}) \left(\frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})}\right)^{\frac{1}{\beta}}, \tag{207}$$

where $\tilde{Z}(\boldsymbol{c}) = \int_{\boldsymbol{x}} p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c}) \left(\frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})}\right)^{\frac{1}{\beta}} d\boldsymbol{x}$ is the normalizing constant. This completes the proof. □

Remark. Note that from (199) to (200), we require $p(\boldsymbol{x}_2|\boldsymbol{c})p(\boldsymbol{x}_1) + p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2) > 0$ for $\forall \boldsymbol{x}_1, \boldsymbol{x}_2$, which holds true if both $p(\boldsymbol{x}|\boldsymbol{c})$ and $p(\boldsymbol{x})$ have full support on $\mathbb{R}^d$. Next, we discuss the corner cases where this assumption doesn't hold.

1. Suppose $p(\boldsymbol{x}_1|\boldsymbol{c}) = 0$ but $p(\boldsymbol{x}_1) > 0$, then we have:

$$\mathcal{L}(w) = \log \text{Sigmoid}\big(r(\boldsymbol{x}_2|\boldsymbol{c}) - r(\boldsymbol{x}_1|\boldsymbol{c})\big)p(\boldsymbol{x}_2|\boldsymbol{c})p(\boldsymbol{x}_1). \tag{208}$$

   In this case, the objective is maximized if $r^*(\boldsymbol{x}_1|\boldsymbol{c}) = -\infty$.

2. Suppose $p(\boldsymbol{x}_1|\boldsymbol{c}) > 0$ but $p(\boldsymbol{x}_1) = 0$, then we have:

$$\mathcal{L}(w) = \log \text{Sigmoid}\big(r(\boldsymbol{x}_1|\boldsymbol{c}) - r(\boldsymbol{x}_2|\boldsymbol{c})\big)p(\boldsymbol{x}_1|\boldsymbol{c})p(\boldsymbol{x}_2). \tag{209}$$

   In this case, the objective is maximized if $r^*(\boldsymbol{x}_1|\boldsymbol{c}) = +\infty$.

3. Suppose $p(\boldsymbol{x}_1|\boldsymbol{c}) = 0$ and $p(\boldsymbol{x}_1) = 0$, then $\boldsymbol{x}_1$ will never be sampled and it does not contribute to the overall objective. In such case, $r^*(\boldsymbol{x}_1|\boldsymbol{c})$ is undefined and will implicitly depend on the practical parameterization of the reward model.

The first two cases are already covered by (191). In particular, case 2 can lead to non-normalizable issue, as in this case $p_{\theta^*}(\cdot|\boldsymbol{c})$ will have infinite density at $\boldsymbol{x}_1$. Although since $p(\boldsymbol{x}_1) = \mathbb{E}_{\boldsymbol{c}}\left[p(\boldsymbol{x}_1|\boldsymbol{c})\right]$, it will not be zero given $p(\boldsymbol{x}_1|\boldsymbol{c}) > 0$. It is highly likely in practice, there exists $\boldsymbol{x}_1$ such that $p(\boldsymbol{x}_1|\boldsymbol{c}) > 0$ but $p(\boldsymbol{x}_1) \approx 0$. In this regime, the ratio $\left(\frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})}\right)^{\frac{1}{\beta}}$ blows up, forcing $p_{\theta^*}(\cdot|\boldsymbol{c})$ to essentially place all its mass on such $\boldsymbol{x}$, which again leads to non-normalizable issue.

Error Model for CC-DPO. As with the MCLR case, CC-DPO recovers the ground-truth conditional distribution under an appropriate error model, as stated in the following corollary.

Corollary 4. *If the base model satisfies:*

$$p_{ref}(\boldsymbol{x}|\boldsymbol{c}) \propto p(\boldsymbol{x}|\boldsymbol{c})^{1-\frac{1}{\beta}} p(\boldsymbol{x})^{\frac{1}{\beta}} \tag{210}$$

then fine-tuning $p_{ref}(\boldsymbol{x}|\boldsymbol{c})$ with the CC-DPO objective (24) recovers the ground-truth conditional distribution $p(\boldsymbol{x}|\boldsymbol{c})$.

Proof. Substituting (210) into (207), we get $p_{\theta^*}(\boldsymbol{x}|\boldsymbol{c}) = p(\boldsymbol{x}|\boldsymbol{c})$, which completes the proof. □

D Theoretical Analysis of CCA

In Section C.2, we've shown that the underlying optimal reward function induced by the CC-DPO objective (192) has the form of the log likelihood-ratio (204). Interestingly, the same reward function can be obtained by minimizing the following optimization problem:

$$r^*(\cdot|\boldsymbol{c}) = \arg\max_{r(\cdot|\boldsymbol{c})} \ \mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{c})} \log \text{Sigmoid}(r(\boldsymbol{x}|\boldsymbol{c})) + \mathbb{E}_{p(\boldsymbol{x})} \log \text{Sigmoid}(-r(\boldsymbol{x}|\boldsymbol{c})) \tag{211}$$

$$= \arg\max_{r(\cdot|\boldsymbol{c})} \ \int_{\boldsymbol{x}} \big(\log \text{Sigmoid}(r(\boldsymbol{x}|\boldsymbol{c}))p(\boldsymbol{x}|\boldsymbol{c}) + \log \text{Sigmoid}(-r(\boldsymbol{x}|\boldsymbol{c}))p(\boldsymbol{x})\big)\mathrm{d}\boldsymbol{x}. \tag{212}$$

Since the objective (212) decomposes pointwise over $\boldsymbol{x}$, the optimal reward function can be obtained by maximizing the integrand for each $\boldsymbol{x}$ independently:

$$\log \text{Sigmoid}(r(\boldsymbol{x}|\boldsymbol{c}))p(\boldsymbol{x}|\boldsymbol{c}) + \log \text{Sigmoid}(-r(\boldsymbol{x}|\boldsymbol{c}))p(\boldsymbol{x}). \tag{213}$$

To find $r^*(\cdot|\boldsymbol{c})$, let's define $w = r(\boldsymbol{x}|\boldsymbol{c})$ and solve the following optimization problem:

$$\max_w \ \mathcal{L}(w) := \log \text{Sigmoid}(w)p(\boldsymbol{x}|\boldsymbol{c}) + \log \text{Sigmoid}(-w)p(\boldsymbol{x}). \tag{214}$$

Note that:

$$\nabla_w \mathcal{L}(w) := \text{Sigmoid}(-w)p(\boldsymbol{x}|\boldsymbol{c}) - \text{Sigmoid}(w)p(\boldsymbol{x}), \tag{215}$$

which implies the stationary point w^* is:

$$w^* = r^*(\boldsymbol{x}|\boldsymbol{c}) = \log \frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})}. \tag{216}$$

Furthermore, since:

$$\nabla_w^2 \mathcal{L}(w) = -\text{Sigmoid}(w)\text{Sigmoid}(-w)p(\boldsymbol{x}|\boldsymbol{c}) - \text{Sigmoid}(w)\text{Sigmoid}(-w)p(\boldsymbol{x}) \tag{217}$$

$$< 0, \tag{218}$$

we know $\mathcal{L}(w)$ is concave and $r^*(\boldsymbol{x}|\boldsymbol{c}) = \log \frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})}$ is the unique maximizer. Hence, we've proved that the optimization problems (192) and (212) lead to the same optimal reward up to an additive constant.

Therefore, by parameterizing the reward function as:

$$r_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c}) = \beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})} \tag{219}$$

, and substitute this in (212), we get the following optimization problem:

$$p_{\boldsymbol{\theta}^*}(\cdot|\boldsymbol{c}) = \arg\max_{p_{\boldsymbol{\theta}}(\cdot|\boldsymbol{c})} \ \mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{c})} \log \text{Sigmoid}\big(\beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})}\big) + \mathbb{E}_{p(\boldsymbol{x})} \log \text{Sigmoid}\big(-\beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})}\big), \tag{220}$$

which has the global maximizer:

$$\beta \log \frac{p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})} = r^*(\boldsymbol{x}|\boldsymbol{c}) = \log \frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})} \tag{221}$$

$$\Rightarrow p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c}) = p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})\big(\frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})}\big)^{\frac{1}{\beta}}. \tag{222}$$

One issue with the optimization objective (220) is that its optimal solution (222) is not a valid probability distribution because it is not normalized. To alleviate this issue, we consider a slightly modified version of (212) by introducing an additional constant λ:

$$r^*(\cdot|\boldsymbol{c}) = \arg\max_{r(\cdot|\boldsymbol{c})} \ \mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{c})} \log \text{Sigmoid}(r(\boldsymbol{x}|\boldsymbol{c})) + \lambda \mathbb{E}_{p(\boldsymbol{x})} \log \text{Sigmoid}(-r(\boldsymbol{x}|\boldsymbol{c})). \tag{223}$$

With similar proof technique, it can be shown that:

$$r^*(\boldsymbol{x}|\boldsymbol{c}) = \log \frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})} + \log \frac{1}{\lambda}. \tag{224}$$

Substituting (219) in (223), we get the following optimization problem:

$$p_{\boldsymbol{\theta}^*}(\cdot|\boldsymbol{c}) = \arg\max_{p_{\boldsymbol{\theta}}(\cdot|\boldsymbol{c})} \mathbb{E}_{p(\boldsymbol{x}|\boldsymbol{c})} \log \text{Sigmoid}\big(\beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})}\big) + \lambda \mathbb{E}_{p(\boldsymbol{x})} \log \text{Sigmoid}\big(-\beta \log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})}\big), \tag{225}$$

which has global maximizer:

$$\beta \log \frac{p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})} = \log \frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})} + \log \frac{1}{\lambda} \tag{226}$$

$$\Rightarrow p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c}) = p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c}) \big(\frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})}\big)^{\frac{1}{\beta}} \big(\frac{1}{\lambda}\big)^{\frac{1}{\beta}}. \tag{227}$$

In this case, $p_{\boldsymbol{\theta}^*}(\boldsymbol{x}|\boldsymbol{c})$ is a valid probability distribution as long as:

$$\lambda^{\frac{1}{\beta}} = \int_{\boldsymbol{x}} p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c}) \big(\frac{p(\boldsymbol{x}|\boldsymbol{c})}{p(\boldsymbol{x})}\big)^{\frac{1}{\beta}} \mathrm{d}\boldsymbol{x}. \tag{228}$$

Note that optimal solution to (225) exactly coincides with that of CC-DPO (24). The formulation in (225), known as the Conditional Contrastive Alignment (CCA), which is essentially a combination of the Noise Contrastive Estimation (NCE) [GH10] and a special model parameterization (224). This objective is first proposed by [Che+25b] for improving the generation quality of visual autoregressive models without relying on CFG. It should be noted that CCA is theoretically correct only when λ is fixed to a particular constant, and it must be tuned as a hyperparameter in practice. In contrast, we demonstrate that CC-DPO fine-tuning recovers the same optimal solution without introducing this additional hyperparameter λ.

E Practical Implementation Details

E.1 Approximating Log-Likelihood with ELBO

Implementing MCLR requires access to the log-likelihood, which is not directly available for diffusion models. We therefore approximate the log-likelihood using the evidence lower bound (ELBO) in (6). In the following, we describe how this approximation is used in MCLR, CC-DPO, and CCA. Before doing so, we state the following fact.

Equivalence between Score Function and MMSE Denoisers. For practical diffusion models, the drift coefficient $f(\cdot)$ in (1) takes the form $f(\boldsymbol{x}, t) = f(t)\boldsymbol{x}$ where $f(\cdot) : \mathbb{R} \to \mathbb{R}$. As a result, the corresponding transition distribution is Gaussian and can be written as [Kar+22]:

$$p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) = \mathcal{N}(\boldsymbol{x}_t; s(t)\boldsymbol{x}, s^2(t)\sigma^2(t)\boldsymbol{I}) \tag{229}$$

, where $\mathcal{N}(\boldsymbol{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the Gaussian density with mean $\boldsymbol{\mu}$ and covariance $\boldsymbol{\Sigma}$ evaluated at $\boldsymbol{x}$, $s(t) = \exp(\int_0^t f(\xi)d\xi)$ and $\sigma(t) = \sqrt{\int_0^t \frac{g^2(\xi)}{s^2(\xi)} d\xi}$. The score of this transition distribution is therefore given by:

$$\nabla_{\boldsymbol{x}_t} \log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) = \frac{s(t)\boldsymbol{x} - \boldsymbol{x}_t}{s^2(t)\sigma^2(t)}. \tag{230}$$

Accordingly, the score network can be parameterized in terms of a denoiser $\mathcal{D}_{\boldsymbol{\theta}}(\cdot)$ as:

$$\boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}, t, \boldsymbol{c}) = \frac{\mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x}_t; s^2(t)\sigma(t), \boldsymbol{c}) - \boldsymbol{x}_t}{s^2(t)\sigma^2(t)}. \tag{231}$$

Without loss of generality, we set $s(t) = 1$, under which the DSM objective in (3) becomes:

$$\begin{aligned} &\frac{1}{2}\int_0^T \mathbb{E}_{\boldsymbol{c}, p(\boldsymbol{x}|\boldsymbol{c}), p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}[w(t)\|\frac{\mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x}_t; \sigma(t), \boldsymbol{c}) - \boldsymbol{x}}{\sigma^2(t)}\|_2^2]dt \\ &= \mathbb{E}_{\boldsymbol{c}, t\sim\mathcal{U}[0,T], p(\boldsymbol{x}|\boldsymbol{c}), p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\left[\tilde{w}(t)\|\mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x}_t; \sigma(t), \boldsymbol{c}) - \boldsymbol{x}\|_2^2\right], \end{aligned} \tag{232}$$

where $\tilde{w}(t) = \frac{Tw(t)}{2\sigma^4(t)}$ This shows that DSM is equivalent to training the MMSE denoiser $\mathcal{D}_{\boldsymbol{\theta}}(\cdot; \sigma(t), \boldsymbol{c})$ for data from class $\boldsymbol{c}$ corrupted by additive Gaussian noise with standard deviation $\sigma(t)$.

In this setting, the score function is related to the MMSE denoiser via Tweedie's formula [Miy+61]:

$$\nabla_{\boldsymbol{x}} \log p_t(\boldsymbol{x} \mid \boldsymbol{c}) = \frac{\mathcal{D}(\boldsymbol{x}; \sigma(t), \boldsymbol{c}) - \boldsymbol{x}}{\sigma^2(t)}, \tag{233}$$

where $\mathcal{D}(\boldsymbol{x}; \sigma(t), \boldsymbol{c})$ denotes the MMSE denoiser and $p_t(\boldsymbol{x}|\boldsymbol{c}) = \int p_{0t}(\boldsymbol{x}|\boldsymbol{x}_0)p_{\text{data}}(\boldsymbol{x}_0)d\boldsymbol{x}_0$ is the marginal distribution at time t. We are now ready to present the practical ELBO-approximated objectives of MCLR, CC-DPO, and CCA for diffusion models.

ELBO-Approximated MCLR for Diffusion Models. Let $\boldsymbol{v}_2 = \nabla_{\boldsymbol{x}_t} \log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c})$ and $\boldsymbol{v}_3 = \nabla_{\boldsymbol{x}_t} \log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \tilde{\boldsymbol{c}})$. Substituting ELBO (6) into (28), the MCLR regularized DSM becomes:

$$\min_{\boldsymbol{\theta}} \ \mathcal{J}_{\text{DSM}}(\boldsymbol{\theta}; g^2(\cdot)) + \eta\, \mathbb{E}_{\substack{\boldsymbol{c}, \tilde{\boldsymbol{c}}, t\sim\mathcal{U}[0,T] \\ p(\boldsymbol{x}|\boldsymbol{c}),\, p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}}\left[g^2(t)\Big(\|\boldsymbol{v}_2\|_2^2 - \|\boldsymbol{v}_3\|_2^2\Big)\right]. \tag{234}$$

Using the denoiser parameterization of the score network, together with a customized training-time noise sampling distribution $p(t)$ and adaptive weighting $w(t)$, the MCLR objective becomes:

$$\mathbb{E}_{\substack{\boldsymbol{c}, \tilde{\boldsymbol{c}}, t\sim p(t) \\ p(\boldsymbol{x}|\boldsymbol{c}),\, p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}}\left[w(t)\big(\|\boldsymbol{x} - \mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x}_t; \sigma(t), \boldsymbol{c})\|_2^2 \ - \ \|\boldsymbol{x} - \mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x}_t; \sigma(t), \tilde{\boldsymbol{c}})\|_2^2\big)\right], \tag{235}$$

which can be approximated with Monte Carlo sampling during training. Rather than parameterizing the noise level through the time variable $t \sim p(t)$, one may equivalently define a training noise distribution directly over σ, denoted by $p(\sigma)$ and a corresponding noise adaptive weighting $w(\sigma)$. Under this formulation, (235) can be rewritten as:

$$\mathbb{E}_{\substack{\boldsymbol{c}, \tilde{\boldsymbol{c}}, \boldsymbol{\epsilon}\sim\mathcal{N}(\boldsymbol{0}, \boldsymbol{I}) \\ \sigma\sim p(\sigma), p(\boldsymbol{x}|\boldsymbol{c})}}\left[w(\sigma)\big(\|\boldsymbol{x} - \mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x} + \sigma\boldsymbol{\epsilon}; \sigma, \boldsymbol{c})\|_2^2 \ - \ \|\boldsymbol{x} - \mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x} + \sigma\boldsymbol{\epsilon}; \sigma, \tilde{\boldsymbol{c}})\|_2^2\big)\right]. \tag{236}$$

Let $\boldsymbol{v}_4 = \nabla_{\boldsymbol{x}_t} \log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t, t, \boldsymbol{c})$ and $\boldsymbol{v}_5 = \nabla_{\boldsymbol{y}_t} \log p_{0t}(\boldsymbol{y}_t|\boldsymbol{y}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{y}_t, t, \boldsymbol{c})$. Similarly, substituting ELBO into (11), we get the following equivalent optimization problem:

$$\min_{\boldsymbol{\theta}} \ \mathcal{J}_{\text{DSM}}(\boldsymbol{\theta}; g^2(\cdot)) + \eta\, \mathbb{E}_{\substack{\boldsymbol{c}, \tilde{\boldsymbol{c}}, t\sim\mathcal{U}[0,T] \\ p(\boldsymbol{x}|\boldsymbol{c}),\, p(\boldsymbol{y}|\tilde{\boldsymbol{c}}) \\ p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}),\, p_{0t}(\boldsymbol{y}_t|\boldsymbol{y})}}\left[g^2(t)\Big(\|\boldsymbol{v}_4\|_2^2 - \|\boldsymbol{v}_5\|_2^2\Big)\right], \tag{237}$$

the denoiser form of which is presented in (17).

ELBO-Approximated CC-DPO for Diffusion Models. The CC-DPO algorithm requires access to the log-likelihood ratio for individual data point. Similar to MCLR, we approximate it with the ELBO:

$$\log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x})}{p_{\text{ref}}(\boldsymbol{x})} \approx \frac{T}{2}\mathbb{E}_{t\sim\mathcal{U}[0,T],p_{0t}(\boldsymbol{x}_t|\boldsymbol{x})}\Big[g^2(t)\big(-\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) - \boldsymbol{s}_{\boldsymbol{\theta}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2 + \tag{238}$$

$$\|\nabla_{\boldsymbol{x}_t}\log p_{0t}(\boldsymbol{x}_t|\boldsymbol{x}) - \boldsymbol{s}_{\text{ref}}(\boldsymbol{x}_t,t,\boldsymbol{c})\|_2^2\big)\Big]. \tag{239}$$

Using the denoiser parameterization of the score network, together with a customized training noise distribution and adaptive weighting, the log-likelihood ratio takes the following form:

$$\log \frac{p_{\boldsymbol{\theta}}(\boldsymbol{x})}{p_{\text{ref}}(\boldsymbol{x})} \approx \mathbb{E}_{\sigma\sim p_{\text{train}}(\sigma),\boldsymbol{\epsilon}\sim\mathcal{N}(\mathbf{0},\boldsymbol{I})}\big[w(\sigma)(-\|\boldsymbol{x}-\mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x}+\sigma\boldsymbol{\epsilon};\sigma,\boldsymbol{c})\|_2^2 + \|\boldsymbol{x}-\mathcal{D}_{\text{ref}}(\boldsymbol{x}+\sigma\boldsymbol{\epsilon};\sigma,\boldsymbol{c})\|_2^2)\big]. \tag{240}$$

Define:

$$\Delta(\boldsymbol{x},\sigma,\boldsymbol{\epsilon},\boldsymbol{c}) := \|\boldsymbol{x}-\mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x}+\sigma\boldsymbol{\epsilon};\sigma,\boldsymbol{c})\|_2^2 - \|\boldsymbol{x}-\mathcal{D}_{\text{ref}}(\boldsymbol{x}+\sigma\boldsymbol{\epsilon};\sigma,\boldsymbol{c})\|_2^2. \tag{241}$$

By substituting (240) into DPO objective (24), we get the following optimization problem:

$$\min_{\boldsymbol{\theta}} \; -\mathbb{E}_{\boldsymbol{c},\tilde{\boldsymbol{c}},\boldsymbol{x}_w\sim p(\boldsymbol{x}|\boldsymbol{c}),\boldsymbol{x}_l\sim p(\boldsymbol{x}|\tilde{\boldsymbol{c}})}\big[\log\text{Sigmoid}(\beta\mathbb{E}_{\sigma,\boldsymbol{\epsilon}}[w(\sigma)(-\Delta(\boldsymbol{x}_w,\sigma,\boldsymbol{\epsilon},\boldsymbol{c})+\Delta(\boldsymbol{x}_l,\sigma,\boldsymbol{\epsilon},\boldsymbol{c}))])\big]. \tag{242}$$

Since the log-Sigmoid function is concave, applying Jensen's inequality by moving the expectation outside yields the following upper bound of objective (242):

$$\min_{\boldsymbol{\theta}} \; -\mathbb{E}_{\boldsymbol{c},\tilde{\boldsymbol{c}},\boldsymbol{x}_w,\boldsymbol{x}_l,\sigma,\boldsymbol{\epsilon}}\big[\log\text{Sigmoid}(\beta w(\sigma)(-\Delta(\boldsymbol{x}_w,\sigma,\boldsymbol{\epsilon},\boldsymbol{c})+\Delta(\boldsymbol{x}_l,\sigma,\boldsymbol{\epsilon},\boldsymbol{c})))\big], \tag{243}$$

which serves as our final training objective.

CCA for Diffusion Models. The Conditional Contrastive Alignment (CCA) [Che+25b] objective takes the following form:

$$\max_{\boldsymbol{\theta}} \; \mathbb{E}_{\boldsymbol{c},p(\boldsymbol{x}|\boldsymbol{c})}\log\text{Sigmoid}\big(\beta\log\frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})}\big) + \lambda\mathbb{E}_{\boldsymbol{c},\tilde{\boldsymbol{c}},p(\boldsymbol{x}|\tilde{\boldsymbol{c}})}\log\text{Sigmoid}\big(-\beta\log\frac{p_{\boldsymbol{\theta}}(\boldsymbol{x}|\boldsymbol{c})}{p_{\text{ref}}(\boldsymbol{x}|\boldsymbol{c})}\big). \tag{244}$$

To adapt this optimization problem to diffusion models, we again approximate the log-likelihood with ELBO, resulting in the following optimization problem:

$$\max_{\boldsymbol{\theta}} \; \mathbb{E}_{\boldsymbol{c},\tilde{\boldsymbol{c}},\boldsymbol{x}_w\sim p(\boldsymbol{x}|\boldsymbol{c}),\boldsymbol{x}_l\sim p(\boldsymbol{x}|\tilde{\boldsymbol{c}})}\Big[\log\text{Sigmoid}\big(\beta\mathbb{E}_{\sigma,\boldsymbol{\epsilon}}\big[-w(\sigma)\Delta(\boldsymbol{x}_w,\sigma,\boldsymbol{\epsilon},\boldsymbol{c})\big]\big) + \lambda\log\text{Sigmoid}\big(\beta\mathbb{E}_{\sigma,\boldsymbol{\epsilon}}\big[w(\sigma)\Delta(\boldsymbol{x}_l,\sigma,\boldsymbol{\epsilon},\boldsymbol{c})\big]\big)\Big]. \tag{245}$$

Since the log-sigmoid function is concave, applying Jensen's inequality by moving the expectation outside yields the following lower bound of objective (245):

$$\max_{\boldsymbol{\theta}} \; \mathbb{E}_{\boldsymbol{c},\boldsymbol{x}_w,\boldsymbol{x}_l,\sigma,\boldsymbol{\epsilon}}\big[\log\text{Sigmoid}\big(-\beta w(\sigma)\Delta(\boldsymbol{x}_w,\sigma,\boldsymbol{\epsilon},\boldsymbol{c})\big) + \lambda\log\text{Sigmoid}\big(\beta w(\sigma)\Delta(\boldsymbol{x}_l,\sigma,\boldsymbol{\epsilon},\boldsymbol{c})\big)\big], \tag{246}$$

which serves as our final training objective.

Remarks on the ELBO Approximation. As discussed in Section B.4, despite standard practice in the literature, when replacing log-likelihood with ELBO, the resulting objectives do not necessarily correspond to the original likelihood formulations, since the ELBO does not enforce the regularity conditions required for the parameterized score function to define a valid score field. Consequently, ELBO-approximated MCLR, DPO and CCA should be viewed as **approximate** likelihood-ratio maximization procedures for diffusion models. This issue does not arise for autoregressive models, where likelihoods are available exactly.

E.2 Building Training Data from a Minibatch

The equivalence of (27) and (28) offer us two choices for constructing the contrastive pairs. The first is constructing tuples of $(\boldsymbol{x}, \boldsymbol{y}, \boldsymbol{c})$, where $\boldsymbol{x} \sim p(\boldsymbol{x}|\boldsymbol{c})$ and $\boldsymbol{y}$ sampled from other random classes. The other is constructing contrastive tuples $(\boldsymbol{x}, \boldsymbol{c}, \tilde{\boldsymbol{c}})$, where $\boldsymbol{x} \sim p(\boldsymbol{x}|\boldsymbol{c})$ and $\tilde{\boldsymbol{c}}$ denotes a randomly chosen alternative class. In our implementation we choose the second method, but our preliminary results demonstrate both achieve similar performance.

Similarly, CC-DPO and CCA require constructing preference-style tuples $(\boldsymbol{x}_w, \boldsymbol{x}_l, \boldsymbol{c})$. In practice, we build these tuples directly from each training minibatch $\{(\boldsymbol{x}_i, \boldsymbol{c}_i)\}_{i=1}^{N}$. We consider two strategies, described below.

Approach 1: Building N pairs. The simplest strategy constructs one contrastive (or preference) tuple per sample in the minibatch. For MCLR, given a sample $(\boldsymbol{x}, \boldsymbol{c})$, we randomly select another label $\tilde{\boldsymbol{c}}$ from the same minibatch, forming a tuple $(\boldsymbol{x}, \boldsymbol{c}, \tilde{\boldsymbol{c}})$. Repeating this process independently for each $\boldsymbol{x}$ (with replacement) yields N tuples in total. The same strategy applies to CC-DPO and CCA, where we construct N tuples $(\boldsymbol{x}_w, \boldsymbol{x}_l, \boldsymbol{c})$ from the minibatch.

Approach 2: Building NK pairs. Alternatively, we can construct multiple contrastive tuples per sample. For MCLR, given each $(\boldsymbol{x}, \boldsymbol{c})$, we randomly select K alternative labels $\{\tilde{\boldsymbol{c}}_k\}_{k=1}^{K}$ from the same minibatch, yielding K tuples $(\boldsymbol{x}, \boldsymbol{c}, \tilde{\boldsymbol{c}}_k)$. Repeating this procedure for all N samples results in NK tuples in total. An analogous strategy is applied to CC-DPO and CCA, producing NK tuples $(\boldsymbol{x}_w, \boldsymbol{x}_l, \boldsymbol{c})$ per minibatch.

Compared to Approach 1, Approach 2 increases the number of training tuples constructed from each minibatch, allowing better exploitation of inter-class contrastive information. In our experiments, we adopt Approach 2 for EDM-based diffusion models, where it consistently yields improved quantitative performance. For VAR and SiT models, we use Approach 1 due to its lower computational overhead and comparable empirical performance.

E.3 Overall Algorithm

We are now ready to present the overall algorithm. The algorithm for MCLR is given in Algorithm 1, while the algorithms for CC-DPO and CCA are presented in Algorithm 2. In Algorithm 1, we compute the MCLR regularization loss without including the standard DSM term. Empirically, we find that optimizing the MCLR term alone yields a better best-case $\text{FD}_{\text{DINOv2}}$ score. One possible explanation is that removing the DSM term makes the contrastive objective the dominant training signal, allowing the model to explore a larger solution space beyond the neighborhood constrained by DSM regularization. Accordingly, the main results report the best $\text{FD}_{\text{DINOv2}}$ scores obtained without the DSM term. Nevertheless, including the DSM term can improve training stability and lead to a more favorable Precision–Recall trade-off as shown in Section G.1.

Algorithm 1 MCLR

Require: Pre-trained base model $\mathcal{D}_{\theta_0}$; noise-adaptive weight function $w(\cdot)$; training noise schedule $p(\sigma)$; learning rate α; training dataset representing $p(\boldsymbol{x}, \boldsymbol{c})$; $K \geq 1$ (number of mismatched class labels per sample in Approach 2)

Initialize $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}_0$
for each training iteration **do**
  Sample a minibatch $\{(\mathbf{x}_i, \mathbf{c}_i)\}_{i=1}^{N}$ with $(\boldsymbol{x}_i, \boldsymbol{c}_i) \sim p(\boldsymbol{x}, \boldsymbol{c})$.
  Sample $\{\sigma_i\}_{i=1}^{N}$ with $\sigma_i \sim p(\sigma)$
  Sample $\{\boldsymbol{\epsilon}_i\}_{i=1}^{N}$ with $\boldsymbol{\epsilon}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$
  Construct contrastive tuples from the minibatch:
  if Approach 1 **then**
    Set $M \leftarrow N$
    For each i, sample $\tilde{\boldsymbol{c}}_i$ from minibatch labels with $\tilde{\boldsymbol{c}}_i \neq \boldsymbol{c}_i$
    Define tuples $\{(\boldsymbol{x}_i, \boldsymbol{c}_i, \tilde{\boldsymbol{c}}_i, \sigma_i, \boldsymbol{\epsilon}_i)\}_{i=1}^{M}$
  end if
  if Approach 2 **then**
    Set $M \leftarrow NK$
    For each i and each $k = 1, \ldots, K$, sample $\tilde{\boldsymbol{c}}_{i,k}$ from minibatch with $\tilde{\boldsymbol{c}}_{i,k} \neq \boldsymbol{c}_i$
    Define tuples $\{(\boldsymbol{x}_{i,k}, \boldsymbol{c}_{i,k}, \tilde{\boldsymbol{c}}_{i,k}, \sigma_{i,k}, \boldsymbol{\epsilon}_{i,k})\}$ where $(\boldsymbol{x}_{i,k}, \boldsymbol{c}_{i,k}) = (\boldsymbol{x}_i, \boldsymbol{c}_i)$ and $(\sigma_{i,k}, \boldsymbol{\epsilon}_{i,k}) = (\sigma_i, \boldsymbol{\epsilon}_i)$
  end if
  $\mathcal{L} \leftarrow 0$
  for each tuple $(\boldsymbol{x}, \boldsymbol{c}, \tilde{\boldsymbol{c}}, \sigma, \boldsymbol{\epsilon})$ **do**
    $\mathcal{L} \mathrel{+}= w(\sigma)\Big(\|\mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x} + \sigma\boldsymbol{\epsilon}; \sigma, \boldsymbol{c}) - \boldsymbol{x}\|_2^2 - \|\mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x} + \sigma\boldsymbol{\epsilon}; \sigma, \tilde{\boldsymbol{c}}) - \boldsymbol{x}\|_2^2\Big)$
  end for
  $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \alpha\nabla_{\boldsymbol{\theta}}\mathcal{L}$
end for
return $\mathcal{D}_{\boldsymbol{\theta}}$

E.4 Hyperparameters

For each method, we carefully tune hyperparameters using grid search to ensure the best performance. The full set of hyperparameter configurations will be made publicly available upon publication.

E.5 Computational Efficiency

MCLR fine-tuning incurs additional cost, but it is modest relative to pretraining. In our runs on 4×A40 40GB GPUs, for EDM2-S (64×64), MCLR fine-tuning took 5h23m to reach best $\text{FD}_{\text{DINOv2}}$ score 56 with Approach 1 (construct N contrastive pairs per batch) and 18h40m to reach best $\text{FD}_{\text{DINOv2}}$ 52 with Approach 2 (construct NK contrastive pairs per batch), starting from a base model with $\text{FD}_{\text{DINOv2}}$ score 95. In contrast, the pretraining stage requires training with 32 A100 GPUs for weeks [Kar+24b]. A comprehensive summary of computational cost can be found within our codebase upon publication.

F Additional Experimental Results

In this section, we present additional experimental results that complement and extend those in the main text. Specifically, we focus on validating the following three claims:

1. MCLR induces progressive class separation, leading to a fidelity–diversity trade-off analogous to that produced by increasing the guidance strength in classifier-free guidance (CFG).
2. MCLR matches or outperforms training-time baselines, including CC-DPO and CCA.
3. MCLR achieves performance comparable to CFG.

F.1 Progressive Class Separation and the Fidelity-Diversity Trade-off.

Training with MCLR induces notable *progressive class separation*, which gives rise to a fidelity–diversity trade-off similar that observed when increasing the guidance scale in CFG. This behavior is quantitatively supported by the following observations.
(i) Increased Precision and Inception Score. As shown in Figure 7(b,d,g) for ImageNet-64, Figure 9(c,f,i) for ImageNet-256 and Figure 8(b,d,g) for ImageNet-512, continued training with MCLR leads to a progressive increase in Inception Score, indicating increasingly class-discriminative generations. This is accompanied by a corresponding increase in Precision, reflecting improved image fidelity.
(ii) Decreased Recall. Conversely, as shown in Figure 7(e,h), Figure 8(e,h), and Figure 9(g,j), excessive training reduces within-class diversity, which manifests as a decrease in Recall (see also Figure 4(e)).

Taken together, these effects result in FD–IS and Precision–Recall trade-offs that closely mirror those induced by CFG, as illustrated in Figure 7(c,f,i), Figure 8(c,f,i), and Figure 9(d,e,h,k). Qualitatively, as shown in Figures 2 and 10 to 21, continued MCLR training leads to gradually enhanced distinct, class-specific visual characteristics in generated samples.

F.2 MCLR Outperforms Training-time Baselines.

Diffusion Models. As shown in Figures 7 and 8, for EDM models trained on ImageNet, MCLR achieves substantially better best-case $\text{FD}_{\text{DINOv2}}$ scores than CCA, CC-DPO. Moreover, MCLR traverses significantly wider FD–IS and Precision–Inception trade-offs region with a faster training speed. In contrast, CCA and CC-DPO often converge early, exhibit zigzag training trajectories, and become trapped in suboptimal local minima, as evidenced by slowly improving or stalled learning curves. **Autoregressive Models.** As shown in Figure 9, MCLR, CC-DPO, and CCA yield comparable performance improvement on VAR-d24. Specifically, we see MCLR consistently achieves higher precision (see Figure 9(f,h,i,k)) and inception score (see Figure 9(c)) in later stages of training.

F.3 MCLR Achieves Comparable Performance as CFG.

Diffusion Models. For EDM models trained on ImageNet, CFG generally exhibits a better FD–IS trade-off and achieves a lower best-case $\text{FD}_{\text{DINOv2}}$ than MCLR. Nevertheless, this gap is moderate for EDM2-L model, where CFG attains a best-case $\text{FD}_{\text{DINOv2}}$ of 39.86 compared to 42.50 for MCLR as shown in Figure 8(c). Moreover, for EDM2-L model, when evaluated using Precision–Recall, MCLR demonstrates competitive and in some regimes superior behavior. As shown in Figure 8(f,g), MCLR matches CFG in the high-recall regime (early training stages) and achieves a substantially

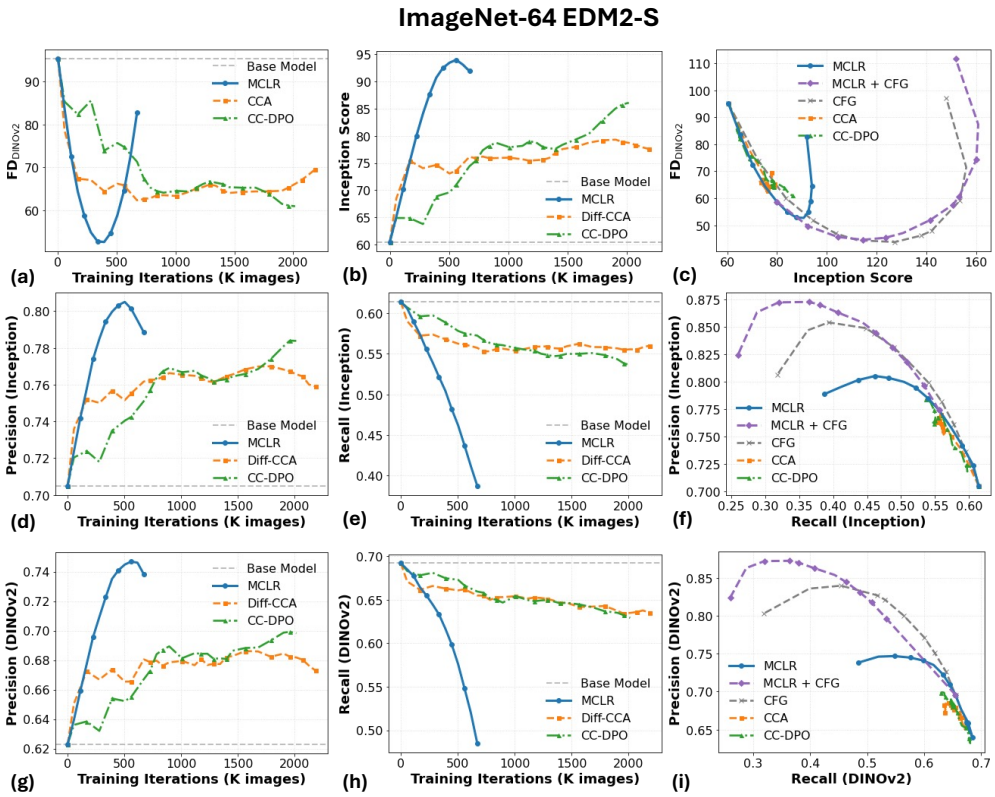


Figure 7: **Quantitative Results for EDM2-S trained on ImageNet-64×64.** (a), (b), (d), (e), (g), (h) show the evolution of FD, Inception Score, Precision (calculated with Inception features), and Recall (calculated with Inception features), Precision (calculated with DINOv2 features), and Recall (calculated with DINOv2 features), respectively, as functions of training iterations. (c) shows the FD–IS trade-offs, while (f), (i) depict the Precision–Recall trade-offs calculated with Inception and DINOv2 features respectively. We evaluate classifier-free guidance (CFG) scales $\gamma \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.7, 0.9, 1, 1.5, 2.0, 3.0\}$.

higher best-case precision in the high-precision regime (later training stages), where CFG begins to generate images with oversaturated colors.

Applying CFG on top of an MCLR-fine-tuned model further alleviates the performance gap between MCLR and CFG, yielding higher best-case Inception Score and Precision on both EDM2-S and EDM2-L models, as shown in Figure 7(c,f,i) and Figure 8(c,f,i).

Qualitatively, MCLR and CFG often produce similar effects, both significantly enhancing class-specific structures in the generated images, as shown in Figures 2 and 10 to 21. This similarity is expected, as both methods improve conditional modeling through inter-class contrastive signals. The key distinction is that MCLR internalizes this mechanism during training, whereas CFG applies it at inference time.

Autoregressive Models. For VAR-d24 model, MCLR achieves a similar FD–IS trade-off to CFG in terms of both $\text{FD}_{\text{DINOv2}}$ and FID, with CFG exhibiting a slightly better best-case FID. Consistent with diffusion models, MCLR outperforms CFG in the Precision–Recall trade-off and achieves a higher best-case precision.

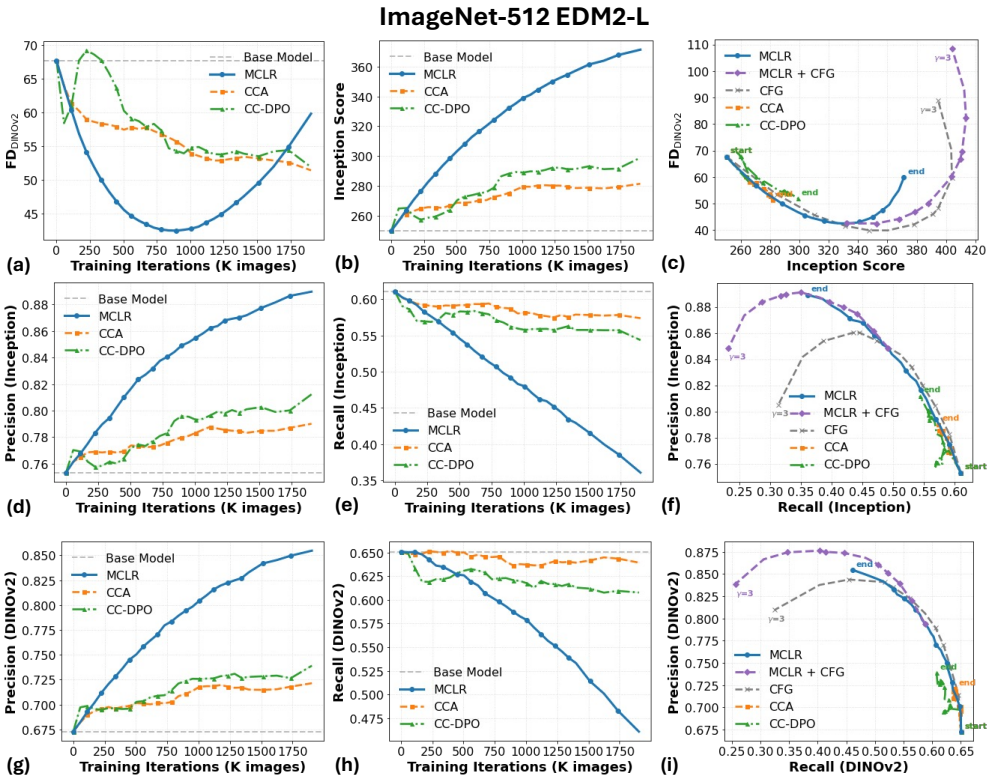


Figure 8: **Quantitative Results for EDM2-L trained on ImageNet-512×512.** (a), (b), (d), (e), (g), (h) show the evolution of FD, Inception Score, Precision (calculated with Inception features), and Recall (calculated with Inception features), Precision (calculated with DINOv2 features), and Recall (calculated with DINOv2 features), respectively, as functions of training iterations. (c) shows the FD–IS trade-offs, while (f), (i) depict the Precision–Recall trade-offs calculated with Inception and DINOv2 features respectively. We evaluate classifier-free guidance (CFG) scales $\gamma \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.7, 0.9, 1, 1.5, 2.0, 3.0\}$.

G Ablation Study

G.1 Effects of DSM and KL Regularization

In this section we study the effects of combining MCLR with either (i) the standard DSM objective or (ii) a KL-style regularizer. Specifically, we fine-tune a REPA-enhanced SiT-XL/2 base model $\mathcal{D}_{\text{ref}}(\cdot)$ pretrained on ImageNet-256 × 256 using the following two objectives:

$$\beta \underbrace{\mathbb{E}_{\substack{\boldsymbol{c},\boldsymbol{\epsilon}\sim\mathcal{N}(\mathbf{0},\boldsymbol{I}) \\ \sigma\sim p(\sigma),p(\boldsymbol{x}|\boldsymbol{c})}} \left[w(\sigma) \left(\|\boldsymbol{x} - \mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x} + \sigma\boldsymbol{\epsilon}; \sigma, \boldsymbol{c})\|_2^2 \right) \right]}_{\text{DSM Objective}}$$
$$+ \underbrace{\mathbb{E}_{\substack{\boldsymbol{c},\tilde{\boldsymbol{c}},\boldsymbol{\epsilon}\sim\mathcal{N}(\mathbf{0},\boldsymbol{I}) \\ \sigma\sim p(\sigma),p(\boldsymbol{x}|\boldsymbol{c}),p(\boldsymbol{y}|\tilde{\boldsymbol{c}})}} \left[w(\sigma) \left(\|\boldsymbol{x} - \mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x} + \sigma\boldsymbol{\epsilon}; \sigma, \boldsymbol{c})\|_2^2 \; - \; \|\boldsymbol{y} - \mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{y} + \sigma\boldsymbol{\epsilon}; \sigma, \boldsymbol{c})\|_2^2 \right) \right]}_{\text{MCLR Regularization}}, \tag{247}$$

$$\underbrace{\beta \mathop{\mathbb{E}}_{\substack{\boldsymbol{c},\boldsymbol{\epsilon}\sim\mathcal{N}(\mathbf{0},\boldsymbol{I})\\ \sigma\sim p(\sigma),p(\boldsymbol{x}|\boldsymbol{c})}}\left[w(\sigma)\left(\|\mathcal{D}_{\text{ref}}(\boldsymbol{x}+\sigma\boldsymbol{\epsilon};\sigma,\boldsymbol{c})-\mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x}+\sigma\boldsymbol{\epsilon};\sigma,\boldsymbol{c})\|_2^2\right)\right]}_{\text{KL Loss}}$$
$$+\underbrace{\mathop{\mathbb{E}}_{\substack{\boldsymbol{c},\tilde{\boldsymbol{c}},\boldsymbol{\epsilon}\sim\mathcal{N}(\mathbf{0},\boldsymbol{I})\\ \sigma\sim p(\sigma),p(\boldsymbol{x}|\boldsymbol{c}),p(\boldsymbol{y}|\tilde{\boldsymbol{c}})}}\left[w(\sigma)\left(\|\boldsymbol{x}-\mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x}+\sigma\boldsymbol{\epsilon};\sigma,\boldsymbol{c})\|_2^2\ -\ \|\boldsymbol{y}-\mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{y}+\sigma\boldsymbol{\epsilon};\sigma,\boldsymbol{c})\|_2^2\right)\right]}_{\text{MCLR Regularization}}, \tag{248}$$

where (248) is an approximation of the KL-regularized MCLR in (12). Although the exact KL objective in (12) should be defined using samples generated by the reference model, we evaluate this term on real data for simplicity.

Figure 22 compares MCLR alone with its DSM- and KL-regularized variants, where $\beta = 1$ is used for both regularized objectives. MCLR alone achieves the best-case $\text{FD}_{\text{DINOv2}}$ score, whereas adding either regularizer improves the FD–IS and Precision–Recall trade-offs and yields a better best-case FID. These results suggest that the additional regularization mainly improves training stability and the overall quality–diversity trade-off, rather than the best achievable $\text{FD}_{\text{DINOv2}}$.

Effect of Regularization Strength. We further study the regularization strength β for the DSM-regularized objective (247), considering $\beta \in \{0, 1, 5\}$. As shown in Figure 23(a,b), MCLR without additional regularization ($\beta = 0$) already substantially improves the reference model on both $\text{FD}_{\text{DINOv2}}$ and FID score, achieving the best $\text{FD}_{\text{DINOv2}}$ score. Adding moderate DSM regularization ($\beta = 1$) further improves the best-case FID score and the Precision–Recall trade-off as shown in Figure 23(b,c). However, an excessively large weight ($\beta = 5$) slows convergence and limits the gains from MCLR.

Smaller values of β generally lead to faster convergence. The best performance is reached within fewer than 1,000 iterations for $\beta = 0$, compared with more than 2,000 iterations for $\beta = 1$ and 19,000 iterations for $\beta = 5$. Without DSM regularization, however, prolonged fine-tuning can lead to performance degradation; in our experiments, this behavior can be mitigated through early stopping.

G.2 Empirical Comparison with Contrastive Flow Matching

We provide an empirical comparison between MCLR and Contrastive Flow Matching (CFM), discussed in Section B.5. Both methods augment the same standard training objective with a contrastive term, while differing in the construction of that term. This structural similarity enables a controlled comparison in which we keep the base model, learning rate, and all other training settings fixed and change only the contrastive objective. We set the contrastive weight to 1 for both methods and vary the weight of the standard DSM/flow matching term over $\beta \in \{0, 1, 5\}$. As shown in Figures 24 to 26, both methods initially improve FID and $\text{FD}_{\text{DINOv2}}$ across all three values of β. However, MCLR consistently yields substantially larger and more sustained improvements, whereas CFM improves only briefly before its performance rapidly deteriorates.

Importantly, although CFM is presented as a method for accelerating diffusion training [Sto+25], this interpretation should be treated with caution. Its early improvements in fidelity metrics are accompanied by a substantial decrease in recall, and therefore do not necessarily indicate faster learning of the target distribution. Since inference-time CFG can similarly improve fidelity at the cost of diversity, establishing genuine training acceleration would require comparisons at matched

precision and recall. Therefore, improvements in FID alone are insufficient to fully substantiate the training-acceleration claim. We therefore advocate evaluating claims of training acceleration at matched fidelity–diversity operating points, rather than relying on FID alone.

G.3 Hyperparameter Sensitivity of CC-DPO

The hyperparameter β in the CC-DPO objective (24) controls the strength of the KL regularization between the base model and the fine-tuned model. We study its effect on EDM2-S trained on ImageNet-64 $\times$ 64. As shown in Figure 27, CC-DPO exhibits relatively stable performance across a wide range of β values. In particular, smaller β values achieve marginally better best-case $\text{FD}_{\text{DINOv2}}$ and Inception Score.

G.4 Hyperparameter Sensitivity of CCA

The hyperparameter λ in the CCA objective (246) controls the strength of the contrastive term, with larger λ placing greater emphasis on penalizing the non-preferred sample $\boldsymbol{x}_l$. We study its effect on EDM2-S trained on ImageNet-64 $\times$ 64, fixing $\beta = 0.001$. As shown in Figure 28, λ has little impact on the best achievable $\text{FD}_{\text{DINOv2}}$; however, smaller λ values lead to more stable and smoother convergence. In contrast, larger λ values can yield slightly higher best-case Inception Scores, but often cause training instability and eventual collapse when training is prolonged.

In practice, jointly tuning β and λ is challenging. Although CCA and CC-DPO share the same theoretical optimal solution, CC-DPO requires tuning only a single hyperparameter β. Empirically, CCA does not outperform CC-DPO under careful tuning as shown in Figure 28. Therefore, we recommend CC-DPO over CCA in practice.

H Discussion on Related Works

H.1 Visual Generation without Guidance

Our work contributes to a growing line of research [Che+25b; Che+25a; Tan+25] that seeks to induce CFG-like effect by modifying the standard training objective, rather than applying CFG at inference time. The most closely related approach is CCA [Che+25b], which aims to learn the gamma-powered distribution (25) in autoregressive models using **Noise Contrastive Estimation (NCE)** [GH10]. While CCA also considers DPO as a baseline and reports inferior performance, our theoretical analysis demonstrates that CC-DPO and CCA admit the same optimal solution, and are therefore fundamentally equivalent at the population level. With a correct implementation, we find that CC-DPO consistently matches or exceeds the empirical performance of CCA while requiring fewer hyperparameters and simpler optimization. We further extend CCA to diffusion models by approximating log-likelihoods via the ELBO and include this variant as a baseline in our experiments. A detailed theoretical and empirical analysis of CCA is provided in Section D. Another relevant direction is **Guidance-Free Training (GFT)** [Che+25a; Tan+25], which aims to directly train diffusion models to reproduce CFG-induced score functions. GFT relies on a reparameterized score function in the same functional form as CFG, hence inherits the same theoretical ambiguities associated with CFG itself. In contrast, the proposed MCLR objective provides a clearer mechanistic interpretation of CFG.

Besides these approaches, several works sought to enhance conditional generative modeling through class-wise contrastive mechanisms. [Yan+24] introduce a contrastive objective to improve the modeling of tail classes; [Lee+25] employ an InfoNCE-style loss [OLV18] to strengthen

text-image alignment in diffusion models; [KMS24] and [YAA25] regularize diffusion models by encouraging class separation in feature space. In contrast to MCLR, CC-DPO and CCA, these approaches are primarily heuristic and do not provide a theoretical characterization of the conditional distribution induced by their objectives.

Lastly, beyond class-wise contrast, another complementary strategy for improving generative modeling is to contrast high-quality ("real") data against low-quality ("synthetic") samples. Zheng et al. [Zhe+25] propose Direct Discriminative Optimization (DDO), which employs the same contrastive objective as CCA, but replaces class-conditioned contrast with a real–synthetic discrimination signal. We include DDO as one of our baselines.

H.2 Inference-Time Alignment via Guidance

Recent works [Fra+25; Jin+25; Che+25c; JWL26] observe that CFG-style inference-time guidance can produce effects resembling those of training-time alignment methods, including reinforcement learning-based approaches. These studies provide empirical evidence that guidance may implicitly induce alignment-like behavior.

However, existing analyses are largely heuristic and often rely on unrealistic assumptions such as the guided score corresponds to that of the gamma-powered distribution (25), an assumption that has been proven incorrect [Kar+24a; BN24]. As a result, a rigorous theoretical connection between CFG and alignment objectives remains incomplete.

By establishing the equivalence between CFG and a weighted MCLR objective, our work provides a formal mechanistic interpretation of CFG as an inference-time alignment algorithm. To the best of our knowledge, this is the first result that rigorously connects CFG-style guidance with likelihood-ratio-based alignment objectives.

I Broader Impact

This work aims to improve the training of conditional diffusion models by reducing reliance on inference-time guidance mechanisms such as classifier-free guidance (CFG). By incorporating contrastive signals directly into the training objective, the proposed method enables standard sampling procedures to achieve CFG-like effects, potentially reducing computational overhead and simplifying deployment in downstream applications. These improvements may benefit a wide range of generative modeling tasks, including image synthesis, data augmentation, and representation learning. At the same time, the proposed approach does not introduce new application domains or capabilities beyond those of existing diffusion models, and thus inherits similar ethical considerations and risks, such as the potential misuse of generated content. As with prior generative models, responsible use, appropriate dataset curation, and adherence to existing guidelines remain essential. Overall, this work contributes to a better theoretical and practical understanding of conditional diffusion models, with the goal of improving their efficiency and interpretability

Algorithm 2 CC-DPO and CCA

Require: Pre-trained base model $\mathcal{D}_{\boldsymbol{\theta}_0}$; noise-adaptive weight function $w(\cdot)$; training noise schedule $p(\sigma)$; learning rate α; KL regularization strength β; CCA hyperparameter λ; training dataset representing $p(\boldsymbol{x}, \boldsymbol{c})$; $K \geq 1$ (number of mismatched class labels per sample in Approach 2)

Initialize $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta}_0$, $\mathcal{D}_{\text{ref}} \leftarrow \mathcal{D}_{\boldsymbol{\theta}_0}$
for each training iteration **do**
  Sample a minibatch $\left\{(\mathbf{x}_i, \mathbf{c}_i)\right\}_{i=1}^{N}$ with $(\boldsymbol{x}_i, \boldsymbol{c}_i) \sim p(\boldsymbol{x}, \boldsymbol{c})$.
  Sample $\{\sigma_i\}_{i=1}^{N}$ with $\sigma_i \sim p(\sigma)$
  Sample $\{\boldsymbol{\epsilon}_i\}_{i=1}^{N}$ with $\boldsymbol{\epsilon}_i \sim \mathcal{N}(\mathbf{0}, \boldsymbol{I})$
  Construct contrastive tuples from the minibatch:
  if Approach 1 **then**
    Set $M \leftarrow N$
    For each i, sample $\tilde{\boldsymbol{x}}_i$ from minibatch samples with label $\tilde{\boldsymbol{c}}_i \neq \boldsymbol{c}_i$
    Define tuples $\{(\boldsymbol{x}_i, \tilde{\boldsymbol{x}}_i, \boldsymbol{c}_i, \sigma_i, \boldsymbol{\epsilon}_i)\}_{i=1}^{M}$
  end if
  if Approach 2 **then**
    Set $M \leftarrow NK$
    For each i and each $k = 1, \ldots, K$, sample $\tilde{\boldsymbol{x}}_{i,k}$ from minibatch samples with label $\tilde{\boldsymbol{c}}_{i,k} \neq \boldsymbol{c}_i$
    Define tuples $\{(\boldsymbol{x}_{i,k}, \tilde{\boldsymbol{x}}_{i,k}, \boldsymbol{c}_{i,k}, \sigma_{i,k}, \boldsymbol{\epsilon}_{i,k})\}$ where $(\boldsymbol{x}_{i,k}, \boldsymbol{c}_{i,k}) = (\boldsymbol{x}_i, \boldsymbol{c}_i)$ and $(\sigma_{i,k}, \boldsymbol{\epsilon}_{i,k}) = (\sigma_i, \boldsymbol{\epsilon}_i)$
  end if
  $\mathcal{L} \leftarrow 0$
  for each tuple $(\boldsymbol{x}, \tilde{\boldsymbol{x}}, \boldsymbol{c}, \sigma, \boldsymbol{\epsilon})$ **do**
    $\Delta(\boldsymbol{x}_w, \sigma, \boldsymbol{\epsilon}, \boldsymbol{c}) := \|\boldsymbol{x} - \mathcal{D}_{\boldsymbol{\theta}}(\boldsymbol{x} + \sigma\boldsymbol{\epsilon}; \sigma, \boldsymbol{c})\|_2^2 - \|\boldsymbol{x} - \mathcal{D}_{\text{ref}}(\boldsymbol{x} + \sigma\boldsymbol{\epsilon}; \sigma, \boldsymbol{c})\|_2^2$
    $\Delta(\boldsymbol{x}_l, \sigma, \boldsymbol{\epsilon}, \boldsymbol{c}) := \|\tilde{\boldsymbol{x}} - \mathcal{D}_{\boldsymbol{\theta}}(\tilde{\boldsymbol{x}} + \sigma\boldsymbol{\epsilon}; \sigma, \boldsymbol{c})\|_2^2 - \|\tilde{\boldsymbol{x}} - \mathcal{D}_{\text{ref}}(\tilde{\boldsymbol{x}} + \sigma\boldsymbol{\epsilon}; \sigma, \boldsymbol{c})\|_2^2$
    if CC-DPO **then**
      $\mathcal{L}$ += $-\log \text{Sigmoid}(\beta w(\sigma)(-\Delta(\boldsymbol{x}_w, \sigma, \boldsymbol{\epsilon}, \boldsymbol{c}) + \Delta(\boldsymbol{x}_l, \sigma, \boldsymbol{\epsilon}, \boldsymbol{c})))$
    end if
    if CCA **then**
      $\mathcal{L}$ −= $\log \text{Sigmoid}\big(-\beta w(\sigma)\Delta(\boldsymbol{x}_w, \sigma, \boldsymbol{\epsilon}, \boldsymbol{c})\big) + \lambda \log \text{Sigmoid}\big(\beta w(\sigma)\Delta(\boldsymbol{x}_l, \sigma, \boldsymbol{\epsilon}, \boldsymbol{c})\big)$
    end if
  end for
  $\boldsymbol{\theta} \leftarrow \boldsymbol{\theta} - \alpha\nabla_{\boldsymbol{\theta}}\mathcal{L}$
end for
return $\mathcal{D}_{\boldsymbol{\theta}}$

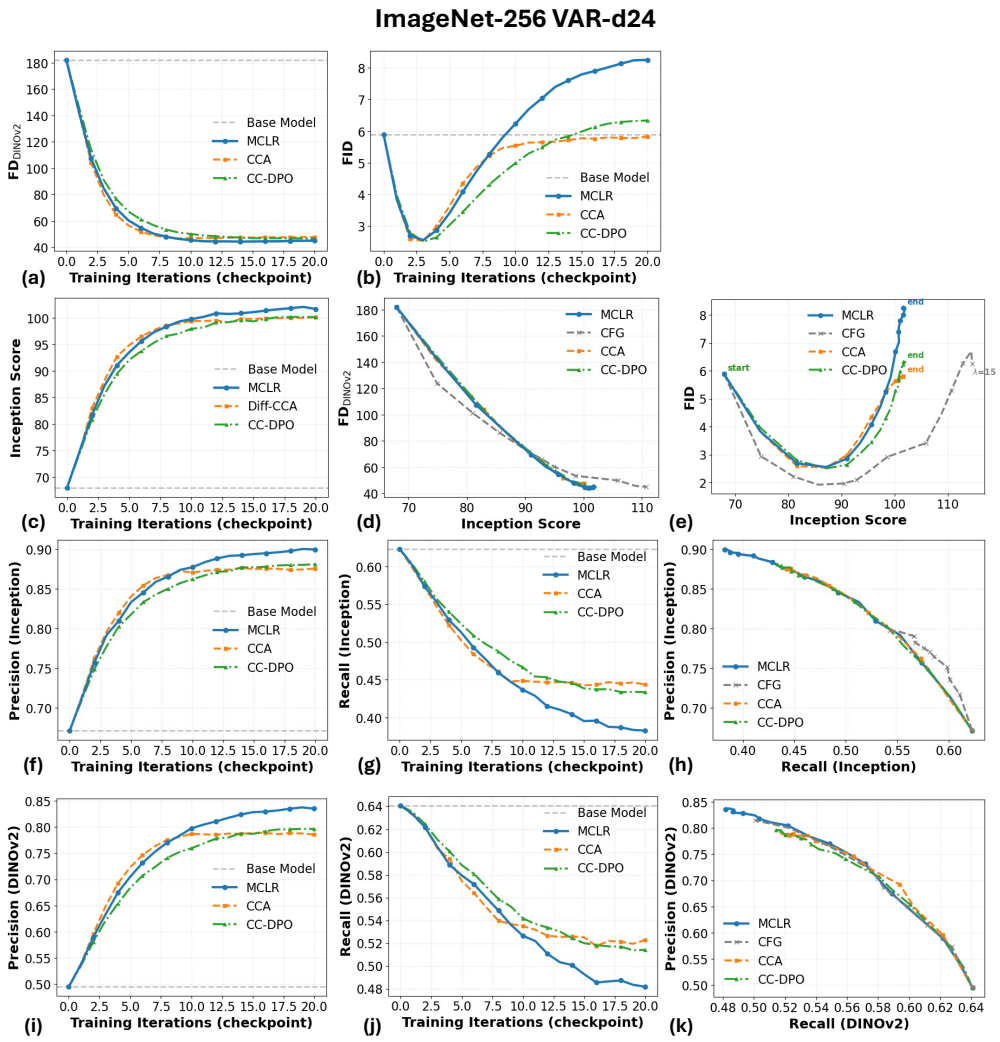


Figure 9: **Quantitative Results for VAR-d24 trained on ImageNet-256×256.** (a), (b), (c), (f), (g), (i), (j) show the evolution of $\text{FD}_{\text{DINOv2}}$, FID, Inception Score, Precision (calculated with Inception features), and Recall (calculated with Inception features), Precision (calculated with DINOv2 features), and Recall (calculated with DINOv2 features), respectively, as functions of training iterations. (d), (e) show the FD–IS trade-offs, while (h), (k) depict the Precision–Recall trade-offs calculated with Inception and DINOv2 features respectively. We evaluate classifier-free guidance (CFG) scales $\gamma \in \{0.5, 0.8, 1.1, 1.5, 1.7, 2.0, 2.5, 3.0, 4.0, 5.0, 7.0, 10.0, 15.0\}$.



Figure 10: **Comparison between CFG and MCLR for Golden Retriever (Class 207).** (a,b) demonstrate the progressive evolution of generated samples on ImageNet-64×64 and ImageNet-512×512, respectively. Increasing the CFG scale γ (top rows) and progressive MCLR training (bottom rows) produce similar effects, both enhancing class-specific structures in the generated images.



Figure 11: **Comparison between CFG and MCLR for Warthog (Class 343).** (a,b) demonstrate the progressive evolution of generated samples on ImageNet-64×64 and ImageNet-512×512, respectively. Increasing the CFG scale γ (top rows) and progressive MCLR training (bottom rows) produce similar effects, both enhancing class-specific structures in the generated images.

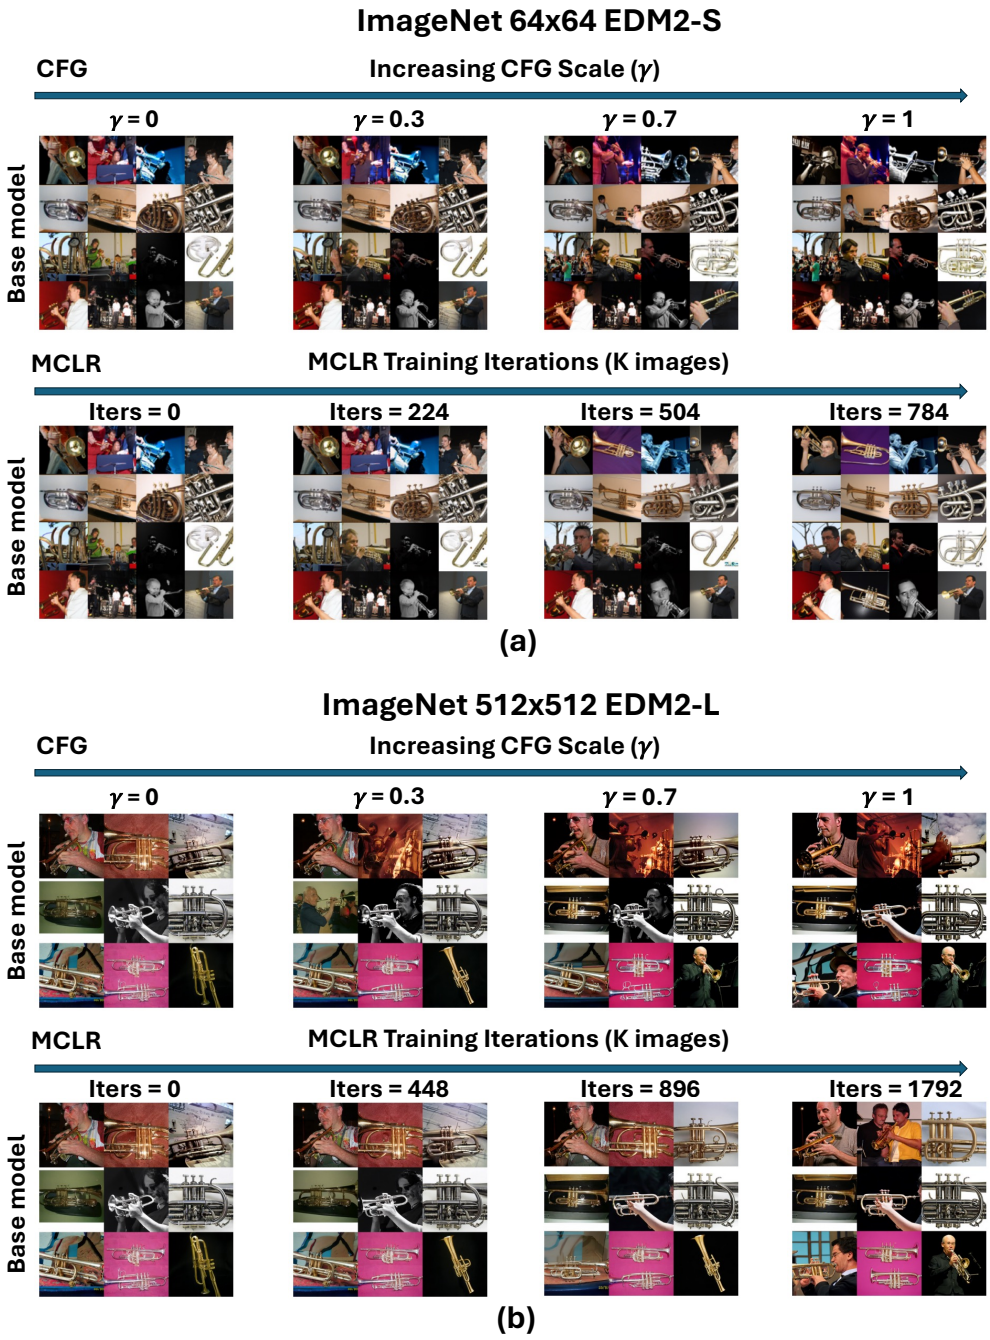


Figure 12: **Comparison between CFG and MCLR for Trumpet (Class 513).** (a,b) demonstrate the progressive evolution of generated samples on ImageNet-64×64 and ImageNet-512×512, respectively. Increasing the CFG scale γ (top rows) and progressive MCLR training (bottom rows) produce similar effects, both enhancing class-specific structures in the generated images.



Figure 13: **Comparison between CFG and MCLR for Cowboy Hat (Class 515).** (a,b) demonstrate the progressive evolution of generated samples on ImageNet-64×64 and ImageNet-512×512, respectively. Increasing the CFG scale γ (top rows) and progressive MCLR training (bottom rows) produce similar effects, both enhancing class-specific structures in the generated images.



Figure 14: **Comparison between CFG and MCLR for Lab Coat (Class 617).** (a,b) demonstrate the progressive evolution of generated samples on ImageNet-64×64 and ImageNet-512×512, respectively. Increasing the CFG scale γ (top rows) and progressive MCLR training (bottom rows) produce similar effects, both enhancing class-specific structures in the generated images.

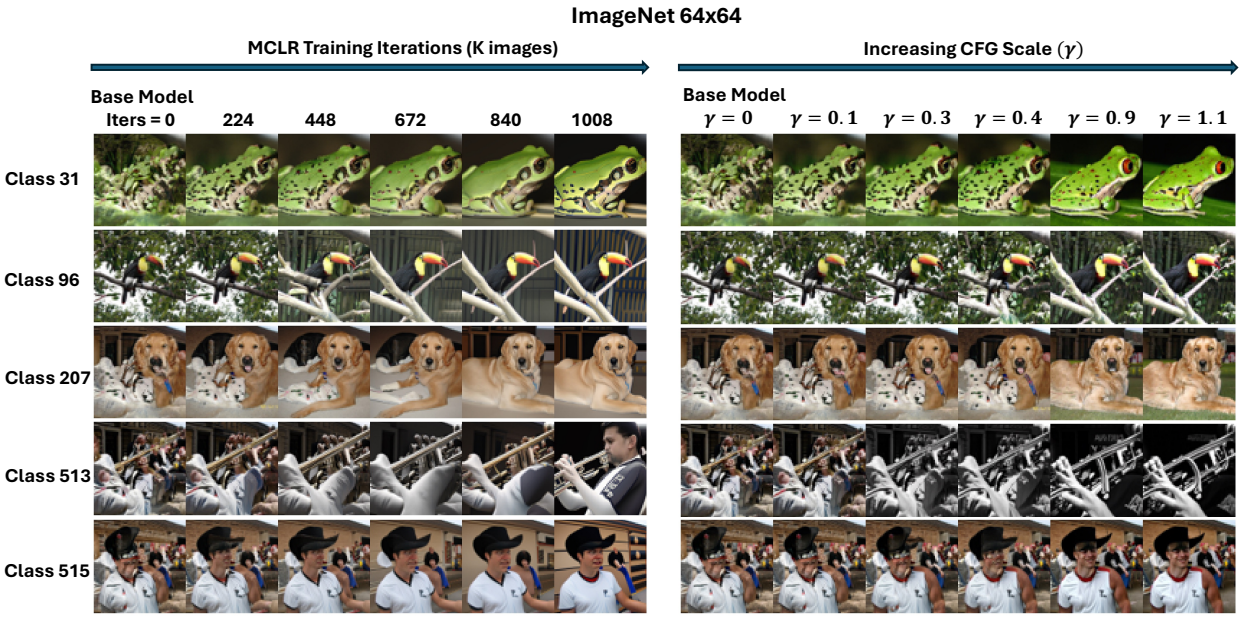


Figure 15: **Comparison between CFG and MCLR.** Left and right figures demonstrate the progressive evolution of generated samples for MCLR and CFG on ImageNet-64×64, respectively, with all images initialized from the same random noise. Increasing the CFG scale γ and progressive MCLR training produce similar effects, both enhancing class-specific structures in the generated images.

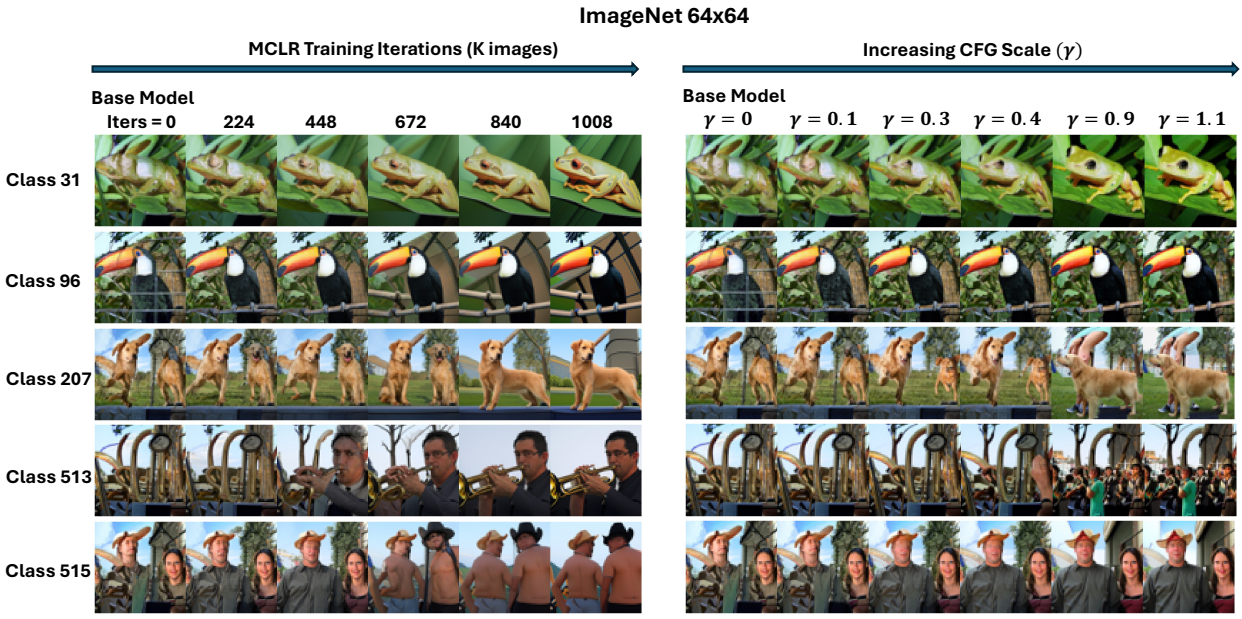


Figure 16: **Comparison between CFG and MCLR.** Same as Figure 15, but with a different initial random noise.

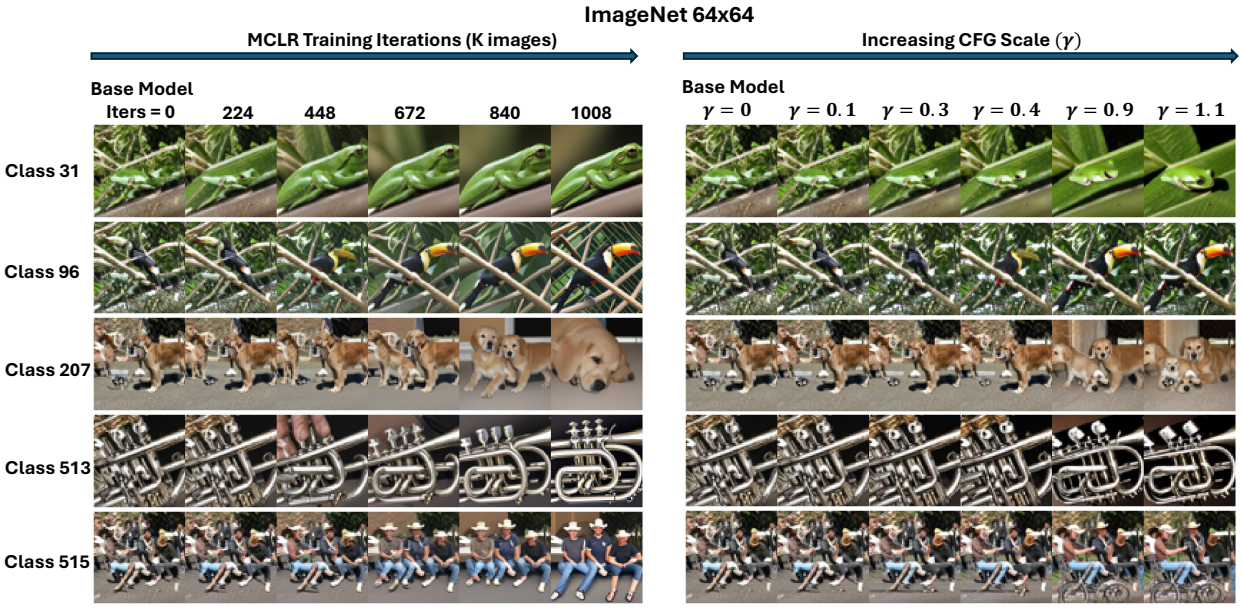


Figure 17: **Comparison between CFG and MCLR.** Same as Figure 15, but with a different initial random noise.

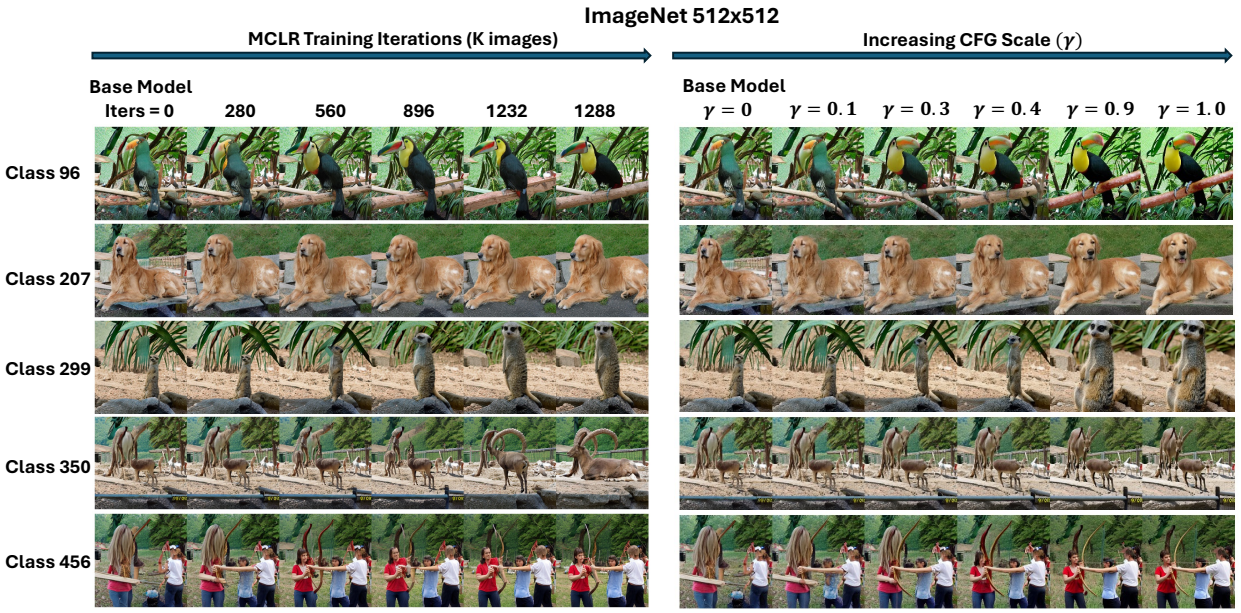


Figure 18: **Comparison between CFG and MCLR.** Same as Figure 15, but for ImageNet-512x512 with a different initial random noise.

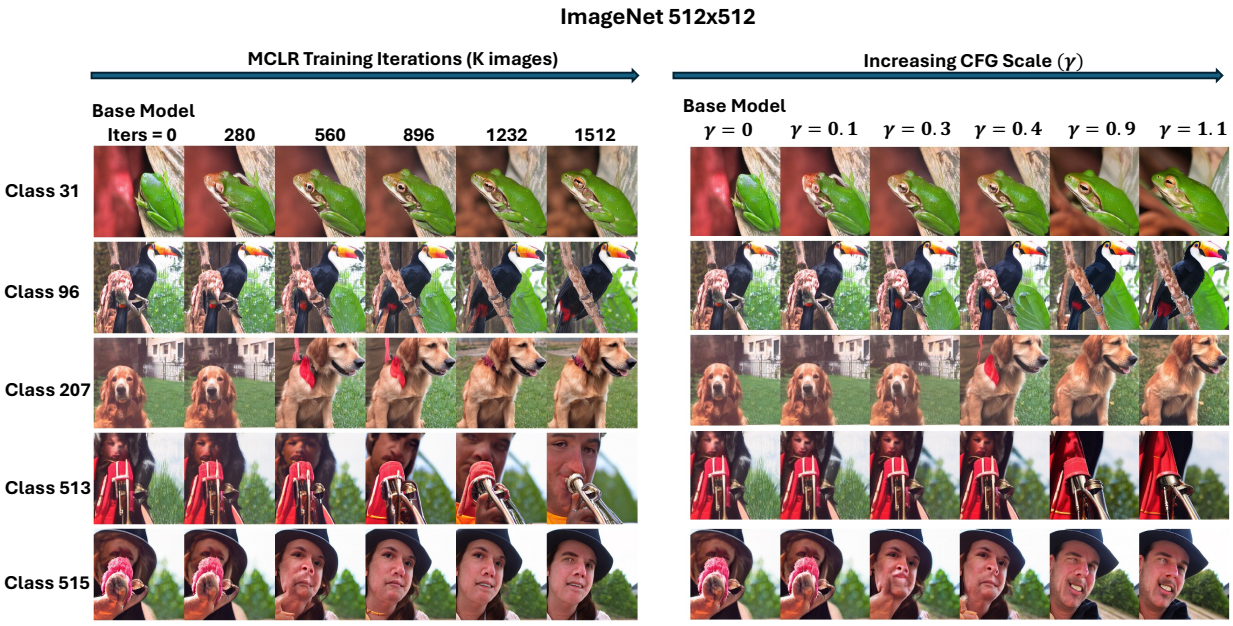


Figure 19: **Comparison between CFG and MCLR.** Same as Figure 15, but for ImageNet-512x512 with a different initial random noise.

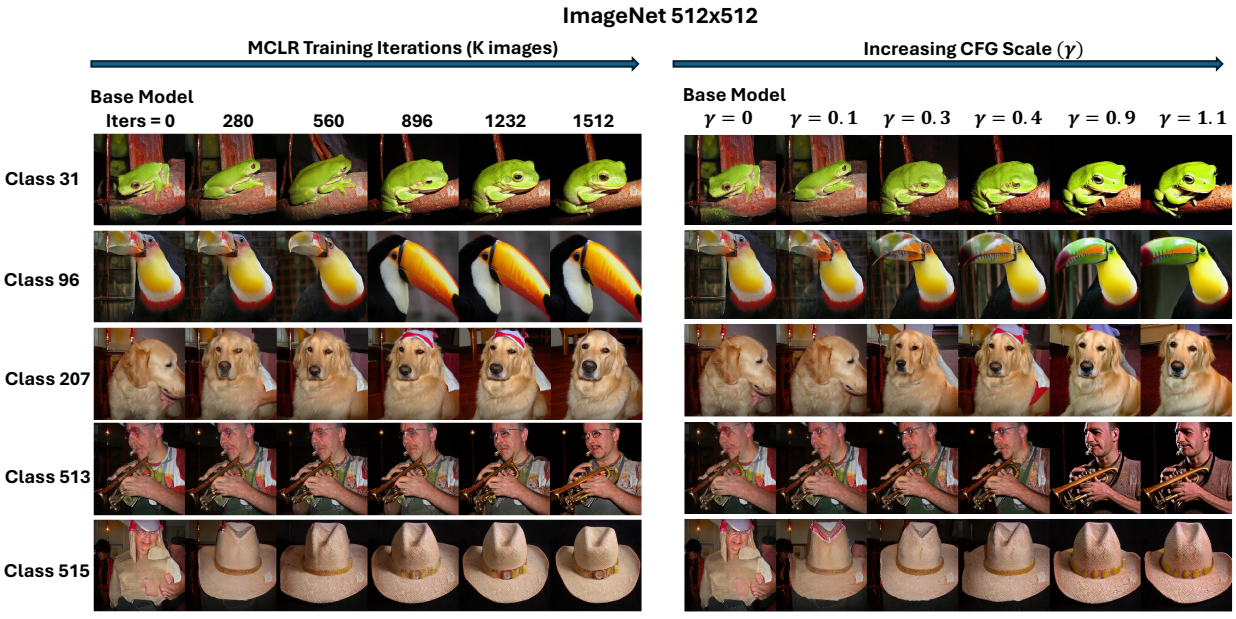


Figure 20: **Comparison between CFG and MCLR.** Same as Figure 15, but for ImageNet-512x512 with a different initial random noise.

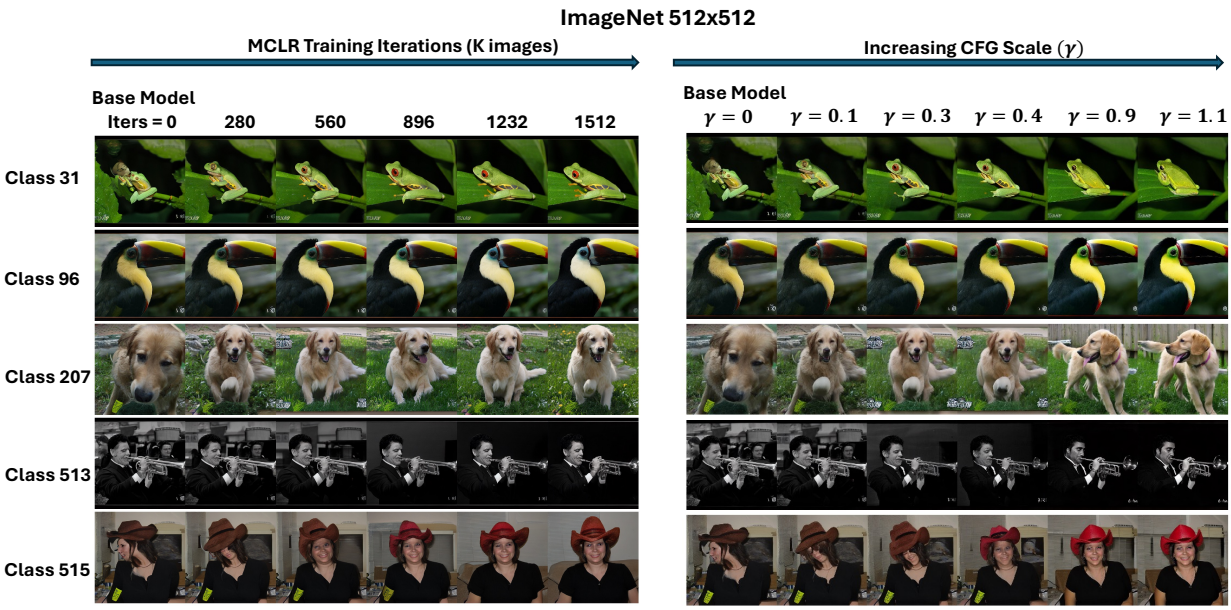


Figure 21: **Comparison between CFG and MCLR.** Same as Figure 15, but for ImageNet-512x512 with a different initial random noise.

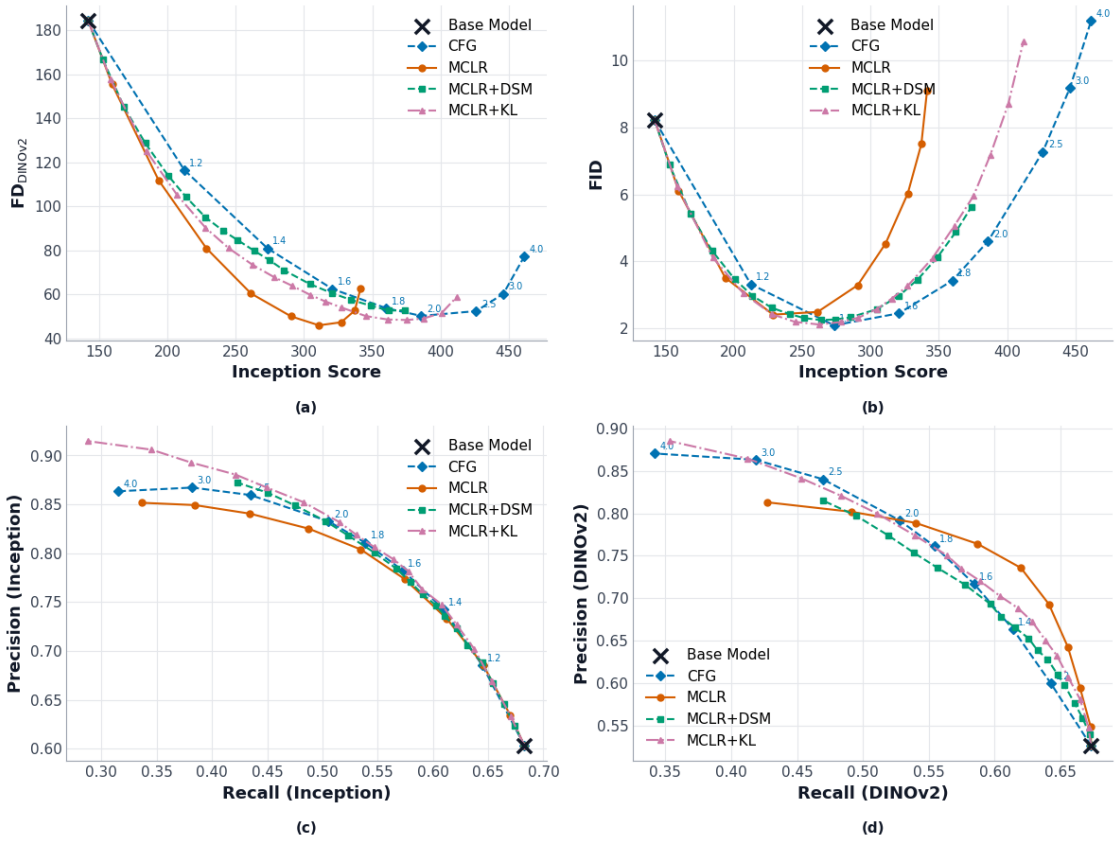


Figure 22: **Quantitative Results for SiT-XL/2 Trained on ImageNet-256×256.** (a), (b), (c), (d) show the FD–IS trade-offs and the Precision–Recall trade-offs calculated with Inception and DINOv2 features respectively. We evaluate classifier-free guidance (CFG) scales $\gamma \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.7, 0.9, 1, 1.5, 2.0, 3.0\}$. For the DSM- and KL-regularized variants, we set $\beta = 1$.

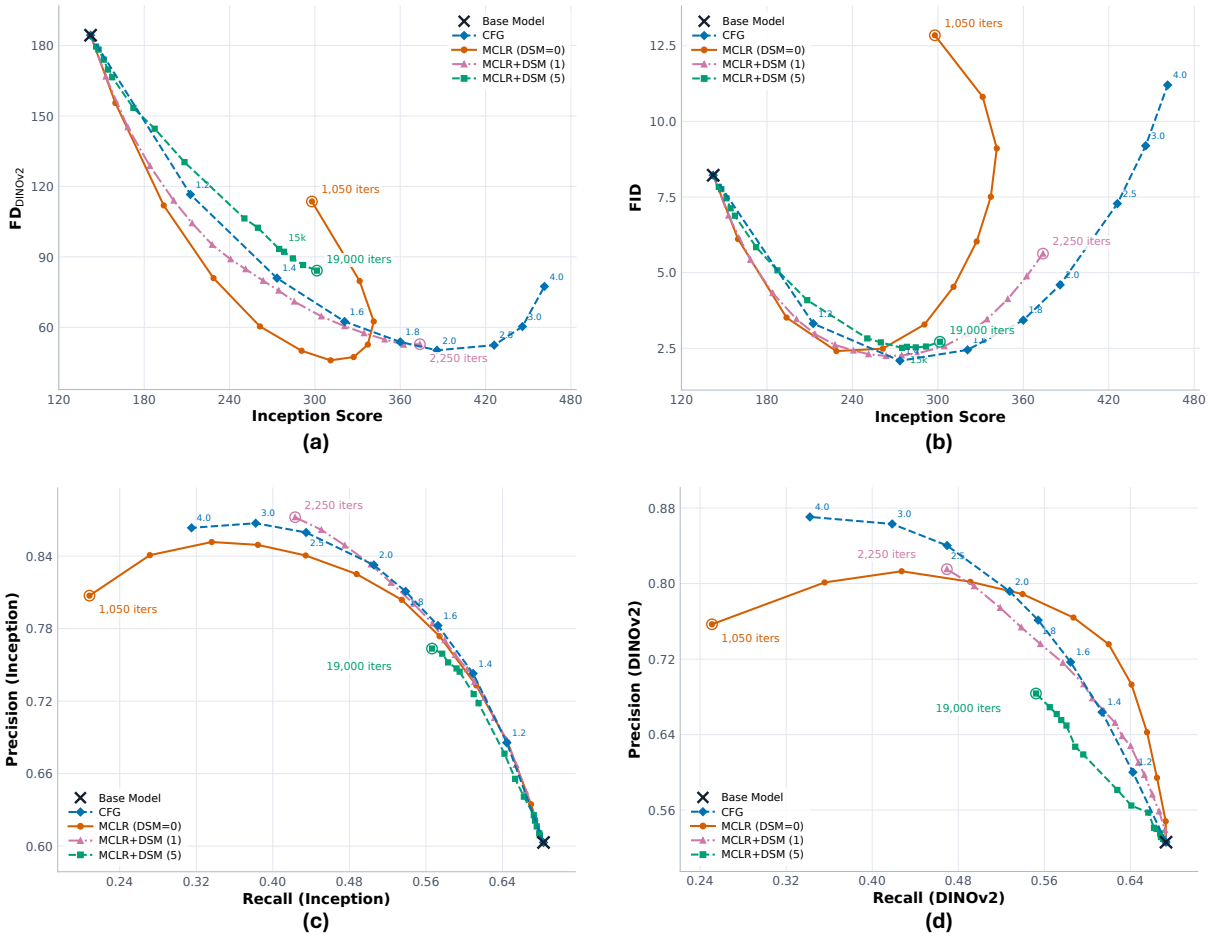


Figure 23: **Quantitative Results for SiT-XL/2 with DSM Regularization Trained on ImageNet-256×256.** (a), (b), (c), (d) show the FD–IS trade-offs and the Precision–Recall trade-offs calculated with Inception and DINOv2 features respectively. We compare MCLR alone ($\beta = 0$) with DSM-regularized variants using $\beta \in \{1, 5\}$, and the annotations indicate the corresponding numbers of fine-tuning iterations. We evaluate classifier-free guidance (CFG) scales $\gamma \in \{0.1, 0.2, 0.3, 0.4, 0.5, 0.7, 0.9, 1, 1.5, 2.0, 3.0\}$.

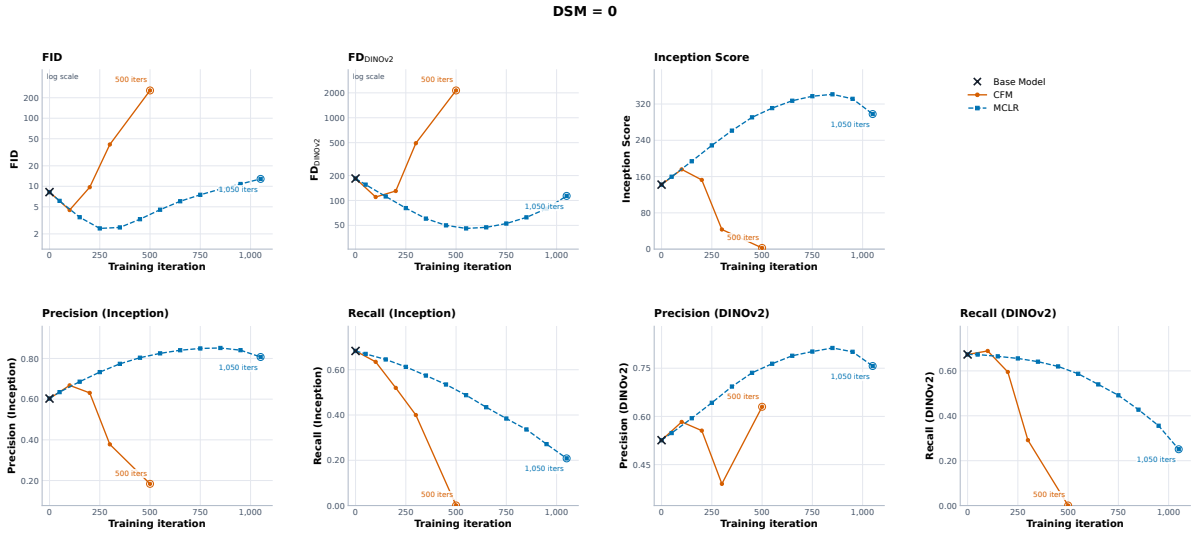


Figure 24: **Training Dynamics of MCLR and CFM without DSM Regularization** ($\beta = 0$)**.** The panels report FID, $\text{FD}_{\text{DINOv2}}$, Inception Score, and Inception- and DINOv2-based precision and recall over fine-tuning iterations. The black cross denotes the base model. MCLR produces larger and more sustained improvements, whereas CFM rapidly degrades after an initial improvement.

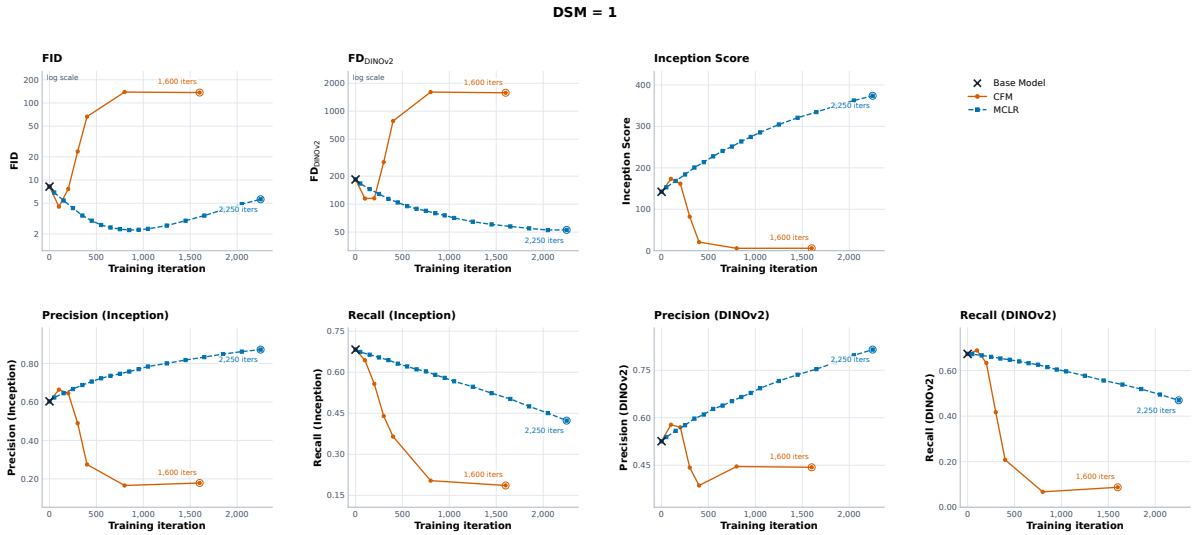


Figure 25: **Training Dynamics of MCLR and CFM with DSM Regularization** ($\beta = 1$)**.** The panels report FID, $\text{FD}_{\text{DINOv2}}$, Inception Score, and Inception- and DINOv2-based precision and recall over fine-tuning iterations. The black cross denotes the base model. MCLR achieves substantially better and more stable performance than CFM.

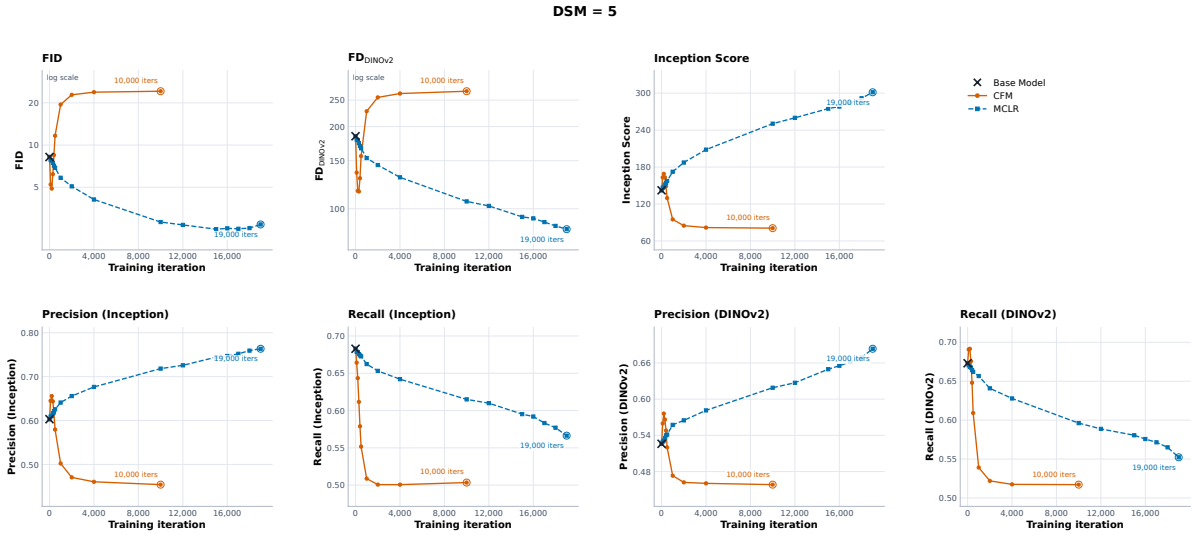


Figure 26: **Training Dynamics of MCLR and CFM with DSM Regularization** ($\beta = 5$)**.** The panels report FID, $\text{FD}_{\text{DINOv2}}$, Inception Score, and Inception- and DINOv2-based precision and recall over fine-tuning iterations. The black cross denotes the base model. MCLR continues to improve over a substantially longer training horizon, whereas CFM eventually deteriorates.

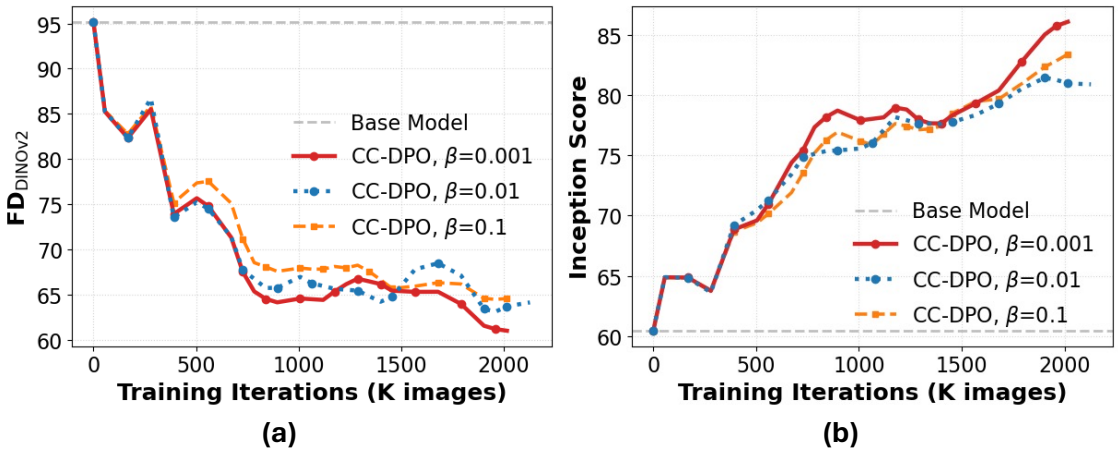


Figure 27: **Effects of β in CC-DPO.** (a,b) compare fine-tuning diffusion models using CC-DPO with different β evaluated on the EDM2-S model trained on ImageNet-64 $\times$ 64. Performance is reported in terms of $\text{FD}_{\text{DINOv2}}$ and Inception Score. CC-DPO achieves stable performance under different β, with a smaller β leading to marginally better performance on selected metrics.

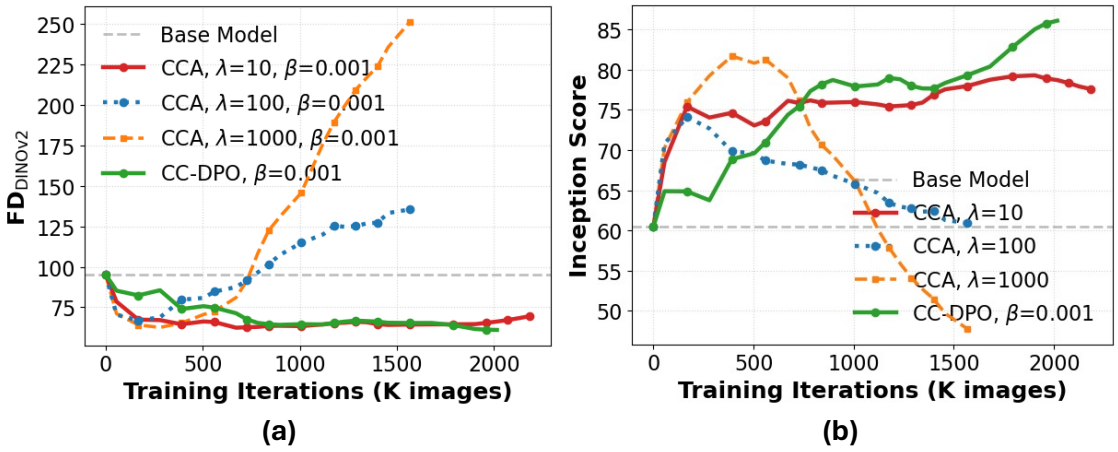


Figure 28: **Effects of λ in CCA.** (a,b) compare fine-tuning diffusion models using CCA with different λ evaluated on the EDM2-S model trained on ImageNet-64 $\times$ 64. Performance is reported in terms of $\text{FD}_{\text{DINOv2}}$ and Inception Score. Note with the same β, CC-DPO achieves comparable performance as CCA while at the same time requiring fewer hyperparameters.