

000 001 002 003 004 005 LEARNING ROBUST DIFFUSION MODELS FROM 006 IMPRECISE SUPERVISION 007 008 009

010 **Anonymous authors**
011 Paper under double-blind review
012
013
014
015
016
017
018
019
020
021
022
023
024

ABSTRACT

025 Conditional diffusion models have achieved remarkable success in various gener-
026 ative tasks recently, but their training typically relies on large-scale datasets that
027 inevitably contain imprecise information in conditional inputs. Such supervision,
028 often stemming from noisy, ambiguous, or incomplete labels, will cause condition
029 mismatch and degrade generation quality. To address this challenge, we propose
030 *DMIS*, a unified framework for training robust Diffusion Models from Imprecise
031 Supervision, which is the first systematic study within diffusion models. Our
032 framework is derived from likelihood maximization and decomposes the objective
033 into generative and classification components: the generative component models
034 imprecise-label distributions, while the classification component leverages a dif-
035 fusion classifier to infer class-posterior probabilities, with its efficiency further
036 improved by an optimized timestep sampling strategy. Extensive experiments
037 on diverse forms of imprecise supervision, covering tasks of image generation,
038 weakly supervised learning, and noisy dataset condensation demonstrate that *DMIS*
039 consistently produces high-quality and class-discriminative samples.
040

041 1 INTRODUCTION

042 Diffusion models (DMs) (Ho et al., 2020; Song et al., 2020; Karras et al., 2022) have emerged as
043 powerful generative frameworks that have unprecedented capabilities in generating realistic data
044 (He et al., 2025; Yang et al., 2024; Ho et al., 2022). With the classifier guidance (Ho & Salimans,
045 2022; Dhariwal & Nichol, 2021), conditional diffusion models (CDMs) extended the capabilities of
046 DMs by conditioning the generation process on additional information, such as text descriptions or
047 class labels. These models have demonstrated remarkable performance in various tasks, including
048 text-to-image synthesis (Rombach et al., 2022; Saharia et al., 2022), image inpainting (Zhao et al.,
049 2024; Corneau et al., 2024), and super-resolution (Esser et al., 2024; Xie et al., 2025).
050

051 Unfortunately, the conditioning information required by CDMs is often imprecise in real-world
052 scenarios. When sourced from the internet or obtained through crowdsourcing, such information
053 can be affected by factors such as privacy constraints or limited annotator expertise, leading to
054 various imperfections. In particular, the conditioning data may contain noise, exhibit ambiguity, or
055 suffer from missing and incomplete annotations. We refer to such cases collectively as imprecise
056 supervision (Chen et al., 2024a), where the provided conditioning information is not fully aligned
057 with the true underlying labels. This includes scenarios such as noisy-label data (Li et al., 2017; Wei
058 et al., 2021), partial-label data (Wang et al., 2025b;a), and supplementary-unlabeled data (He et al.,
059 2023). These forms of imprecise supervision can introduce incorrect inductive biases during training
060 and severely affect the reliability and generalization of CDMs.
061

062 To address this, several recent studies have proposed adaptations of diffusion models to handle
063 imprecise supervision, such as noise-robust diffusion models (Na et al., 2024; Li et al., 2024) and
064 positive-unlabeled diffusion models (Takahashi et al., 2025). However, these approaches often focus
065 on specific types of imprecise supervision. Moreover, many of them rely on strong external priors to
066 guide the learning process. For example, Na et al. (2024) estimated a noise transition matrix using
067 external noisy-label learning methods, and Li et al. (2024) required risk confidence scores associated
068 with noisy samples. These diffusion-based methods not only rely on prior knowledge from data
069 or previous techniques, but are also designed with task-specific architectures for particular types
070 of supervision. Such reliance and structural complexity limit their applicability and efficiency in
071

054 practice. There remains a need for a unified framework that can robustly train CDMs under diverse
 055 forms of imprecise supervision without requiring strong prior assumptions.
 056

057 In this paper, to train a robust CDM in a unified manner, we first formulate the overall learning
 058 objective as a likelihood maximization problem (Section 4.1). Then we decompose this objective
 059 into a generative term that models the imprecise data distribution (Section 4.2) and a classification
 060 term that infers posterior label probabilities from imprecise supervision (Section 4.3). During
 061 generative modeling, we show that the imprecise-label conditional score can be expressed as a linear
 062 combination of clean-label conditional scores, weighted by the corresponding posterior probabilities.
 063 Building on this insight, we propose a weighted denoising score matching objective, which enables
 064 the model to achieve label-conditioned learning without requiring clean annotations. Finally, to
 065 reduce the time complexity of posterior inference, we further introduce an efficient timestep sampling
 066 strategy (Section 5). Extensive experiments across multiple tasks, including image generation, weakly
 067 supervised learning, and noisy dataset condensation show that CDMs trained with our framework
 068 not only achieve strong generative quality but also produce class-discriminative samples. Our
 069 contributions are summarized as follows:
 070

- 071 • We propose a unified diffusion framework for training CDMs under diverse forms of
 072 imprecise supervision, which is the first exploration in the diffusion model field.
- 073 • To improve efficiency, we develop an optimized timestep sampling strategy for diffusion
 074 classifiers that greatly reduces the computation cost without compromising performance.
- 075 • Building on this framework, we pioneer the study of noisy dataset condensation, a practical
 076 yet previously unexplored setting, and establish a solid baseline for future research.
- 077 • Extensive experiments on image generation, weakly supervised learning, and noisy dataset
 078 condensation demonstrate the effectiveness and versatility of our unified framework.

079 2 RELATED WORK

080 2.1 ROBUST DIFFUSION MODELS

081 Training conditional diffusion models under limited or imperfect supervision is still relatively under-
 082 explored. Recent work has begun to address specific forms of weak or noisy information, such as
 083 noise-robust diffusion models (Na et al., 2024; Li et al., 2024) that mitigate corrupted labels, and
 084 positive-unlabeled diffusion models (Takahashi et al., 2025) that combine positive samples with
 085 large unlabeled corpora to approximate conditional distributions. We take a different perspective:
 086 rather than tailoring objectives to a single type of imperfect label, we formulate a unified conditional
 087 score-learning framework that can be instantiated under multiple imprecise-label regimes.
 088

089 2.2 IMPRECISE LABEL LEARNING

090 Imprecise label learning studies supervision that is incomplete, ambiguous, or corrupted relative to
 091 clean ground-truth labels. Canonical settings include partial-label learning (Feng et al., 2020; Wu
 092 et al., 2022; Tian et al., 2023; Lv et al., 2020; Wang et al., 2025b), where each instance is associated
 093 with a candidate label set containing the true label; semi-supervised learning (Berthelot et al., 2019b;
 094 Zhang et al., 2021a; Yang et al., 2022; Wang et al., 2022c), where only a subset of samples are
 095 labeled; and noisy-label learning (Han et al., 2018; Wei et al., 2021; Han et al., 2020), where observed
 096 labels are corrupted versions of the true labels. Beyond these settings, mixture imprecise-label
 097 learning (Chen et al., 2024a; Zhang et al., 2020; Wei et al., 2023; Shukla et al., 2023; Xie et al., 2024)
 098 combines several forms of imprecision in a single framework. Our work can be viewed as lifting
 099 these ideas from discriminative prediction to conditional score modeling, providing a generative view
 100 of learning with heterogeneous imprecise supervision.
 101

102 2.3 DATASET CONDENSATION

103 Dataset distillation (DD) (Wang et al., 2018) compresses a large labeled dataset into a compact
 104 synthetic set that preserves task-relevant information, thereby reducing training cost while maintaining
 105 competitive accuracy. Bi-level optimization methods learn synthetic data whose training signals
 106

108 match those of the original data via gradient, trajectory, or meta-model matching (Zhao et al., 2021;
 109 Kim et al., 2022; Cazenavette et al., 2022a; Cui et al., 2023; Wang et al., 2018; Loo et al., 2022), often
 110 achieving high fidelity at nontrivial computational cost. Distribution-matching approaches instead
 111 align statistics in pixel, feature, or kernel space (Wang et al., 2022b; Sajedi et al., 2023; Xue et al.,
 112 2025; Yin et al., 2024; Sun et al., 2024; Shao et al., 2024; Yin & Shen, 2024), enabling more scalable
 113 DD. While DD primarily targets data efficiency under clean labels, our framework instead focuses on
 114 robustness to imprecise supervision.

3 BACKGROUND

Diffusion Models. Let $\mathcal{X} \subseteq \mathbb{R}^d$ denote the d -dimensional input space. Given a clean input $\mathbf{x} := \mathbf{x}_0$ from the real data distribution with density $q(\mathbf{x}_0)$, the forward diffusion process corrupts the data into a sequence of noisy samples $\{\mathbf{x}_t\}_{t=1}^T$ ¹ by gradually adding Gaussian noise with a fixed scaling schedule $\{\alpha_t\}_{t=1}^T$ and a fixed noise schedule $\{\sigma_t\}_{t=1}^T$, as defined by

$$q(\mathbf{x}_t | \mathbf{x}_0) = \mathcal{N}(\mathbf{x}_t; \alpha_t \mathbf{x}_0, \sigma_t^2 \mathbf{I}), \quad (1)$$

where \mathbf{I} denotes the identity matrix and $\mathcal{N}(\mathbf{x}; \boldsymbol{\mu}, \boldsymbol{\Sigma})$ denotes the Gaussian density with mean $\boldsymbol{\mu}$ and covariance matrix $\boldsymbol{\Sigma}$. Assuming that the signal-to-noise ratio $\text{SNR}(t) = \alpha_t^2 / \sigma_t^2$ decreases monotonically over time, the sample \mathbf{x}_t becomes increasingly noisier during the forward process. The scaling and noise schedules are prescribed such that \mathbf{x}_T nearly follows an isotropic Gaussian distribution. The reverse process for Eq. (1) is defined as a Markov chain, which aims to approximate $q(\mathbf{x}_0)$ by gradually denoising from the standard Gaussian distribution $p(\mathbf{x}_T) = \mathcal{N}(\mathbf{x}_T; \mathbf{0}, \mathbf{I})$:

$$p_\theta(\mathbf{x}_{0:T}) = p(\mathbf{x}_T) \prod_{t=1}^T p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t), \quad (2)$$

$$p_\theta(\mathbf{x}_{t-1} | \mathbf{x}_t) = \mathcal{N}(\mathbf{x}_{t-1}; \boldsymbol{\mu}_\theta(\mathbf{x}_t, t), \tilde{\sigma}_t^2 \mathbf{I}), \quad (3)$$

where $\boldsymbol{\mu}_\theta$ is generally parameterized by a time-conditioned score prediction network $\mathbf{s}_\theta(\mathbf{x}_t, t)$ (Song et al., 2020; 2021; Song & Ermon, 2019; 2020):

$$\boldsymbol{\mu}_\theta(\mathbf{x}_t, t) = \frac{\alpha_{t-1}}{\alpha_t} \left[\mathbf{x}_t + \left(\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2 \right) \mathbf{s}_\theta(\mathbf{x}_t, t) \right]. \quad (4)$$

The reverse process can be learned by optimizing the variational lower bound on log-likelihood as

$$\log p_\theta(\mathbf{x}) \geq -\mathbb{E}_t \left[w_t \left\| \mathbf{s}_\theta(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t) \right\|_2^2 \right] + C_1, \quad (5)$$

where \mathbb{E} denotes the expectation, $w_t = \frac{\sigma_t^2}{2} \left(\frac{\sigma_t^2 \alpha_{t-1}^2}{\sigma_{t-1}^2 \alpha_t^2} - 1 \right)$, and C_1 is a constant that is typically small and can be dropped (Song et al., 2020). The expectation term is called the *score matching loss* (Kingma et al., 2021), where $\nabla \log q_t(\mathbf{x}_t)$ is the gradient of data density at \mathbf{x}_t in data space.

The above definition can be reformulated to match other commonly used diffusion models, such as those in Ho et al. (2020), Karras et al. (2022) and Song et al. (2020). The corresponding conversions are detailed in Appendix B.1. For clarity, we adopt the the elucidated diffusion model (EDM) (Karras et al., 2022) as the default diffusion model throughout this paper, as it offers a unified structure and well-optimized parameterization.

Imprecise Supervision. Imprecise-label data typically refers to settings where the true label is not directly available, and instead only imprecise label information is provided. Let $\mathcal{Y} = [c] := \{1, \dots, c\}$ represent the label space with c distinct classes. In this work, we primarily focus on three representative forms of imprecise supervision that have been widely studied in the literature:

- *Partial-label data*, where each instance X is associated with a candidate label set $S \subset [c]$ that is guaranteed to contain the true label Y , i.e., $p(Y \in S | X, S) = 1$. This setting is widely studied in partial-label learning (Tian et al., 2023).
- *Supplementary-unlabeled data*, consisting of a small labeled subset (X^1, Y^1) together with a large number of unlabeled samples (X^u, \emptyset) . This scenario is the focus of semi-supervised learning (Yang et al., 2022), which aims to exploit unlabeled data to improve generalization.

¹We use the subscript t of the sample \mathbf{x} to denote the noisy version of the sample at timestep t .

- 162 • *Noisy-label data*, where the observed label \hat{Y} is a corrupted version of the underlying true label
 163 Y , modeled by a conditional distribution $p(\hat{Y}|X, Y)$. This gives rise to noisy-label learning (Han
 164 et al., 2020), which seeks to build models robust to label corruption.

166 **4 METHODOLOGY**

169 In this section, we first introduce the unified learning objective that integrates generative and classifi-
 170 cation components. Then we elaborate on the formulation and optimization of these components.

172 **4.1 UNIFIED LEARNING OBJECTIVE**

174 To robustly learn a diffusion model with learnable parameters θ under imprecise supervision (denoted
 175 as $Z \subseteq \mathcal{Y}$), we treat the true label Y as a latent variable and maximize the likelihood of the joint
 176 distribution of the input X and Z . By the maximum likelihood principle, our objective is to find

$$177 \theta^* = \arg \max_{\theta} \log p_{\theta}(X, Z) = \arg \max_{\theta} \log \sum_Y p_{\theta}(X, Y, Z), \quad (6)$$

179 where θ^* denotes the optimal parameter. Eq. (6) involves the log of the marginalization over latent
 180 variables and cannot generally be solved in closed form. To circumvent this intractability, we instead
 181 maximize a variational lower bound on the marginal log-likelihood:

$$182 \theta^n = \arg \max_{\theta} \mathbb{E}_{p_{\phi}(Y|X, Z)} [\log p_{\theta}(X, Y, Z)] \\ 183 = \arg \max_{\theta} \left\{ \log p_{\theta}(X|Z) + \mathbb{E}_{p_{\phi}(Y|X, Z)} [\log p_{\theta}(Y|X, Z)] \right\}, \quad (7)$$

186 where θ^n denotes the n -th estimate of θ , and ϕ is instantiated as the exponential moving average
 187 (EMA) of θ over its 1st through $(n-1)$ iterates. A complete derivation of this variational lower
 188 bound is provided in Appendix B.3. From Eq. (7), we can observe that maximizing the marginal
 189 likelihood can be performed from generative and classification perspectives. The former focuses on
 190 modeling the data distribution conditioned on the imprecise supervision, while the latter aims to infer
 191 the posterior distribution based on the feature and the imprecise label. In this paper, we adopt the
 192 commonly used class-conditional setting, where the generation of the imprecise label Z is assumed
 193 to be independent of the input X given the true label Y (Yao et al., 2020; Wen et al., 2021).

195 **4.2 GENERATIVE OBJECTIVE: MODELING THE IMPRECISE DATA DISTRIBUTION**

197 Since samples are assumed to be independent of each other, we present the analysis in this and the
 198 following subsections using a single sample (\mathbf{x}, z) for notational clarity, with the final objective
 199 computed over the entire dataset. Following the standard formulation of diffusion models in Eq. (5),
 200 we parameterize the conditional generative process $p_{\theta}(\mathbf{x}|z)$ using a score network $\mathbf{s}_{\theta}(\mathbf{x}_t, z, t)$. The
 201 corresponding variational lower bound on the conditional log-likelihood is given by

$$202 \log p_{\theta}(\mathbf{x}_0|z) \geq -\mathbb{E}_t \left[w_t \left\| \mathbf{s}_{\theta}(\mathbf{x}_t, z, t) - \nabla_{\mathbf{x}_t} \log q_{t|0}(\mathbf{x}_t|\mathbf{x}_0, z) \right\|_2^2 \right] + C_2, \quad (8)$$

204 where C_2 is another constant. Directly optimizing the score network with this objective on imprecise-
 205 label data would lead it to converge to the score of the imprecise conditional distribution.

207 **Remark 1.** Let $\hat{\theta}$ denote the parameters obtained by maximizing the lower bound in Eq. (8) using
 208 denoising score matching. In this case, the learned score function satisfies $\mathbf{s}_{\hat{\theta}}(\mathbf{x}_t, z, t) = \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t|z)$
 209 for all $\mathbf{x}_t \in \mathcal{X}$, $z \subseteq \mathcal{Y}$, and $t \in [T]$. However, since $q_t(\mathbf{x}_t|z)$ corresponds to the imprecise-label
 210 density, the resulting generation is biased and thus fails to fully recover the true data distribution. The
 211 derivation and visualization of this bias is deferred to Appendix B.5.

212 Therefore, to align the learned score with the clean-label conditional score, we propose modifying
 213 the objective to correct the gradient signal from score matching (Kingma et al., 2021). **Building on**
 214 **the linear relationship between clean- and noisy-label conditional scores modeled by Na et al. (2024)**,
 215 we further derive an explicit relationship that connects imprecise-label conditional scores to their
 216 clean-label counterparts.

216 **Theorem 1.** *Under the class-conditional setting, for all $\mathbf{x}_t \in \mathcal{X}$, $z \subseteq \mathcal{Y}$, and $t \in [T]$,*

$$218 \quad \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | z) = \sum_{y=1}^c p(y | \mathbf{x}_t, z) \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | y). \quad (9)$$

220 The formal proof is in Appendix B.6. Since $p(y | \mathbf{x}_t, z) \geq 0$ and $\sum_{y=1}^c p(y | \mathbf{x}_t, z) = 1$, Theorem 1
 221 implies that the imprecise-label conditional score can be expressed as a convex combination of the
 222 clean-label conditional scores, weighted by $p(y | \mathbf{x}_t, z)$. These weights represent the model’s posterior
 223 probability over labels given \mathbf{x}_t and z , implicitly requiring the model to perform classification during
 224 training. To our knowledge, this is the first work to explicitly reveal and exploit the classification
 225 capability of diffusion models within the training process under imprecise supervision.
 226

227 According to Remark 1, directly optimizing the denoising score matching objective in Eq. (8) drives
 228 the score network to approximate the imprecise-label conditional score. However, Theorem 1 shows
 229 that this score can be decomposed as a convex combination of clean-label conditional scores, weighted
 230 by the posterior probability $p(y | \mathbf{x}_t, z)$. Motivated by this insight, we propose a new training objective
 231 that supervises the clean-label score network $\mathbf{s}_\theta(\mathbf{x}_t, y, t)$ through a reweighted aggregation of its
 232 posterior outputs. The resulting weighted denoising score matching loss is

$$233 \quad \mathcal{L}_{\text{Gen}}(\theta) = \mathbb{E}_t \left[w_t \left\| \sum_{y=1}^c p(y | \mathbf{x}_t, z) \mathbf{s}_\theta(\mathbf{x}_t, y, t) - \nabla_{\mathbf{x}_t} \log q_{t|0}(\mathbf{x}_t | \mathbf{x}_0, z) \right\|_2^2 \right]. \quad (10)$$

235 This loss encourages the weighted aggregation of clean-label scores to approximate the imprecise
 236 score derived from data, thereby enabling label-conditioned learning without the need for explicit
 237 clean annotations. The following Proposition 1, with proof provided in Appendix B.7, guarantees
 238 that the optimal solution recovers the clean-label conditional scores:
 239

240 **Proposition 1.** *Let $\theta_{\text{Gen}}^* = \arg \min_{\theta} \mathcal{L}_{\text{Gen}}(\theta)$ be the minimizer of Eq. (10). Then, for all $\mathbf{x}_t \in \mathcal{X}$,
 241 $z \subseteq \mathcal{Y}$, and $t \in [T]$, the learned score function satisfies $\mathbf{s}_{\theta_{\text{Gen}}^*}(\mathbf{x}_t, y, t) = \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | y)$.*

242 4.3 CLASSIFICATION OBJECTIVE: INFERRING LABELS FROM IMPRECISE SIGNALS

244 We assume the class prior to be uniform, i.e., $p(y) = 1/c$. To infer the class-posterior probability
 245 $p_\theta(y | \mathbf{x}_t)$, we adopt a diffusion-based approximation as defined below:

247 **Definition 1** (Approximated Posterior Noised Diffusion Classifier (Chen et al., 2024b)). Assuming
 248 the uniform prior $p(y)$, the class-posterior probability for a noisy input \mathbf{x}_t under a conditional
 249 diffusion model can be derived using Bayes’ rule, as follows:

$$250 \quad p_\theta(y | \mathbf{x}_t) = \frac{p_\theta(\mathbf{x}_t | y)}{\sum_{y'} p_\theta(\mathbf{x}_t | y')} = \frac{\exp\{\log p_\theta(\mathbf{x}_t | y)\}}{\sum_{y'} \exp\{\log p_\theta(\mathbf{x}_t | y')\}}. \quad (11)$$

253 Here, following Chen et al. (2024b), the conditional likelihood $\log p_\theta(\mathbf{x}_t | y)$ is approximated by the
 254 conditional evidence lower bound (ELBO), given by

$$255 \quad \log p_\theta(\mathbf{x}_t | y) \approx - \sum_{\tau=t+1}^{T-1} w_\tau \mathbb{E}_{q(\mathbf{x}_\tau | \mathbf{h}_\theta(\mathbf{x}_t, y, t))} \left[\left\| \mathbf{h}_\theta(\mathbf{x}_\tau, y, \tau) - \mathbf{x}_0 \right\|_2^2 \right], \quad (12)$$

258 where $\mathbf{h}_\theta(\mathbf{x}_\tau, y, \tau) = \frac{\mathbf{x}_\tau}{\alpha_\tau} + \frac{\sigma_\tau^2}{\alpha_\tau} \mathbf{s}_\theta(\mathbf{x}_\tau, y, \tau)$ and $w_\tau = \frac{\sigma_\tau^2 + \sigma_{\text{data}}^2}{\sigma_\tau^2 \sigma_{\text{data}}^2} \cdot \frac{P_{\text{std}}^{-1}}{\sigma_\tau \sqrt{2\pi}} \exp\left\{-\frac{(\log \sigma_\tau - P_{\text{mean}})^2}{2P_{\text{std}}^2}\right\}$.²

260 This diffusion classifier can be extended to non-uniform priors by incorporating $p(y)$ into the logits
 261 of class y , where $p(y)$ is estimated from the training set (Luo et al., 2024; Wang et al., 2022a), as
 262 detailed in Appendix C.2. As training proceeds, the conditional ELBO converges towards the true
 263 distribution $q_t(\mathbf{x}_t | y)$, thereby yielding increasingly accurate posterior estimates. For convenience,
 264 we denote the class probability of a noisy input \mathbf{x}_t with the diffusion classifier as $f(\mathbf{x}_t)$.

265 To derive the classification loss, we transform the maximization problem of the classification term in
 266 Eq. (7) into the minimization of the negative log-likelihood. We show that the resulting objective,
 267 i.e., $-\sum_Y p_\phi(Y | X, Z) \log p_\theta(Y | X, Z)$, naturally aligns closely with prior work (Lv et al., 2020;
 268 Tarvainen & Valpola, 2017; Liu et al., 2020) and has been shown to be effective in practice.
 269

²As specified in the EDM (Karras et al., 2022), we use $\sigma_{\text{data}} = 0.5$, $P_{\text{mean}} = -1.2$ and $P_{\text{std}} = 1.2$.

270 **Partial-label data.** For partial-label data, the imprecise label Z is given as a candidate set S that is
 271 guaranteed to include the true label. In this case, the posterior distribution $p_\theta(Y|X, S)$ is restricted
 272 to assign non-zero probability only to labels within the candidate set. Accordingly, for each sample
 273 (\mathbf{x}, s) , we compute the classification loss from Eq. (7) as

$$275 \quad \mathcal{L}_{\text{Cls}}^{\text{PL}}(\mathbf{x}) = - \sum_{y \in \mathcal{Y}} p_\phi(y|\mathbf{x}, s) \log p_\theta(y|\mathbf{x}, s) = - \sum_{y \in s} \tilde{f}_\phi^{\text{PL}}(\mathbf{x})_y \log f_\theta(\mathbf{x})_y, \quad (13)$$

277 where $\tilde{f}_\phi^{\text{PL}}(\mathbf{x})$ denotes the normalized probability over s such that $\sum_{y \in s} \tilde{f}_\phi^{\text{PL}}(\mathbf{x})_y = 1$ and $\tilde{f}_\phi^{\text{PL}}(\mathbf{x})_y =$
 278 0 for all $y \notin s$. Eq. (13) can be interpreted as an EMA-stabilized variant of the method called
 279 progressive identification (PRODEN) (Lv et al., 2020), where EMA predictions serve as soft pseudo-
 280 targets.

281 **Supplementary-unlabeled data.** In this scenario, the training set consists of a small portion of
 282 labeled data and a larger number of unlabeled data. This setting can be regarded as a special case of
 283 the partial-label formulation: labeled instances are assigned singleton candidate sets containing the
 284 ground-truth label, while unlabeled instances are associated with the full label space. Accordingly,
 285 the classification loss for each instance is defined as

$$287 \quad \mathcal{L}_{\text{Cls}}^{\text{SU}}(\mathbf{x}) = - \sum_{y \in \mathcal{Y}} p_\phi(y|\mathbf{x}, z) \log p_\theta(y|\mathbf{x}, z) = - \sum_{y \in \mathcal{Y}} \tilde{f}_\phi^{\text{SU}}(\mathbf{x})_y \log f_\theta(\mathbf{x})_y, \quad (14)$$

289 where $\tilde{f}_\phi^{\text{SU}}(\mathbf{x})$ denotes the pseudo-target distribution: for labeled samples, it reduces to a one-hot
 290 vector of the ground-truth label, while for unlabeled samples, it corresponds to the EMA model’s
 291 prediction over the entire label set. This loss can thus be viewed as an EMA-stabilized self-training
 292 objective (Tarvainen & Valpola, 2017), a widely used strategy in semi-supervised learning that
 293 leverages unlabeled data through soft pseudo-labels.

294 **Noisy-label data.** In practice, accurately distinguishing clean labels from noisy ones is often
 295 difficult, making it challenging to retain reliable supervision while applying self-training for label
 296 refinement. To mitigate this, we leverage the memorization effect in noisy-label learning, where
 297 neural networks typically fit clean labels before overfitting to noise (Han et al., 2020). Drawing
 298 inspiration from the noisy-label learning method called early learning regularization (ELR) (Liu
 299 et al., 2020), we propose a simpler yet effective loss function that retains its core idea, defined as

$$301 \quad \mathcal{L}_{\text{Cls}}^{\text{NL}}(\mathbf{x}) = - \sum_{y \in \mathcal{Y}} \text{sg}(\mathbf{r}(\mathbf{x}))_y \log f_\theta(\mathbf{x})_y, \quad \mathbf{r}(\mathbf{x}) = \hat{\mathbf{y}} - \frac{f_\theta(\mathbf{x}) \odot (\langle f_\theta(\mathbf{x}), f_\phi(\mathbf{x}) \rangle \mathbf{1} - f_\phi(\mathbf{x}))}{1 - \langle f_\theta(\mathbf{x}), f_\phi(\mathbf{x}) \rangle}, \quad (15)$$

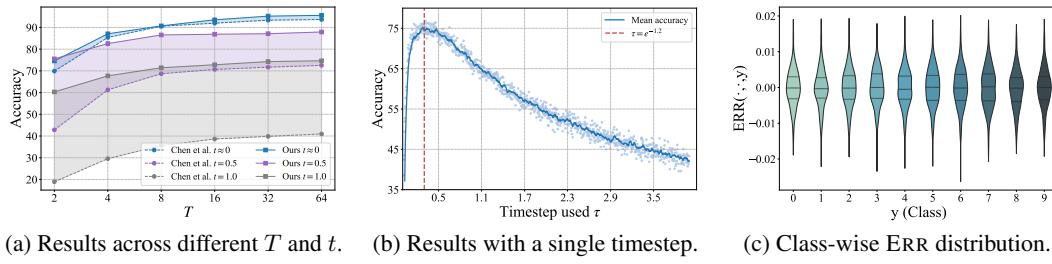
303 where $\hat{\mathbf{y}}$ denotes the one-hot vector of the noisy label \hat{y} , $\text{sg}(\cdot)$ is the stop-gradient operator³, \odot is the
 304 Hadamard product, and $\langle \cdot, \cdot \rangle$ denotes the inner product. This formulation inherits the core principle
 305 of ELR, stabilizing training through soft pseudo-targets derived from the EMA model. It effectively
 306 amplifies the gradient contribution of cleanly labeled samples while suppressing the influence of
 307 mislabeled ones, which we further analyze in detail in Appendix C.1.

309 5 TIME COMPLEXITY REDUCTION

311 The oracle diffusion classifier requires repeated calculations of the conditional ELBO across all classes
 312 to make a prediction, resulting in a substantial computation cost. To address this issue, Chen et al.
 313 (2024b) showed that when estimating ELBO with Monte Carlo sampling, reusing the same \mathbf{x}_τ across
 314 classes and selecting timesteps at uniform intervals is sufficient for effective classification. However,
 315 our experiments reveal that this strategy is empirically suboptimal as illustrated in Figure 1(a). We
 316 identify the core reason to be the model’s varying discriminative ability across different timesteps,
 317 with notable disparities in performance, as shown in Figure 1(b) where the accuracy is evaluated
 318 using only a single timestep. Specifically, when the timestep τ is small, the added noise is negligible,
 319 leading to reconstructions with low label sensitivity. Conversely, when the timestep τ is large, the
 320 input becomes overwhelmed by noise, rendering the predictions highly unreliable.

321 To this end, we aim to identify a compact subset of timesteps that enables efficient ELBO estimation
 322 while maintaining sufficient classification performance. Let $p(\tau)$ be a probability density function

323 ³The stop-gradient operator $\text{sg}(\cdot)$ returns its input but blocks gradient flow, i.e., $\nabla_{\mathbf{x}} \text{sg}(\mathbf{r}(\mathbf{x})) = 0$.



324
325
326
327
328
329
330
331
332
333
334
335
336
337
338
339
340
341
342
343
344
345
346
347
348
349
350
351
352
353
354
355
356
357
358
359
360
361
362
363
364
365
366
367
368
369
370
371
372
373
374
375
376
377

(a) Results across different T and t . (b) Results with a single timestep. (c) Class-wise ERR distribution.

Figure 1: (a): Test accuracy (%) comparison on CIFAR-10 dataset under time complexity reduction technique from Chen et al. (2024c) and ours. (b): Test accuracy (%) on CIFAR-10 dataset evaluated with only a single timestep per class. (c): Violin plot of class-wise $\text{ERR}(\cdot, \cdot, y)$ computed across samples using a fixed subinterval length Δ . Wider regions of the violin indicate higher density.

over the interval $\tau \in (0, +\infty)$, satisfying $\int_0^{+\infty} p(\tau) d\tau = 1$. Our objective is to select a subinterval $\tau \in [l, r]$ such that

$$\underset{0 \leq l \leq r}{\text{minimize}} \left\| \mathbb{E}_{\tau \sim p(\tau | \tau \in [l, r])} [\hbar(\tau, y)] - \mathbb{E}_{\tau \sim p(\tau)} [\hbar(\tau, y)] \right\|_2^2, \quad (16)$$

where $\hbar(\tau, y) = w_\tau \mathbb{E}_{\mathbf{x}_\tau} [\|\mathbf{h}_\theta(\mathbf{x}_\tau, y, \tau) - \mathbf{x}_0\|_2^2]$. Eq. (16) formalizes the goal of finding a representative range where the expected reconstruction error closely matches that of the full distribution. To strike a compromise between signal and noise within the selected subinterval, we propose choosing it around the median of $p(\tau)$, so that signal-dominant early timesteps and noise-dominant later timesteps complement each other. This strategy yields a more stable and representative approximation, especially when $p(\tau)$ is skewed. Therefore, we provide a formal characterization of the subinterval construction for the EDM by the following theorem.

Theorem 2. Consider an EDM where τ is sampled from a log-normal distribution, i.e., $\ln(\tau) \sim \mathcal{N}(\mu, \sigma^2)$, where $\mu \in \mathbb{R}$ and $\sigma^2 > 0$. Given a fixed subinterval length Δ , a sampling range centered around the median of $p(\tau)$ can be constructed by solving the following equation for the left boundary l :

$$l = \text{Solve}_\tau (F(\tau) + F(\tau + \Delta) - 1 = 0), \quad r = l + \Delta,$$

where $\text{Solve}_\tau(\cdot)$ denotes a numerical root-finding algorithm over τ , such as the Brent method (Brent, 2013), and $F(\cdot)$ is the cumulative distribution function of $p(\tau)$.

The proof of Theorem 2 as well as a similar conclusion for denoising diffusion probabilistic model (DDPM) (Ho et al., 2020) are provided in Appendix B.8. Notably, our finding aligns with the effective timestep hypothesis proposed in Li et al. (2023) for the DDPM setting. Furthermore, based on Eq. (16), we can derive a necessary condition that any theoretically optimal subinterval must satisfy, as formalized in the following theorem:

Theorem 3 (Necessary Condition for Optimal Subinterval). Given (l^*, r^*) be an optimal subinterval of the support of $p(\tau)$, a necessary condition for attaining the theoretical minimum of the squared error objective in Eq. (16) is

$$\text{ERR}(l^*, r^*, y) = \mathbb{E}_{\tau \sim p(\tau) | \tau \in [l^*, r^*]} [\hbar(\tau, y)] - \frac{\hbar(l^*, y) + \hbar(r^*, y)}{2} = 0. \quad (17)$$

The proof of Theorem 3 can be found in Appendix B.9. Based on Theorem 3, we empirically present the class-wise distribution of $\text{ERR}(\cdot, \cdot, y)$ across samples in Figure 1(c), where the errors are generally concentrated around zero, supporting the effectiveness of our proposed time complexity reduction strategy. Notably, when the subinterval is reduced to a single sampling point, choosing the median of $p(\tau)$ (i.e., e^{μ}) yields the best classification performance as shown in Figure 1(b). This observation is consistent with our earlier hypothesis regarding the informativeness of the median timestep. In practical posterior inference, we combine timestep subinterval reduction strategy with \mathbf{x}_τ reuse technique (Chen et al., 2024c) to further improve inference efficiency.

Table 1: Generative results on CIFAR-10 and ImageNette under various settings. ‘uncond’ and ‘cond’ indicate unconditional and conditional metrics. **Bold** numbers indicate better performance.

Metric	CIFAR-10	Noisy-label supervision				Partial-label supervision				Suppl-unlabeled supervision				Clean		
		Sym-40%		Asym-40%		Random		Class-50%		Random-1%		Random-10%				
		Vanilla	DMIS	Vanilla	DMIS	Vanilla	DMIS	Vanilla	DMIS	Vanilla	DMIS	Vanilla	DMIS			
CIFAR-10	cond	FID	(↓)	3.33	3.47	3.23	3.10	7.76	2.26	11.75	2.77	3.16	3.12	2.93	2.89	2.05
		IS	(↑)	9.56	9.68	9.02	9.73	9.09	9.80	9.62	9.68	10.03	10.57	9.80	9.83	10.61
		Density	(↑)	101.39	109.75	100.06	109.69	103.21	106.49	108.76	109.06	97.19	108.18	99.96	108.87	112.59
		Coverage	(↑)	81.12	81.21	80.71	81.30	68.45	82.69	64.90	81.52	78.44	81.00	81.85	82.00	83.27
		CW-FID	(↓)	29.84	13.85	14.70	13.24	27.18	10.65	32.44	11.56	16.25	16.12	11.84	11.77	9.83
		CW-Density	(↑)	72.98	107.23	90.85	107.07	102.04	105.75	102.43	108.66	89.99	100.73	96.29	107.94	111.70
		CW-Coverage	(↑)	73.39	80.11	79.63	79.65	65.45	82.09	61.45	81.24	75.03	76.84	80.80	81.12	83.91
ImageNette	cond	FID	(↓)	14.11	13.44	13.93	13.91	79.13	72.62	91.28	79.12	23.88	19.26	14.32	12.84	11.52
		IS	(↑)	12.69	13.21	12.51	13.73	9.19	9.40	9.27	9.11	12.23	13.72	12.80	13.16	13.81
		Density	(↑)	109.31	112.52	111.66	106.78	95.33	99.83	94.29	102.58	115.94	125.68	105.27	109.23	117.23
		Coverage	(↑)	76.62	76.81	78.32	79.81	21.44	32.48	16.69	22.30	53.53	55.39	73.79	75.55	80.12
		CW-FID	(↓)	80.31	60.12	62.26	58.20	157.76	63.58	163.45	67.92	71.66	70.27	49.22	44.31	40.20
		CW-Density	(↑)	73.99	81.12	93.53	94.58	93.38	95.83	91.50	95.21	115.90	118.69	103.41	115.67	120.35
		CW-Coverage	(↑)	67.89	71.94	74.18	75.82	19.76	24.35	15.88	18.93	51.73	52.15	72.61	74.85	78.48

6 EXPERIMENTS

We present experiments on three tasks including image generation, weakly supervised learning, and dataset condensation to demonstrate the utility and versatility of our method. Evaluations are performed on three benchmark datasets widely used for both generation and classification, covering image resolutions from 28×28 (Fashion-MNIST (Xiao et al., 2017)) and 32×32 (CIFAR-10 (Krizhevsky et al., 2009)) to 64×64 (ImageNette (Deng et al., 2009)). As a baseline, we refer to the model trained with the generative objective in Eq. (8) as the *Vanilla* method. The training hyperparameters are kept consistent with those used in the EDM model (Karras et al., 2022).

Dataset construction. For partial-label data, we generate synthetic candidate label sets using both class-dependent (Wen et al., 2021) and random generation models (Feng et al., 2020). In the class-dependent setting, we construct a transition matrix that maps each true label to a set of semantically similar labels, where each similar label is included in the candidate set with probability 50%. In contrast, the random setting assigns each incorrect label an equal probability 50% of being included in the candidate set. For supplementary-unlabeled data, we follow a standard semi-supervised setup by randomly selecting 10% and 1% of the training data classwise as labeled samples, and treating the remaining data as unlabeled. For noisy-label data, we consider both symmetric and asymmetric noise. In the symmetric case, labels are uniformly flipped to any incorrect class, whereas in the asymmetric case, they are flipped to semantically similar classes according to a predefined mapping. In both cases, the corruption probability is referred to as the noise rate, which is set to 40%.

6.1 TASK1: IMAGE GENERATION

Setup. We evaluate the trained CDMs using four unconditional metrics, including Fréchet Inception Distance (FID) (Heusel et al., 2017), Inception Score (IS) (Salimans et al., 2016), Density, and Coverage (Naeem et al., 2020), as well as three conditional metrics, namely CW-FID, CW-Density, and CW-Coverage (Chao et al., 2022). The Class-Wise (CW) metrics are computed per class and then averaged. Detailed descriptions of these metrics are provided in the Appendix E.1.

Results. Table 10 reports the generative performance of the *Vanilla* model and our proposed *DMIS* model on various settings. It can be seen that our model outperforms the baseline across almost all cases with respect to both unconditional and conditional metrics. The performance gap is especially

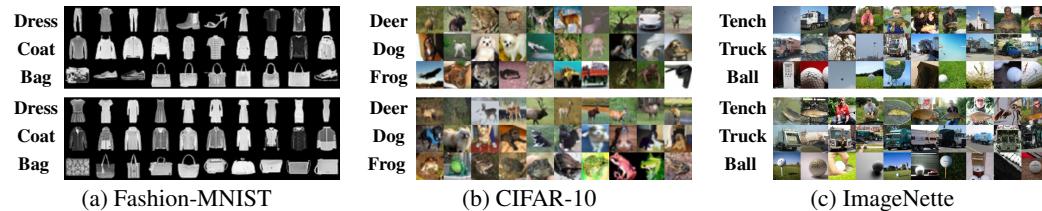


Figure 2: Comparison of conditionally generated images from *Vanilla* (top) and our *DMIS* model (bottom), each trained with 40% symmetric noise on Fashion-MNIST, CIFAR-10, and ImageNette.

432
 433 Table 2: Classification results (test accuracy, %) on Fashion-MNIST, CIFAR-10, and ImageNette
 434 datasets under various types of imprecise supervision (\spadesuit : partial-label, \heartsuit : supplementary-unlabeled,
 435 \clubsuit : noisy-label). **Bold** numbers indicate the best performance.⁵

436 Dataset \spadesuit	437 Type	438 PRODEN	439 IDGP	440 PiCO	441 CRDPLL	442 DIRM	443 Vanilla	444 DMIS ^{CE}	445 DMIS
446 F-MNIST	447 Random	448 93.31 ± 0.07	449 92.26 ± 0.25	450 93.32 ± 0.12	451 94.03 ± 0.14	452 94.11 ± 0.22	453 80.20 ± 1.29	454 84.24 ± 0.37	455 94.27 ± 0.55
	447 Class-50%	448 93.44 ± 0.21	449 93.07 ± 0.16	450 93.32 ± 0.33	451 93.80 ± 0.23	452 93.99 ± 0.24	453 66.03 ± 1.43	454 78.45 ± 0.46	455 94.20 ± 0.15
446 CIFAR-10	447 Random	448 90.02 ± 0.22	449 89.65 ± 0.53	450 86.40 ± 0.89	451 92.74 ± 0.26	452 93.48 ± 0.14	453 60.25 ± 0.17	454 91.47 ± 0.15	455 94.70 ± 0.49
	447 Class-50%	448 90.44 ± 0.44	449 90.83 ± 0.34	450 87.51 ± 0.66	451 92.89 ± 0.27	452 93.22 ± 0.37	453 56.34 ± 0.50	454 90.52 ± 0.35	455 93.53 ± 0.12
446 ImageNette	447 Random	448 84.75 ± 0.13	449 84.07 ± 0.26	450 82.15 ± 0.23	451 84.31 ± 0.25	452 87.90 ± 0.11	453 56.04 ± 0.61	454 84.49 ± 0.05	455 89.31 ± 0.21
	447 Class-50%	448 83.50 ± 0.60	449 82.18 ± 0.13	450 84.41 ± 0.93	451 88.08 ± 0.34	452 87.47 ± 0.17	453 59.47 ± 0.51	454 82.34 ± 0.27	455 88.42 ± 0.43
446 Dataset \heartsuit	447 Type	448 Dash	449 CoMatch	450 FlexMatch	451 SimMatch	452 SoftMatch	453 Vanilla	454 DMIS ^{CE}	455 DMIS
446 F-MNIST	447 Random-1%	448 84.73 ± 0.09	449 85.31 ± 0.29	450 84.43 ± 0.30	451 84.69 ± 0.17	452 84.72 ± 0.23	453 78.37 ± 0.72	454 82.92 ± 0.17	455 85.92 ± 0.13
	447 Random-10%	448 91.16 ± 0.20	449 90.52 ± 0.12	450 90.69 ± 0.03	451 91.18 ± 0.13	452 91.22 ± 0.11	453 90.50 ± 1.00	454 91.07 ± 0.18	455 92.97 ± 0.21
446 CIFAR-10	447 Random-1%	448 70.14 ± 0.69	449 61.45 ± 1.46	450 70.72 ± 0.93	451 73.33 ± 1.02	452 73.74 ± 0.82	453 53.49 ± 0.15	454 75.30 ± 0.17	455 76.40 ± 0.54
	447 Random-10%	448 81.50 ± 0.68	449 77.79 ± 0.53	450 81.35 ± 0.48	451 82.90 ± 0.43	452 88.66 ± 0.60	453 85.13 ± 0.12	454 89.85 ± 0.08	455 92.47 ± 0.39
446 ImageNette	447 Random-1%	448 57.68 ± 2.19	449 63.88 ± 0.78	450 61.39 ± 0.70	451 58.12 ± 2.66	452 58.50 ± 2.31	453 49.55 ± 0.99	454 62.64 ± 0.24	455 68.23 ± 0.19
	447 Random-10%	448 74.66 ± 0.81	449 73.20 ± 0.46	450 73.08 ± 0.13	451 76.12 ± 0.45	452 75.75 ± 0.25	453 74.70 ± 0.53	454 71.39 ± 0.45	455 77.30 ± 0.15
446 Dataset \clubsuit	447 Type	448 CE	449 Mixup	450 Coteaching	451 ELR	452 PENCIL	453 Vanilla	454 DMIS ^{CE}	455 DMIS
446 F-MNIST	447 Sym-40%	448 76.18 ± 0.26	449 92.21 ± 0.03	450 92.17 ± 0.34	451 93.13 ± 0.13	452 90.85 ± 0.58	453 90.11 ± 1.24	454 87.76 ± 0.57	455 93.40 ± 0.40
	447 Asym-40%	448 82.01 ± 0.06	449 92.01 ± 1.02	450 92.78 ± 0.25	451 92.82 ± 0.09	452 91.77 ± 0.69	453 85.41 ± 0.96	454 83.39 ± 0.24	455 93.20 ± 0.30
446 CIFAR-10	447 Sym-40%	448 67.22 ± 0.26	449 84.26 ± 0.64	450 86.54 ± 0.57	451 85.68 ± 0.13	452 85.91 ± 0.26	453 80.22 ± 0.10	454 84.75 ± 0.36	455 88.63 ± 0.12
	447 Asym-40%	448 76.98 ± 0.42	449 83.21 ± 0.85	450 79.38 ± 0.39	451 81.32 ± 0.31	452 84.89 ± 0.49	453 86.31 ± 0.10	454 84.21 ± 0.18	455 88.83 ± 0.33
446 ImageNette	447 Sym-40%	448 58.43 ± 0.77	449 76.65 ± 1.62	450 66.55 ± 1.00	451 84.33 ± 2.86	452 81.94 ± 1.26	453 55.86 ± 1.95	454 80.47 ± 0.56	455 84.12 ± 0.18
	447 Asym-40%	448 71.81 ± 0.38	449 77.16 ± 0.71	450 75.12 ± 0.50	451 73.51 ± 0.31	452 77.20 ± 1.15	453 53.91 ± 1.07	454 77.21 ± 0.19	455 79.30 ± 0.27

451
 452 pronounced under partial-label supervision. These results indicate that *DMIS* not only enhances
 453 the quality of samples but also produces generative distributions that more closely align with the
 454 true data distribution. Furthermore, Figure 2 compares conditionally generated samples from the
 455 *Vanilla* and *DMIS* models across different datasets. Compared to the *Vanilla* model which often
 456 produces samples misaligned with the class, our model produces images of higher visual fidelity and
 457 class-conditional generations that more accurately reflect the intended semantic categories.
 458

459 6.2 TASK2: WEAKLY SUPERVISED LEARNING

460
461 Setup. We evaluate our method under three weakly supervised scenarios. In partial-label learning,
 462 we compare against approaches including *PRODEN* (Lv et al., 2020), *IDGP* (Qiao et al., 2023),
 463 *PiCO* (Wang et al., 2023), *CRDPLL* (Wu et al., 2022) and *DIRK* (Wu et al., 2024). For semi-
 464 supervised learning, we adopt *Dash* (Xu et al., 2021), *CoMatch* (Li et al., 2021a), *FlexMatch* (Zhang
 465 et al., 2021a), *SimMatch* (Zheng et al., 2022) and *SoftMatch* (Chen et al., 2023) as comparison
 466 methods. For noisy-label learning, we compare with *Coteaching* (Han et al., 2018), *ELR* (Liu
 467 et al., 2020), *PENCIL* (Yi & Wu, 2019), as well as standard normal cross-entropy (*CE*) training and
 468 *Mixup* (Zhang et al., 2018). To ensure a fair comparison, the discriminative classifier is implemented
 469 as Wide-ResNet-40-10 with 55.84M parameters, while our generative model contains 55.73M
 470 parameters, and all models are trained from scratch without pre-training.

471
472 Results. The classification results for weakly supervised learning are reported in Table 2. Overall,
 473 our method *DMIS*, evaluated via a diffusion classifier, achieves the best performance, demonstrating
 474 the stronger generalization capability of diffusion models over prior discriminative approaches. Inter-
 475 interestingly, the *Vanilla* method still outperforms several baselines, particularly in the noisy-label setting,
 476 suggesting that the vanilla denoising score matching objective still acts as an implicit regularizer
 477 against label noise. Moreover, compared to standard *CE* training, the regenerate-classification variant
 478 *DMIS*^{CE} improves accuracy by up to 11.58%, 17.53%, and 22.13% on Fashion-MNIST, CIFAR-10,
 479 and ImageNette dataset, respectively, showing that the regenerated dataset effectively mitigates label
 480 imprecision and yields cleaner supervision for downstream discriminative training.

481 6.3 TASK3: NOISY DATASET CONDENSATION

482 While the task of dataset condensation has achieved remarkable progress recently, existing methods
 483 are typically developed under the assumption of clean labels. However, label noise is inevitable and
 484 cannot be fully eliminated in practice. Therefore, exploring how to condense a clean dataset from
 485

⁵*DMIS*^{CE} denotes regenerate-classification results, i.e., we regenerate datasets of the same size under conditional sampling and then train a discriminative model on them using standard *CE* loss.

486
 487 **Table 3: Classification results (test accuracy, %) on noisy-label Fashion-MNIST, CIFAR-10, and**
 488 **ImageNette datasets. ‘IPC’ indicates the number of images per class in the condensed dataset. **Bold****
 489 **numbers indicate the best performance.**

Dataset	Type	IPC	Random	DC	DSA	DM	MTT	RDED	SRE2L	DMIS
F-MNIST	Sym-40%	10	34.42 \pm 0.69	22.85 \pm 1.69	42.07 \pm 2.49	57.06 \pm 1.52	9.03 \pm 3.81	18.57 \pm 1.06	15.80 \pm 0.38	70.18\pm0.37
		50	52.36 \pm 0.60	35.64 \pm 2.26	55.22 \pm 1.51	68.23 \pm 0.47	10.91 \pm 0.82	23.19 \pm 0.74	19.51 \pm 0.96	80.73\pm0.07
		100	55.14 \pm 0.06	30.46 \pm 1.74	41.30 \pm 0.85	73.21 \pm 0.69	13.73 \pm 3.96	25.43 \pm 0.21	19.66 \pm 1.91	84.26\pm0.02
	Asym-40%	10	48.28 \pm 0.34	53.17 \pm 1.59	57.15 \pm 2.37	63.27 \pm 1.60	8.75 \pm 0.82	18.42 \pm 1.62	16.45 \pm 1.96	65.02\pm1.85
		50	69.44 \pm 0.17	49.21 \pm 0.69	77.20 \pm 0.34	76.39 \pm 0.57	8.76 \pm 2.11	22.31 \pm 0.67	27.07 \pm 0.35	79.65\pm0.63
		100	70.80 \pm 0.91	36.95 \pm 0.57	80.24 \pm 0.54	78.43 \pm 0.63	12.59 \pm 1.22	24.03 \pm 0.97	26.52 \pm 1.46	83.22\pm0.33
CIFAR-10	Sym-40%	10	16.30 \pm 0.96	18.11 \pm 1.02	18.06 \pm 1.72	23.71 \pm 0.40	12.06 \pm 0.46	19.85 \pm 0.88	13.12 \pm 1.04	27.83\pm0.98
		50	26.59 \pm 0.70	20.63 \pm 0.22	28.76 \pm 0.57	29.50 \pm 0.56	17.96 \pm 2.10	34.64 \pm 0.58	14.23 \pm 1.67	46.47\pm0.41
		100	31.19 \pm 0.74	19.91 \pm 0.54	29.45 \pm 0.34	32.26 \pm 0.75	18.04 \pm 3.55	44.03 \pm 0.21	14.21 \pm 0.93	56.53\pm0.03
	Asym-40%	10	24.89 \pm 1.65	18.51 \pm 1.35	22.23 \pm 1.80	26.53 \pm 0.07	9.62 \pm 1.45	23.48 \pm 0.65	14.64 \pm 1.03	24.94\pm0.49
		50	40.95 \pm 0.59	25.97 \pm 0.97	40.81 \pm 0.29	43.09 \pm 0.76	16.54 \pm 1.88	39.12 \pm 0.13	16.03 \pm 0.21	47.77\pm0.78
		100	47.49 \pm 0.64	27.76 \pm 0.72	42.96 \pm 0.84	51.61 \pm 0.60	17.67 \pm 2.53	44.45 \pm 0.19	17.55 \pm 0.91	55.89\pm0.39
ImageNette	Sym-40%	10	23.09 \pm 0.19	15.89 \pm 0.73	27.70 \pm 1.25	28.83 \pm 0.73	33.60 \pm 0.53	21.15 \pm 1.05	25.03 \pm 1.17	34.36\pm1.05
		50	33.83 \pm 0.28	24.62 \pm 0.73	32.07 \pm 1.01	42.66 \pm 1.27	38.39 \pm 1.67	35.87 \pm 0.39	35.37 \pm 0.82	44.93\pm0.28
		100	40.04 \pm 0.71	22.81 \pm 1.22	36.05 \pm 1.76	43.25 \pm 2.13	39.61 \pm 1.52	35.87 \pm 0.39	41.74 \pm 1.37	56.23\pm0.84
	Asym-40%	10	26.54 \pm 0.88	19.26 \pm 0.98	30.62 \pm 2.09	33.40 \pm 0.48	33.65 \pm 1.29	26.23 \pm 0.06	25.74 \pm 2.21	37.09\pm0.29
		50	47.91 \pm 0.61	31.68 \pm 2.15	43.41 \pm 1.24	50.97 \pm 1.61	38.71 \pm 1.24	32.75 \pm 0.43	35.29 \pm 0.14	55.20\pm0.46
		100	59.10 \pm 1.41	29.19 \pm 0.21	53.79 \pm 0.84	60.70 \pm 1.88	37.69 \pm 1.29	35.48 \pm 0.22	42.37 \pm 0.34	68.97\pm0.12

501
 502 noisy-label data is natural and meaningful. To the best of our knowledge, this is the first work to
 503 investigate dataset condensation under noisy supervision, which we term *noisy dataset condensation*.
 504

505 **Setup.** During condensation, we employ our trained CDMs to synthesize images according to the spec-
 506 ified IPC. For evaluation, we compare against both hard-label-based methods, including *DC* (Zhao
 507 et al., 2021), *DSA* (Zhao & Bilen, 2021), *DM* (Zhao & Bilen, 2023), and *MTT* (Cazenavette et al.,
 508 2022b), as well as soft-label-based methods, namely *RDED* (Sun et al., 2024) and *SRE2L* (Yin et al.,
 509 2024). Following common protocols (Sun et al., 2024; Yin et al., 2024), we adopt ResNet-18 as the
 510 backbone during condensation and evaluate the condensed datasets on a test set using ResNet-34.

511 **Results.** Table 3 presents the results of noisy dataset condensation, with qualitative visualizations
 512 provided in Appendix E.5. Our method consistently surpasses prior approaches across datasets and
 513 noise types. These results highlight the advantage of generative condensation: rather than memorizing
 514 noisy labels, *DMIS* implicitly denoises them during generation, leading to cleaner condensed datasets.
 515 Notably, unlike the trends observed in clean dataset condensation, distribution-matching methods
 516 (e.g., *DM*) achieve the second-best results in this noisy setting, suggesting that distribution alignment
 517 helps regularize the effect of label noise. Moreover, instance-selection methods generally outperform
 518 synthetic-generation methods (e.g., *Random* vs. *DC/DSA/MTT* and *RDED* vs. *SRE2L*), indicating
 519 that discarding noisy samples during condensation is also an effective strategy to mitigate label noise.
 520 Collectively, these findings not only demonstrate the superiority of our approach but also provide
 521 useful insights for future work on noisy dataset condensation.

522 7 CONCLUSION

523 In this paper, we addressed the challenge of training CDMs under imprecise supervision, a setting
 524 that frequently arises in real-world applications. We introduced a unified framework that formulates
 525 the learning problem as likelihood maximization and decomposes it into generative and classification
 526 components. Based on this formulation, we proposed a weighted denoising score matching objective
 527 that enables label-conditioned learning without clean annotations, and developed an efficient timestep
 528 sampling strategy to reduce the computational cost of posterior inference. Extensive experiments
 529 across image generation, weakly supervised learning, and noisy dataset condensation verified the
 530 effectiveness and versatility of our approach. Beyond establishing strong baselines, our work also
 531 pioneers the study of noisy dataset condensation, opening new opportunities for future exploration in
 532 robust and scalable diffusion modeling under weak supervision.

533 REFERENCES

534 David Berthelot, Nicholas Carlini, Ekin D Cubuk, Alex Kurakin, Kihyuk Sohn, Han Zhang, and
 535 Colin Raffel. Remixmatch: Semi-supervised learning with distribution matching and augmentation
 536 anchoring. In *International Conference on Learning Representations*, 2019a.

- 540 David Berthelot, Nicholas Carlini, Ian Goodfellow, Nicolas Papernot, Avital Oliver, and Colin A
 541 Raffel. Mixmatch: A holistic approach to semi-supervised learning. In *Advances in Neural*
 542 *Information Processing Systems*, pp. 5050–5060, 2019b.
- 543
- 544 Richard P Brent. *Algorithms for minimization without derivatives*. Courier Corporation, 2013.
- 545
- 546 George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A Efros, and Jun-Yan Zhu. Dataset
 547 distillation by matching training trajectories. In *IEEE/CVF Conference on Computer Vision and*
 548 *Pattern Recognition*, pp. 4750–4759, 2022a.
- 549
- 550 George Cazenavette, Tongzhou Wang, Antonio Torralba, Alexei A Efros, and Jun-Yan Zhu. Dataset
 551 distillation by matching training trajectories. In *IEEE/CVF Conference on Computer Vision and*
 552 *Pattern Recognition*, pp. 4750–4759, 2022b.
- 553
- 554 Chen-Hao Chao, Wei-Fang Sun, Bo-Wun Cheng, Yi-Chen Lo, Chia-Che Chang, Yu-Lun Liu, Yu-Lin
 555 Chang, Chia-Ping Chen, and Chun-Yi Lee. Denoising likelihood score matching for conditional
 556 score-based data generation. In *International Conference on Learning Representations*, 2022.
- 557
- 558 Hao Chen, Ran Tao, Yue Fan, Yidong Wang, Jindong Wang, Bernt Schiele, Xing Xie, Bhiksha Raj,
 559 and Marios Savvides. Softmatch: Addressing the quantity-quality tradeoff in semi-supervised
 560 learning. In *International Conference on Learning Representations*, 2023.
- 561
- 562 Hao Chen, Ankit Shah, Jindong Wang, Ran Tao, Yidong Wang, Xiang Li, Xing Xie, Masashi
 563 Sugiyama, Rita Singh, and Bhiksha Raj. Imprecise label learning: A unified framework for learning
 564 with various imprecise label configurations. In *Advances in Neural Information Processing Systems*,
 565 pp. 59621–59654, 2024a.
- 566
- 567 Huanran Chen, Yinpeng Dong, Shitong Shao, Zhongkai Hao, Xiao Yang, Hang Su, and Jun Zhu.
 568 Diffusion models are certifiably robust classifiers. In *Advances in Neural Information Processing*
 569 *Systems*, pp. 50062–50097, 2024b.
- 570
- 571 Huanran Chen, Yinpeng Dong, Zhengyi Wang, Xiao Yang, Chengqi Duan, Hang Su, and Jun Zhu.
 572 Robust classification via a single diffusion model. In *Proceedings of the International Conference*
 573 *on Machine Learning*, pp. 6643–6665, 2024c.
- 574
- 575 Ciprian Corneanu, Raghudeep Gadde, and Aleix M Martinez. Latentpaint: Image inpainting in latent
 576 space with diffusion models. In *Proceedings of the IEEE/CVF winter conference on applications*
 577 *of computer vision*, pp. 4334–4343, 2024.
- 578
- 579 Justin Cui, Ruochen Wang, Si Si, and Cho-Jui Hsieh. Scaling up dataset distillation to imagenet-1k
 580 with constant memory. In *Proceedings of the International Conference on Machine Learning*, pp.
 581 6565–6590, 2023.
- 582
- 583 Jia Deng, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei. Imagenet: A large-scale hier-
 584 archical image database. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*,
 585 pp. 248–255, 2009.
- 586
- 587 Prafulla Dhariwal and Alexander Quinn Nichol. Diffusion models beat gans on image synthesis. In
 588 *Advances in Neural Information Processing Systems*, pp. 8780–8794, 2021.
- 589
- 590 Nicolas Dufour, Victor Besnier, Vicky Kalogeiton, and David Picard. Don’t drop your samples!
 591 coherence-aware training benefits conditional diffusion. In *IEEE/CVF Conference on Computer*
 592 *Vision and Pattern Recognition*, 2024.
- 593
- 594 Bradley Efron. Tweedie’s formula and selection bias. *Journal of the American Statistical Association*,
 595 pp. 1602–1614, 2011.
- 596
- 597 Patrick Esser, Sumith Kulal, Andreas Blattmann, Rahim Entezari, Jonas Müller, Harry Saini, Yam
 598 Levi, Dominik Lorenz, Axel Sauer, Frederic Boesel, et al. Scaling rectified flow transformers
 599 for high-resolution image synthesis. In *Proceedings of the International Conference on Machine*
 600 *Learning*, pp. 12606–12633, 2024.

- 594 Lei Feng, Jiaqi Lv, Bo Han, Miao Xu, Gang Niu, Xin Geng, Bo An, and Masashi Sugiyama. Provably
 595 consistent partial-label learning. In *Advances in Neural Information Processing Systems*, pp.
 596 10948–10960, 2020.
- 597 Bo Han, Quanming Yao, Xingrui Yu, Gang Niu, Miao Xu, Weihua Hu, Ivor Tsang, and Masashi
 598 Sugiyama. Co-teaching: Robust training of deep neural networks with extremely noisy labels. In
 599 *Advances in Neural Information Processing Systems*, pp. 8536–8546, 2018.
- 600 Bo Han, Quanming Yao, Tongliang Liu, Gang Niu, Ivor W Tsang, James T Kwok, and Masashi
 601 Sugiyama. A survey of label-noise representation learning: Past, present and future. *arXiv preprint*
 602 *arXiv:2011.04406*, 2020.
- 603 Chunming He, Chengyu Fang, Yulun Zhang, Longxiang Tang, Jinfu Huang, Kai Li, Zhenhua Guo,
 604 Xiu Li, and Sina Farsiu. Reti-diff: Illumination degradation image restoration with retinex-based
 605 latent diffusion model. In *International Conference on Learning Representations*, 2025.
- 606 Wei He, Kai Han, Ying Nie, Chengcheng Wang, and Yunhe Wang. Species196: A one-million
 607 semi-supervised dataset for fine-grained species recognition. In *Advances in Neural Information
 608 Processing Systems*, pp. 44957–44975, 2023.
- 609 Martin Heusel, Hubert Ramsauer, Thomas Unterthiner, Bernhard Nessler, and Sepp Hochreiter. Gans
 610 trained by a two time-scale update rule converge to a local nash equilibrium. In *Advances in Neural
 611 Information Processing Systems*, pp. 6626–6637, 2017.
- 612 Jonathan Ho and Tim Salimans. Classifier-free diffusion guidance. In *NeurIPS 2021 Workshop on
 613 Deep Generative Models and Downstream Applications*, 2022.
- 614 Jonathan Ho, Ajay Jain, and Pieter Abbeel. Denoising diffusion probabilistic models. In *Advances in
 615 Neural Information Processing Systems*, pp. 6840–6851, 2020.
- 616 Jonathan Ho, Tim Salimans, Alexey Gritsenko, William Chan, Mohammad Norouzi, and David J
 617 Fleet. Video diffusion models. In *Advances in Neural Information Processing Systems*, pp.
 618 8633–8646, 2022.
- 619 Takuhiro Kaneko, Yoshitaka Ushiku, and Tatsuya Harada. Label-noise robust generative adversarial
 620 networks. In *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 2467–2476,
 621 2019.
- 622 Tero Karras, Miika Aittala, Timo Aila, and Samuli Laine. Elucidating the design space of diffusion-
 623 based generative models. In *Advances in Neural Information Processing Systems*, pp. 26565–26577,
 624 2022.
- 625 Jang-Hyun Kim, Jinuk Kim, Seong Joon Oh, Sangdoo Yun, Hwanjun Song, Joonhyun Jeong, Jung-
 626 Woo Ha, and Hyun Oh Song. Dataset condensation via efficient synthetic-data parameterization.
 627 In *Proceedings of the International Conference on Machine Learning*, pp. 11102–11118, 2022.
- 628 Diederik Kingma, Tim Salimans, Ben Poole, and Jonathan Ho. Variational diffusion models. In
 629 *Advances in Neural Information Processing Systems*, pp. 21696–21707, 2021.
- 630 Alex Krizhevsky, Geoffrey Hinton, et al. Learning multiple layers of features from tiny images. 2009.
- 631 Dong-Hyun Lee et al. Pseudo-label: The simple and efficient semi-supervised learning method for
 632 deep neural networks. In *Workshop on challenges in representation learning, ICML*, pp. 896, 2013.
- 633 Alexander C Li, Mihir Prabhudesai, Shivam Duggal, Ellis Brown, and Deepak Pathak. Your diffusion
 634 model is secretly a zero-shot classifier. In *Proceedings of the IEEE/CVF International Conference
 635 on Computer Vision*, pp. 2206–2217, 2023.
- 636 Junnan Li, Caiming Xiong, and Steven CH Hoi. Comatch: Semi-supervised learning with contrastive
 637 graph regularization. In *Proceedings of the IEEE/CVF International Conference on Computer
 638 Vision*, pp. 9475–9484, 2021a.
- 639 Wen Li, Limin Wang, Wei Li, Eirikur Agustsson, and Luc Van Gool. Webvision database: Visual
 640 learning and understanding from web data. *arXiv preprint arXiv:1708.02862*, 2017.

- 648 Xuefeng Li, Tongliang Liu, Bo Han, Gang Niu, and Masashi Sugiyama. Provably end-to-end label-
 649 noise learning without anchor points. In *Proceedings of the International Conference on Machine*
 650 *Learning*, pp. 6403–6413, 2021b.
- 651
- 652 Yangming Li, Max Ruiz Luyten, and Mihaela van der Schaar. Risk-sensitive diffusion: Robustly
 653 optimizing diffusion models with noisy samples. In *International Conference on Learning Representations*, 2024.
- 654
- 655 Sheng Liu, Jonathan Niles-Weed, Narges Razavian, and Carlos Fernandez-Granda. Early-learning reg-
 656 ularization prevents memorization of noisy labels. In *Advances in Neural Information Processing*
 657 *Systems*, pp. 20331–20342, 2020.
- 658
- 659 Noel Loo, Ramin Hasani, Alexander Amini, and Daniela Rus. Efficient dataset distillation using
 660 random feature approximation. In *Advances in Neural Information Processing Systems*, pp.
 661 13877–13891, 2022.
- 662
- 663 Calvin Luo. Understanding diffusion models: A unified perspective. *arXiv preprint arXiv:2208.11970*,
 664 2022.
- 665
- 666 Wenshui Luo, Shuo Chen, Tongliang Liu, Bo Han, Gang Niu, Masashi Sugiyama, Dacheng Tao, and
 667 Chen Gong. Estimating per-class statistics for label noise learning. *IEEE Transactions on Pattern*
 668 *Analysis and Machine Intelligence*, 47(1):305–322, 2024.
- 669
- 670 Jiaqi Lv, Miao Xu, Lei Feng, Gang Niu, Xin Geng, and Masashi Sugiyama. Progressive identification
 671 of true labels for partial-label learning. In *Proceedings of the International Conference on Machine*
 672 *Learning*, pp. 6500–6510, 2020.
- 673
- 674 Eran Malach and Shai Shalev-Shwartz. Decoupling” when to update” from” how to update”. In
 675 *Advances in Neural Information Processing Systems*, pp. 960–970, 2017.
- 676
- 677 Takeru Miyato and Masanori Koyama. cgans with projection discriminator. In *International*
 678 *Conference on Learning Representations*, 2018.
- 679
- 680 Takeru Miyato, Shin-ichi Maeda, Masanori Koyama, and Shin Ishii. Virtual adversarial training: a
 681 regularization method for supervised and semi-supervised learning. *IEEE Transactions on Pattern*
 682 *Analysis and Machine Intelligence*, 41(8):1979–1993, 2018.
- 683
- 684 Byeonghu Na, Yeongmin Kim, HeeSun Bae, Jung Hyun Lee, Se Jung Kwon, Wanmo Kang, and
 685 Il-Chul Moon. Label-noise robust diffusion models. In *International Conference on Learning*
 686 *Representations*, 2024.
- 687
- 688 Muhammad Ferjad Naeem, Seong Joon Oh, Youngjung Uh, Yunjey Choi, and Jaejun Yoo. Reliable
 689 fidelity and diversity metrics for generative models. In *Proceedings of the International Conference*
 690 *on Machine Learning*, pp. 7176–7185, 2020.
- 691
- 692 Zhuoshi Pan, Yuguang Yao, Gaowen Liu, Bingquan Shen, H Vicky Zhao, Ramana Rao Kompella,
 693 and Sijia Liu. From trojan horses to castle walls: Unveiling bilateral backdoor effects in diffusion
 694 models. In *NeurIPS 2023 Workshop on Backdoors in Deep Learning-The Good, the Bad, and the*
 695 *Ugly*, 2023.
- 696
- 697 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
 698 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, et al. Pytorch: An imperative style,
 699 high-performance deep learning library. In *Advances in Neural Information Processing Systems*,
 700 2019.
- 701
- Congyu Qiao, Ning Xu, and Xin Geng. Decompositional generation process for instance-dependent
 702 partial label learning. In *International Conference on Learning Representations*, 2023.
- 703
- Robin Rombach, Andreas Blattmann, Dominik Lorenz, Patrick Esser, and Bjorn Ommer. High-
 704 resolution image synthesis with latent diffusion models. In *IEEE/CVF Conference on Computer*
 705 *Vision and Pattern Recognition*, pp. 23464–23473, 2022.

- 702 Chitwan Saharia, William Chan, Saurabh Saxena, Lala Li, Jay Whang, Emily L Denton, Kamyar
 703 Ghasemipour, Raphael Gontijo Lopes, Burcu Karagol Ayan, Tim Salimans, et al. Photorealistic
 704 text-to-image diffusion models with deep language understanding. pp. 36479–36494, 2022.
- 705
- 706 Ahmad Sajedi, Samir Khaki, Ehsan Amjadian, Lucy Z. Liu, Yuri A. Lawryshyn, and Konstantinos N.
 707 Plataniotis. Datadam: Efficient dataset distillation with attention matching. In *Proceedings of the*
 708 *IEEE/CVF International Conference on Computer Vision*, pp. 17097–17107, 2023.
- 709
- 710 Tim Salimans and Jonathan Ho. Progressive distillation for fast sampling of diffusion models. In
 711 *International Conference on Learning Representations*, 2022.
- 712
- 713 Tim Salimans, Ian Goodfellow, Wojciech Zaremba, Vicki Cheung, Alec Radford, and Xi Chen.
 714 Improved techniques for training gans. In *Advances in Neural Information Processing Systems*, pp.
 715 2226–2234, 2016.
- 716
- 717 Shitong Shao, Zikai Zhou, Huanran Chen, and Zhiqiang Shen. Elucidating the design space of dataset
 718 condensation. In *Advances in Neural Information Processing Systems*, pp. 99161–99201, 2024.
- 719
- 720 Vinay Shukla, Zhe Zeng, Kareem Ahmed, and Guy Van den Broeck. A unified approach to count-
 721 based weakly supervised learning. In *Advances in Neural Information Processing Systems*, pp.
 722 38709–38722, 2023.
- 723
- 724 Kihyuk Sohn, David Berthelot, Nicholas Carlini, Zizhao Zhang, Han Zhang, Colin A Raffel, Ekin Do-
 725 gus Cubuk, Alexey Kurakin, and Chun-Liang Li. Fixmatch: Simplifying semi-supervised learning
 726 with consistency and confidence. In *Advances in Neural Information Processing Systems*, pp.
 727 596–608, 2020.
- 728
- 729 Yang Song and Stefano Ermon. Generative modeling by estimating gradients of the data distribution.
 730 In *Advances in Neural Information Processing Systems*, pp. 11895–11907, 2019.
- 731
- 732 Yang Song and Stefano Ermon. Improved techniques for training score-based generative models. In
 733 *Advances in Neural Information Processing Systems*, pp. 12438–12448, 2020.
- 734
- 735 Yang Song, Jascha Sohl-Dickstein, Diederik P Kingma, Abhishek Kumar, Stefano Ermon, and Ben
 736 Poole. Score-based generative modeling through stochastic differential equations. In *International*
 737 *Conference on Learning Representations*, 2020.
- 738
- 739 Yang Song, Conor Durkan, Iain Murray, and Stefano Ermon. Maximum likelihood training of score-
 740 based diffusion models. In *Advances in Neural Information Processing Systems*, pp. 1415–1428,
 741 2021.
- 742
- 743 Peng Sun, Bei Shi, Daiwei Yu, and Tao Lin. On the diversity and realism of distilled dataset: An
 744 efficient dataset distillation paradigm. In *IEEE/CVF Conference on Computer Vision and Pattern*
 745 *Recognition*, pp. 9390–9399, 2024.
- 746
- 747 Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. Rethinking
 748 the inception architecture for computer vision. In *IEEE/CVF Conference on Computer Vision and*
 749 *Pattern Recognition*, pp. 2818–2826, 2016.
- 750
- 751 Hiroshi Takahashi, Tomoharu Iwata, Atsutoshi Kumagai, Yuuki Yamanaka, and Tomoya Yamashita.
 752 Positive-unlabeled diffusion models for preventing sensitive data generation. In *International*
 753 *Conference on Learning Representations*, 2025.
- 754
- 755 Antti Tarvainen and Harri Valpola. Mean teachers are better role models: Weight-averaged consis-
 756 tency targets improve semi-supervised deep learning results. In *Advances in Neural Information*
 757 *Processing Systems*, pp. 1195–1204, 2017.
- 758
- 759 Yingjie Tian, Xiaotong Yu, and Saiji Fu. Partial label learning: Taxonomy, analysis and outlook.
 760 *Neural Networks*, 161:708–734, 2023.
- 761
- Pascal Vincent. A connection between score matching and denoising autoencoders. *Neural computa-
 762 tion*, 23(7):1661–1674, 2011.

- 756 Haobo Wang, Mingxuan Xia, Yixuan Li, Yuren Mao, Lei Feng, Gang Chen, and Junbo Zhao. Solar:
 757 Sinkhorn label refinery for imbalanced partial-label learning. In *Advances in Neural Information*
 758 *Processing Systems*, pp. 8104–8117, 2022a.
- 759
- 760 Haobo Wang, Ruixuan Xiao, Yixuan Li, Lei Feng, Gang Niu, Gang Chen, and Junbo Zhao. Pico+:
 761 Contrastive label disambiguation for robust partial label learning. *IEEE Transactions on Pattern*
 762 *Analysis and Machine Intelligence*, pp. 3183–3198, 2023.
- 763 Hsiu-Hsuan Wang, Tan-Ha Mai, Nai-Xuan Ye, Wei-I Lin, and Hsuan-Tien Lin. Climage: Human-
 764 annotated datasets for complementary-label learning. *Transactions on Machine Learning Research*,
 765 2025, 2025a.
- 766
- 767 Kai Wang, Bo Zhao, Xiangyu Peng, Zheng Zhu, Shuo Yang, Shuo Wang, Guan Huang, Hakan
 768 Bilen, Xinchao Wang, and Yang You. Cafe: Learning to condense dataset by aligning features. In
 769 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 12196–12205, 2022b.
- 770 Tongzhou Wang, Jun-Yan Zhu, Antonio Torralba, and Alexei A Efros. Dataset distillation. *arXiv*
 771 *preprint arXiv:1811.10959*, 2018.
- 772
- 773 Wei Wang, Dong-Dong Wu, Jindong Wang, Gang Niu, Min-Ling Zhang, and Masashi Sugiyama.
 774 Realistic evaluation of deep partial-label learning algorithms. In *International Conference on*
 775 *Learning Representations*, 2025b.
- 776 Yidong Wang, Hao Chen, Yue Fan, Wang Sun, Ran Tao, Wenxin Hou, Renjie Wang, Linyi Yang, Zhi
 777 Zhou, Lan-Zhe Guo, et al. Usb: A unified semi-supervised learning benchmark for classification.
 778 In *Advances in Neural Information Processing Systems*, pp. 3938–3961, 2022c.
- 779
- 780 Yidong Wang, Hao Chen, Qiang Heng, Wenxin Hou, Yue Fan, Zhen Wu, Jindong Wang, Marios
 781 Savvides, Takahiro Shinozaki, Bhiksha Raj, et al. Freematch: Self-adaptive thresholding for
 782 semi-supervised learning. In *International Conference on Learning Representations*, 2022d.
- 783 Yisen Wang, Xingjun Ma, Zaiyi Chen, Yuan Luo, Jinfeng Yi, and James Bailey. Symmetric cross
 784 entropy for robust learning with noisy labels. In *Proceedings of the IEEE/CVF International*
 785 *Conference on Computer Vision*, pp. 322–330, 2019.
- 786
- 787 Hongxin Wei, Lei Feng, Xiangyu Chen, and Bo An. Combating noisy labels by agreement: A joint
 788 training method with co-regularization. In *IEEE/CVF Conference on Computer Vision and Pattern*
 789 *Recognition*, pp. 13726–13735, 2020.
- 790
- 791 Jiaheng Wei, Zhaowei Zhu, Hao Cheng, Tongliang Liu, Gang Niu, and Yang Liu. Learning with
 792 noisy labels revisited: A study using real-world human annotations. In *International Conference*
 793 *on Learning Representations*, 2021.
- 794
- 795 Jiaheng Wei, Zhaowei Zhu, Hao Cheng, Tongliang Liu, Gang Niu, and Yang Liu. Learning with
 796 noisy labels revisited: A study using real-world human annotations. In *International Conference*
 797 *on Learning Representations*, 2022.
- 798
- 799 Zixi Wei, Lei Feng, Bo Han, Tongliang Liu, Gang Niu, Xiaofeng Zhu, and Heng Tao Shen. A
 800 universal unbiased method for classification from aggregate observations. In *Proceedings of the*
 801 *International Conference on Machine Learning*, pp. 36804–36820, 2023.
- 802
- 803 Hongwei Wen, Jingyi Cui, Hanyuan Hang, Jiabin Liu, Yisen Wang, and Zhouchen Lin. Leveraged
 804 weighted loss for partial label learning. In *Proceedings of the International Conference on Machine*
 805 *Learning*, pp. 11091–11100, 2021.
- 806
- 807 Dong-Dong Wu, Deng-Bao Wang, and Min-Ling Zhang. Revisiting consistency regularization for
 808 deep partial label learning. In *Proceedings of the International Conference on Machine Learning*,
 809 pp. 24212–24225, 2022.
- 810
- 811 Dong-Dong Wu, Deng-Bao Wang, and Min-Ling Zhang. Distilling reliable knowledge for instance-
 812 dependent partial label learning. In *Proceedings of the AAAI Conference on Artificial Intelligence*,
 813 pp. 15888–15896, 2024.

- 810 Shiyu Xia, Jiaqi Lv, Ning Xu, and Xin Geng. Ambiguity-induced contrastive learning for instance-
 811 dependent partial label learning. In *Proceedings of the International Joint Conference on Artificial*
 812 *Intelligence*, pp. 3615–3621, 2022.
- 813 Han Xiao, Kashif Rasul, and Roland Vollgraf. Fashion-mnist: a novel image dataset for benchmarking
 814 machine learning algorithms. *arXiv preprint arXiv:1708.07747*, 2017.
- 815 Enze Xie, Junsong Chen, Junyu Chen, Han Cai, Haotian Tang, Yujun Lin, Zhekai Zhang, Muyang Li,
 816 Ligeng Zhu, Yao Lu, et al. Sana: Efficient high-resolution image synthesis with linear diffusion
 817 transformers. In *International Conference on Learning Representations*, 2025.
- 818 Qizhe Xie, Zihang Dai, Eduard Hovy, Thang Luong, and Quoc Le. Unsupervised data augmentation
 819 for consistency training. In *Advances in Neural Information Processing Systems*, pp. 6256–6268,
 820 2020.
- 821 Zheng Xie, Yu Liu, Hao-Yuan He, Ming Li, and Zhi-Hua Zhou. Weakly supervised auc optimization:
 822 A unified partial auc approach. *IEEE Transactions on Pattern Analysis and Machine Intelligence*,
 823 46(7):4780–4795, 2024.
- 824 Ning Xu, Biao Liu, Jiaqi Lv, Congyu Qiao, and Xin Geng. Progressive purification for instance-
 825 dependent partial label learning. In *Proceedings of the International Conference on Machine*
 826 *Learning*, pp. 38551–38565, 2023.
- 827 Yi Xu, Lei Shang, Jinxing Ye, Qi Qian, Yu-Feng Li, Baigui Sun, Hao Li, and Rong Jin. Dash: Semi-
 828 supervised learning with dynamic thresholding. In *Proceedings of the International Conference on*
 829 *Machine Learning*, pp. 11525–11536, 2021.
- 830 Eric Xue, Yijiang Li, Haoyang Liu, Peiran Wang, Yifan Shen, and Haohan Wang. Towards adversari-
 831 ally robust dataset distillation by curvature regularization. In *Proceedings of the AAAI Conference*
 832 *on Artificial Intelligence*, pp. 9041–9049, 2025.
- 833 Muqiao Yang, Chunlei Zhang, Yong Xu, Zhongweiyang Xu, Heming Wang, Bhiksha Raj, and Dong
 834 Yu. Usee: Unified speech enhancement and editing with conditional diffusion models. In *IEEE*
 835 *International Conference on Acoustics, Speech and Signal Processing*, pp. 7125–7129, 2024.
- 836 Xiangli Yang, Zixing Song, Irwin King, and Zenglin Xu. A survey on deep semi-supervised learning.
 837 *IEEE transactions on knowledge and data engineering*, 35(9):8934–8954, 2022.
- 838 Yu Yao, Tongliang Liu, Bo Han, Mingming Gong, Jiankang Deng, Gang Niu, and Masashi Sugiyama.
 839 Dual t: Reducing estimation error for transition matrix in label-noise learning. pp. 7260–7271,
 840 2020.
- 841 Kun Yi and Jianxin Wu. Probabilistic end-to-end noise correction for learning with noisy labels. In
 842 *IEEE/CVF Conference on Computer Vision and Pattern Recognition*, pp. 7017–7025, 2019.
- 843 Zeyuan Yin and Zhiqiang Shen. Dataset distillation via curriculum data synthesis in large data era.
 844 *Transactions on Machine Learning Research*, 2024.
- 845 Zeyuan Yin, Eric Xing, and Zhiqiang Shen. Squeeze, recover and relabel: Dataset condensation at
 846 imangenet scale from a new perspective. In *Advances in Neural Information Processing Systems*,
 847 pp. 73582–73603, 2024.
- 848 Xingrui Yu, Bo Han, Jiangchao Yao, Gang Niu, Ivor Tsang, and Masashi Sugiyama. How does
 849 disagreement help generalization against label corruption? In *Proceedings of the International*
 850 *Conference on Machine Learning*, pp. 7164–7173, 2019.
- 851 Bowen Zhang, Yidong Wang, Wenxin Hou, Hao Wu, Jindong Wang, Manabu Okumura, and Takahiro
 852 Shinozaki. Flexmatch: Boosting semi-supervised learning with curriculum pseudo labeling. In
 853 *Advances in Neural Information Processing Systems*, pp. 18408–18419, 2021a.
- 854 Fei Zhang, Lei Feng, Bo Han, Tongliang Liu, Gang Niu, Tao Qin, and Masashi Sugiyama. Ex-
 855 ploring class activation value for partial-label learning. In *International Conference on Learning*
 856 *Representations*, 2021b.

- 864 Hongyi Zhang, Moustapha Cisse, Yann N Dauphin, and David Lopez-Paz. mixup: Beyond empirical
865 risk minimization. In *International Conference on Learning Representations*, 2018.
866
- 867 Yivan Zhang, Nontawat Charoenphakdee, Zhenguo Wu, and Masashi Sugiyama. Learning from
868 aggregate observations. In *Advances in Neural Information Processing Systems*, pp. 7993–8005,
869 2020.
- 870 Zhilu Zhang and Mert Sabuncu. Generalized cross entropy loss for training deep neural networks
871 with noisy labels. In *Advances in Neural Information Processing Systems*, pp. 8792–8802, 2018.
872
- 873 Bo Zhao and Hakan Bilen. Dataset condensation with differentiable siamese augmentation. In
874 *Proceedings of the International Conference on Machine Learning*, pp. 12674–12685, 2021.
- 875 Bo Zhao and Hakan Bilen. Dataset condensation with distribution matching. In *Proceedings of the*
876 *IEEE/CVF winter conference on applications of computer vision*, pp. 6514–6523, 2023.
- 877 Bo Zhao, Konda Reddy Mopuri, and Hakan Bilen. Dataset condensation with gradient matching. In
878 *International Conference on Learning Representations*, 2021.
- 879
- 880 Chen Zhao, Weiling Cai, Chenyu Dong, and Chengwei Hu. Wavelet-based fourier information
881 interaction with frequency diffusion adjustment for underwater image restoration. In *IEEE/CVF*
882 *Conference on Computer Vision and Pattern Recognition*, pp. 8281–8291, 2024.
- 883
- 884 Mingkai Zheng, Shan You, Lang Huang, Fei Wang, Chen Qian, and Chang Xu. Simmatch: Semi-
885 supervised learning with similarity matching. In *IEEE/CVF Conference on Computer Vision and*
886 *Pattern Recognition*, pp. 14471–14481, 2022.
- 887
- 888
- 889
- 890
- 891
- 892
- 893
- 894
- 895
- 896
- 897
- 898
- 899
- 900
- 901
- 902
- 903
- 904
- 905
- 906
- 907
- 908
- 909
- 910
- 911
- 912
- 913
- 914
- 915
- 916
- 917

918	CONTENTS	
919		
920		
921	1 Introduction	1
922		
923	2 Related Work	2
924	2.1 Robust Diffusion Models	2
925	2.2 Imprecise Label Learning	2
926	2.3 Dataset Condensation	2
927		
928		
929	3 Background	3
930		
931		
932	4 Methodology	4
933	4.1 Unified Learning Objective	4
934	4.2 Generative Objective: Modeling the Imprecise Data Distribution	4
935	4.3 Classification Objective: Inferring Labels from Imprecise Signals	5
936		
937		
938	5 Time Complexity Reduction	6
939		
940		
941	6 Experiments	8
942	6.1 Task1: Image Generation	8
943	6.2 Task2: Weakly Supervised Learning	9
944	6.3 Task3: Noisy Dataset Condensation	9
945		
946		
947	7 Conclusion	10
948		
949	A Notation and Definitions	20
950		
951	B Proof	20
952	B.1 Connections among Different Diffusion Models.	20
953	B.2 Derivation of Eq. (5)	21
954	B.3 Derivation of Varitional Lower Bound Eq. (7)	23
955		
956	B.4 Derivation of Conditional ELBO in Eq. (8)	24
957		
958	B.5 Derivation of Remark 1	24
959		
960	B.6 Proof of Theorem 1	26
961		
962	B.7 Proof of Proposition 1	27
963		
964	B.8 Proof of Theorem 2	27
965		
966	B.9 Proof of Theorem 3	28
967		
968	C Discussion	29
969	C.1 Analysis of Early-Learning Regularization in Eq. (15)	29
970	C.2 Class-Prior Estimation in Imprecise-Label Datasets	30
971		
	D Implementation Details	31

972	E Experiments	32
973	E.1 Evaluation Metrics	32
974	E.2 Full Results in Weakly Supervised Learning	32
975	E.3 Integration with Existing Imprecise-Label Correctors	34
976	E.4 Comparison of Accuracy Curves between <i>DMIS</i> and <i>Vanilla</i>	34
977	E.5 Visualization of Noisy Condensed Datasets	34
978	E.6 Additional Results on Dataset Condensation under Different Forms of Imprecise	
979	Supervision	35
980		
981		
982		
983		
984	F Additional Experiments Results	37
985	F.1 The Full Results of Table 1.	37
986	F.2 Comparison against Other Noise-robust Diffusion Methods.	37
987	F.3 Top- k truncation for large label spaces	38
988	F.4 Experiments beyond synthetic class-conditional noise	38
989		
990		
991		
992		
993		
994		
995		
996		
997		
998		
999		
1000		
1001		
1002		
1003		
1004		
1005		
1006		
1007		
1008		
1009		
1010		
1011		
1012		
1013		
1014		
1015		
1016		
1017		
1018		
1019		
1020		
1021		
1022		
1023		
1024		
1025		

1026 **The Use of Large Language Models (LLMs).** LLMs were only used for language polishing and
 1027 proofreading. No part of the technical content, experiments, or analysis was generated by LLMs.
 1028

1029 **A NOTATION AND DEFINITIONS**

1030 We present the notation table for each symbol used in this paper in Table 4.
 1031

1032 **Table 4: List of common mathematical symbols used in this paper.**

1035 Symbol	1036 Definition
1037 \mathbf{x}	A sample of training data
1038 z	Imprecise label associated with a sample
1039 s	Candidate label set for a sample
1040 y	Class index label
1041 c	Total number of classes
1042 \mathcal{X}	Input space from which \mathbf{x} is drawn
1043 \mathcal{Y}	Label space from which y is drawn
1044 X	Random variable for training instances
1045 Y	Random variable for true labels
1046 Z	Random variable for imprecise labels
1047 S	Random variable for partial labels
1048 \hat{Y}	Random variable for noisy labels
1049 X^l	Set of labeled data instances
1050 X^u	Set of unlabeled data instances
1051 Y^l	Set of labels corresponding to X^l
1052 \emptyset	Empty label set
1053 θ	Parameters of the diffusion model to be optimized
1054 ϕ	Exponential moving average of θ over training iteration
1055 $\mathbf{0}$	Zero vector
1056 \mathbf{I}	Identity matrix
1057 \mathbf{x}_t	Noisy version of the sample at timestep t
1058 τ	Continuous timestep variable
1059 α_t	Scaling factor at timestep t
1060 σ_t	Noise scale at timestep t
1061 l	Left boundary of a subsampled timestep interval
1062 r	Right boundary of a subsampled timestep interval
1063 Δ	Length of a subsampled timestep interval
1064 $q(\cdot)$	Real Data distribution
1065 $q(\cdot \cdot)$	Real conditional data distribution
1066 $p(\cdot)$	Marginal probability distribution
1067 $p(\cdot \cdot)$	Model-inferred conditional distribution
1068 $F(\cdot)$	Cumulative distribution function of $p(\cdot)$
1069 $f(\cdot)$	Diffusion classifier
1070 $\mathbf{s}(\cdot, \cdot)$	Time-conditioned score prediction network
1071 $\mathcal{N}(\cdot, \cdot)$	Gaussian distribution

1072 **B PROOF**

1073 **B.1 CONNECTIONS AMONG DIFFERENT DIFFUSION MODELS.**

1074 The diffusion model we define in this paper can be reformulated to align with other common diffusion
 1075 frameworks, such as DDPM (Ho et al., 2020), SMLD (Song & Ermon, 2019), VE-SDE (Song et al.,
 1076 2020) and VP-SDE (Song et al., 2020), as well as with approaches like x-prediction (Ho et al., 2020),
 1077 v-prediction (Salimans & Ho, 2022), and ϵ -prediction (Ho et al., 2020). This demonstrates that
 1078 our formulation is compatible with diverse diffusion paradigms while facilitating unified theoretical
 1079 analysis. To better demonstrate this transformation, we present the following pseudocodes.

1080 **Algorithm 1** Our models to EDM
1081 **Require:** A score network \mathbf{s}_θ , a noisy input \mathbf{x}_t , noise level t , linear schedule $\{\alpha_i\}_{i=1}^T$ and $\{\sigma_i\}_{i=1}^T$
1082 1: Calculate the denoised image \mathbf{x}_0 using \mathbf{s}_θ : $\mathbf{x}_0 = (\mathbf{x}_t + \sigma_t^2 \mathbf{s}_\theta(\mathbf{x}_t/\alpha_t, \sigma_t/\alpha_t))/\alpha_t$
1083 2: **if** performing \mathbf{x}_0 -prediction **then**
1084 3: **return** \mathbf{x}_0 .
1085 4: **end if**
1086 5: Calculate the noise component ϵ : $\epsilon = \frac{\mathbf{x}_t - \alpha_t \mathbf{x}_0}{\sigma_t}$
1087 6: **if** performing ϵ -prediction **then**
1088 7: **return** ϵ .
1089 8: **end if**
1090 9: Calculate the noise component v : $v = \alpha_t \epsilon - \sigma_t \mathbf{x}_0$
1091 10: **if** performing v -prediction **then**
1092 11: **return** v .
1093 12: **end if**

1094
1095 **DDPM.** DDPM define a sequence $\{\beta_t\}_{t=0}^T$ and $\mathbf{x}_t = \sqrt{\prod_{i=0}^t (1 - \beta_i)} \mathbf{x}_0 + \sqrt{1 - \prod_{i=0}^t (1 - \beta_i)} \epsilon$,
1096 which can be seen as a special case of Eq. (1) where we can set $\alpha_t = \sqrt{\prod_{i=0}^t (1 - \beta_i)}$ and $\sigma_t = \sqrt{1 - \prod_{i=0}^t (1 - \beta_i)}$.
1097
1098
1099

1100 **SMLD.** SMLD defines a noise schedule $\sigma(t)_{t=0}^T$ and $\mathbf{x}_t = \mathbf{x}_0 + \sigma(t) \epsilon$, with $\sigma(1) < \sigma(2) < \dots < \sigma(T)$. In this setup, Eq. (1) reduces to $\alpha_t = 1$, $\sigma_t = \sigma(t)$.
1101
1102

1103 **VP-SDE.** VP-SDE is the continuous case of DDPM, which define a stochastic differential equation
1104 (SDE) as

1105
1106
$$dX_t = -\frac{1}{2} \beta(t) X_t dt + \sqrt{\beta(t)} dW_t, \quad t \in [0, 1],$$

1107
1108 where $\beta(t) = \beta_{t:T} \cdot T$. In this setup, $\alpha_t = \sqrt{\exp\left(-\int_0^t \beta(s) ds\right)}$, $\sigma_t = 1 - \exp\left(-\int_0^t \beta(s) ds\right)$.
1109
1110

1111 **VE-SDE.** VE-SDE is the continuous case of SMLD, whose forward process of VE-SDE is defined as

1112
1113
$$dX_t = \sqrt{\frac{d\sigma(t)^2}{dt}} dW_t.$$

1114
1115 In this setup, $\alpha_t = 1$ and $\sigma_t = \sqrt{\sigma^2(t) - \sigma^2(0)}$.
1116

1117 While the models above each define their own specific frameworks for the diffusion process,
1118 EDM (Karras et al., 2022) proposes a unified structure and optimizes the parameters choice within
1119 the diffusion process, making it both robust and adaptable. Therefore, for our implementation, we
1120 adopt EDM as the foundational diffusion model. In EDM, the scaling and noise schedules are a
1121 special case of VE-SDE, where the variance of the noise is given by $\sigma(t) = t$. Accordingly, we use
1122 $\mathbf{s}_\theta(\mathbf{x}/\alpha_t, \sigma_t/\alpha_t)$ to obtain the predicted score, as shown in Algorithm 1.

1123 B.2 DERIVATION OF EQ. (5)

1124 Maximizing the variational lower bound, or equivalently evidence lower bound (ELBO), to optimize
1125 the diffusion model is a common approach. To avoid redundant proofs, we directly use the conclusion
1126 from Eq. (58) in Luo (2022) as below:

1127
1128
$$\log p_\theta(\mathbf{x}) \geq \mathbb{E}_q[-D_{\text{KL}}(q(\mathbf{x}_T|\mathbf{x}_0) \| p(\mathbf{x}_T)) + \log p_\theta(\mathbf{x}_0|\mathbf{x}_1) - \sum_{t>1} D_{\text{KL}}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) \| p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t))]$$

1129

1130 Although each KL divergence term $D_{\text{KL}}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) \| p_\theta(\mathbf{x}_{t-1}|\mathbf{x}_t))$ is difficult to minimize for
1131 arbitrary posteriors, we can leverage the Gaussian transition assumption to make optimization
1132 tractable. By Bayes rule, we have:

1133
$$q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) = \frac{q(\mathbf{x}_t|\mathbf{x}_{t-1}, \mathbf{x}_0)q(\mathbf{x}_{t-1}|\mathbf{x}_0)}{q(\mathbf{x}_t|\mathbf{x}_0)}$$

1134 As we already know that $q(\mathbf{x}_t | \mathbf{x}_0)$ and $q(\mathbf{x}_{t-1} | \mathbf{x}_0)$ from Eq. (1), $q(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{x}_0)$ can be derived from
 1135 its equivalent form $q(\mathbf{x}_t | \mathbf{x}_{t-1})$ as follows:

$$\begin{aligned}
 1137 \quad \mathbf{x}_t &= \alpha_t \mathbf{x}_0 + \sigma_t \boldsymbol{\epsilon}_0 \\
 1138 &= \alpha_t \left(\frac{\mathbf{x}_{t-1} - \sigma_{t-1} \boldsymbol{\epsilon}_0^*}{\alpha_{t-1}} \right) + \sigma_t \boldsymbol{\epsilon}_0 \\
 1139 &= \frac{\alpha_t}{\alpha_{t-1}} \mathbf{x}_{t-1} + \sigma_t \boldsymbol{\epsilon}_0 - \frac{\alpha_t}{\alpha_{t-1}} \sigma_{t-1} \boldsymbol{\epsilon}_0^* \\
 1140 &= \frac{\alpha_t}{\alpha_{t-1}} \mathbf{x}_{t-1} + \sqrt{\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2} \boldsymbol{\epsilon}_{t-1} \\
 1141 &= \mathcal{N}(\mathbf{x}_t; \frac{\alpha_t}{\alpha_{t-1}} \mathbf{x}_{t-1}, \sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2 \mathbf{I}) \\
 1142 & \\
 1143 & \\
 1144 & \\
 1145 & \\
 1146 & \\
 1147
 \end{aligned}$$

1148 Now, knowing the forms of $q(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{x}_0)$, we can proceed to calculate the form of $q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0)$
 1149 by substituting into the Bayes rule expansion:

$$\begin{aligned}
 1150 \quad q(\mathbf{x}_{t-1} | \mathbf{x}_t, \mathbf{x}_0) &= \frac{q(\mathbf{x}_t | \mathbf{x}_{t-1}, \mathbf{x}_0) q(\mathbf{x}_{t-1} | \mathbf{x}_0)}{q(\mathbf{x}_t | \mathbf{x}_0)} \\
 1151 &= \frac{\mathcal{N}(\mathbf{x}_t; \frac{\alpha_t}{\alpha_{t-1}} \mathbf{x}_{t-1}, \sqrt{\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2} \mathbf{I}) \mathcal{N}(\mathbf{x}_{t-1}; \alpha_{t-1} \mathbf{x}_0, \sigma_{t-1}^2 \mathbf{I})}{\mathcal{N}(\mathbf{x}_t; \alpha_t \mathbf{x}_0, \sigma_t^2 \mathbf{I})} \\
 1152 & \\
 1153 & \\
 1154 & \\
 1155 & \\
 1156 & \propto \exp \left\{ - \left[\frac{(\mathbf{x}_t - \frac{\alpha_t}{\alpha_{t-1}} \mathbf{x}_{t-1})^2}{2(\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2)} + \frac{(\mathbf{x}_{t-1} - \alpha_{t-1} \mathbf{x}_0)^2}{2\sigma_{t-1}^2} - \frac{(\mathbf{x}_t - \alpha_t \mathbf{x}_0)^2}{2\sigma_t^2} \right] \right\} \\
 1157 & \\
 1158 & \\
 1159 & \\
 1160 & = \exp \left\{ - \frac{1}{2} \left[\frac{-2 \frac{\alpha_t}{\alpha_{t-1}} \mathbf{x}_t \mathbf{x}_{t-1} + (\frac{\alpha_t}{\alpha_{t-1}})^2 \mathbf{x}_{t-1}^2}{\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2} + \frac{\mathbf{x}_{t-1}^2 - 2\alpha_{t-1} \mathbf{x}_{t-1} \mathbf{x}_0}{\sigma_{t-1}^2} + C(\mathbf{x}_t, \mathbf{x}_0) \right] \right\} \\
 1161 & \\
 1162 & \\
 1163 & \propto \exp \left\{ - \frac{1}{2} \left[\left(\frac{(\frac{\alpha_t}{\alpha_{t-1}})^2}{\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2} + \frac{1}{\sigma_{t-1}^2} \right) \mathbf{x}_{t-1}^2 - 2 \left(\frac{\frac{\alpha_t}{\alpha_{t-1}} \mathbf{x}_t}{\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2} + \frac{\alpha_{t-1} \mathbf{x}_0}{\sigma_{t-1}^2} \right) \mathbf{x}_{t-1} \right] \right\} \\
 1164 & \\
 1165 & \\
 1166 & \\
 1167 & = \exp \left\{ - \frac{1}{2} \left[\frac{\sigma_{t-1}^2 (\frac{\alpha_t}{\alpha_{t-1}})^2 + (\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2) \mathbf{x}_{t-1}^2}{(\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2) \sigma_{t-1}^2} - 2 \left(\frac{\frac{\alpha_t}{\alpha_{t-1}} \mathbf{x}_t}{\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2} + \frac{\alpha_{t-1} \mathbf{x}_0}{\sigma_{t-1}^2} \right) \mathbf{x}_{t-1} \right] \right\} \\
 1168 & \\
 1169 & \\
 1170 & = \exp \left\{ - \frac{1}{2} \left[\frac{\sigma_t^2}{(\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2) \sigma_{t-1}^2} \mathbf{x}_{t-1}^2 - 2 \left(\frac{\frac{\alpha_t}{\alpha_{t-1}} \mathbf{x}_t}{\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2} + \frac{\alpha_{t-1} \mathbf{x}_0}{\sigma_{t-1}^2} \right) \mathbf{x}_{t-1} \right] \right\} \\
 1171 & \\
 1172 & \\
 1173 & \\
 1174 & = \exp \left\{ - \frac{1}{2} \left(\frac{\sigma_t^2}{(\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2) \sigma_{t-1}^2} \right) \left[\mathbf{x}_{t-1}^2 - 2 \left(\frac{\frac{\alpha_t}{\alpha_{t-1}} \mathbf{x}_t}{\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2} + \frac{\alpha_{t-1} \mathbf{x}_0}{\sigma_{t-1}^2} \right) \mathbf{x}_{t-1} \right] \right\} \\
 1175 & \\
 1176 & \\
 1177 & \\
 1178 & \\
 1179 & \\
 1180 & = \exp \left\{ - \frac{1}{2} \left(\frac{1}{\frac{(\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2) \sigma_{t-1}^2}{\sigma_t^2}} \right) \left[\mathbf{x}_{t-1}^2 - 2 \left(\frac{\frac{\alpha_t}{\alpha_{t-1}} \mathbf{x}_t \sigma_{t-1}^2 + (\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2) \alpha_{t-1} \mathbf{x}_0}{\sigma_t^2} \right) \mathbf{x}_{t-1} \right] \right\} \\
 1181 & \\
 1182 & \\
 1183 & \\
 1184 & \propto \mathcal{N}(\mathbf{x}_{t-1}; \frac{\frac{\alpha_t}{\alpha_{t-1}} \mathbf{x}_t \sigma_{t-1}^2 + (\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2) \alpha_{t-1} \mathbf{x}_0}{\sigma_t^2}, \frac{(\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2) \sigma_{t-1}^2}{\sigma_t^2} \mathbf{I}) \\
 1185 & \\
 1186 & \\
 1187
 \end{aligned}$$

1188 where in the fourth Equation, $C(\mathbf{x}_t, \mathbf{x}_0)$ is a constant term with respect to \mathbf{x}_{t-1} computed as a
 1189 combination of only \mathbf{x}_t , \mathbf{x}_0 , and α values. We have therefore shown that at each step, $\mathbf{x}_{t-1} \sim$

1188 $q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0)$ is normally distributed, with mean $\mu_q(\mathbf{x}_t, \mathbf{x}_0)$ that is a function of \mathbf{x}_t and \mathbf{x}_0 , and
 1189 variance $\Sigma_q(t)$ as a function of α and σ coefficients. These coefficients are known and fixed at each
 1190 timestep; they are either set permanently when modeled as hyperparameters, or treated as the current
 1191 inference output of a network that seeks to model them.

1192 We can then set the variances of the two Gaussians to match exactly, optimizing the KL Divergence
 1193 term reduces to minimizing the difference between the means of the two distributions:
 1194

$$\begin{aligned} 1195 \arg \min_{\theta} D_{\text{KL}}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) \| p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t)) \\ 1196 \\ 1197 = \arg \min_{\theta} \frac{1}{2\sigma_q^2(t)} \left[\|\mu_{\theta}(\mathbf{x}_t, t) - \mu_q(\mathbf{x}_t, \mathbf{x}_0)\|_2^2 \right], \end{aligned} \quad (18)$$

1199 where $\sigma_q^2(t) = \frac{(\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2) \sigma_{t-1}^2}{\sigma_t^2}$, the derivation is the same as in Eq. (92) in Luo (2022), so we skip
 1200 the derivation here. To derive the score matching function, we appeal to Tweedie's Formula Efron
 1201 (2011), which states $\mathbb{E}[\mu_z|z] = z + \Sigma_z \nabla_z \log q(z)$ for a given Gaussian variable $z \sim \mathcal{N}(z; \mu_z, \Sigma_z)$.
 1202 In this case, we apply it to predict the true posterior mean of \mathbf{x}_t given its samples. We can obtain:
 1203

$$\begin{aligned} 1204 \mathbb{E}[\mu_{\mathbf{x}_t}|\mathbf{x}_t] &= \mathbf{x}_t + \sigma_t^2 \nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t) = \alpha_t \mathbf{x}_0 \\ 1205 \\ 1206 \therefore \mathbf{x}_0 &= \frac{\mathbf{x}_t + \sigma_t^2 \nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t)}{\alpha_t} \end{aligned} \quad (19)$$

1208 Then, we can plug Eq. (19) into our ground-truth denoising transition mean $\mu_q(\mathbf{x}_t, \mathbf{x}_0)$ once again
 1209 and derive a new form:
 1210

$$\begin{aligned} 1211 \mu_q(\mathbf{x}_t, \mathbf{x}_0) &= \frac{\frac{\alpha_t}{\alpha_{t-1}} \sigma_{t-1}^2 \mathbf{x}_t + (\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2) \alpha_{t-1} \cdot \frac{\mathbf{x}_t + \sigma_t^2 \nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t)}{\alpha_t}}{\sigma_t^2} \\ 1212 \\ 1213 &= \frac{\frac{\alpha_t}{\alpha_{t-1}} \sigma_{t-1}^2}{\sigma_t^2} \mathbf{x}_t + \frac{\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2}{\sigma_t^2 \frac{\alpha_t}{\alpha_{t-1}}} \mathbf{x}_t + \frac{(\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2) \sigma_t^2 \nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t)}{\sigma_t^2 \frac{\alpha_t}{\alpha_{t-1}}} \\ 1214 \\ 1215 &= \frac{\alpha_{t-1}}{\alpha_t} \mathbf{x}_t + \left(\frac{\alpha_{t-1}}{\alpha_t} \sigma_t^2 - \frac{\alpha_t}{\alpha_{t-1}} \sigma_{t-1}^2 \right) \nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t) \\ 1216 \\ 1217 &= \frac{\alpha_{t-1}}{\alpha_t} \left[\mathbf{x}_t + \left(\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2 \right) \mathbf{s}_{\theta}(\mathbf{x}_t, t) \right] \end{aligned} \quad (20)$$

1218 Finally, we plug Eq. (20) into our optimization function Eq. (18), and we can get:
 1219

$$\begin{aligned} 1220 \arg \min_{\theta} D_{\text{KL}}(q(\mathbf{x}_{t-1}|\mathbf{x}_t, \mathbf{x}_0) \| p_{\theta}(\mathbf{x}_{t-1}|\mathbf{x}_t)) \\ 1221 \\ 1222 = \arg \min_{\theta} \frac{1}{2\sigma_q^2(t)} \left[\|\mu_{\theta}(\mathbf{x}_t, t) - \mu_q(\mathbf{x}_t, \mathbf{x}_0)\|_2^2 \right] \\ 1223 \\ 1224 = \arg \min_{\theta} \frac{1}{2 \frac{(\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2) \sigma_{t-1}^2}{\sigma_t^2}} \cdot \left(\frac{\alpha_{t-1}}{\alpha_t} \right)^2 \cdot (\sigma_t^2 - \frac{\alpha_t^2}{\alpha_{t-1}^2} \sigma_{t-1}^2)^2 \|\mathbf{s}_{\theta}(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t)\|_2^2 \\ 1225 \\ 1226 = \frac{\sigma_t^2}{2} \left(\frac{\sigma_t^2 \alpha_{t-1}^2}{\sigma_{t-1}^2 \alpha_t^2} - 1 \right) \|\mathbf{s}_{\theta}(\mathbf{x}_t, t) - \nabla_{\mathbf{x}_t} \log q(\mathbf{x}_t)\|_2^2 \end{aligned}$$

1233 B.3 DERIVATION OF VARIATIONAL LOWER BOUND EQ. (7)

1234 To model $\log p_{\theta}(X, Z)$, we introduce an auxiliary distribution $Q(Y)$ over the latent variable Y :

$$\begin{aligned} 1235 \log p_{\theta}(X, Z) &= \int Q(Y) \log p_{\theta}(X, Z) dY \\ 1236 \\ 1237 &= \int Q(Y) \log p_{\theta}(X, Z) \frac{p_{\theta}(Y|X, Z)}{p_{\theta}(Y|X, Z)} dY \\ 1238 \\ 1239 &= \int Q(Y) \log \frac{p_{\theta}(X, Y, Z)}{Q(Y)} dY - \int Q(Y) \log \frac{p_{\theta}(Y|X, Z)}{Q(Y)} dY, \end{aligned}$$

where the first term is the ELBO and the second term is the KL divergence $\mathcal{D}_{\text{KL}}(Q(Y)||p_{\theta}(Y|X, Z))$. Since the KL divergence is non-negative, maximizing the ELBO provides a valid surrogate for maximizing $\log p_{\theta}(X, Z)$. Replacing $Q(Y)$ with $p_{\phi}(Y|X, Z)$ at each iteration will obtain as follows:

$$\begin{aligned}
 \theta^* &= \arg \max_{\theta} \log p_{\theta}(X, Z) \\
 &= \arg \max_{\theta} \mathbb{E}_{p_{\phi}(Y|X, Z)} [\log p_{\theta}(X, Y, Z)] \\
 &= \arg \max_{\theta} \mathbb{E}_{p_{\phi}(Y|X, Z)} [\log p_{\theta}(X|Z) + \log p_{\theta}(Y|X, Z) + \log p_{\theta}(Z)] \\
 &= \arg \max_{\theta} \mathbb{E}_{p_{\phi}(Y|X, Z)} [\log p_{\theta}(X|Z)] + \mathbb{E}_{p_{\phi}(Y|X, Z)} [\log p_{\theta}(Y|X, Z)] \\
 &= \arg \max_{\theta} \log p_{\theta}(X|Z) + \mathbb{E}_{p_{\phi}(Y|X, Z)} [\log p_{\theta}(Y|X, Z)].
 \end{aligned}$$

which is exactly the variational lower bound presented in Eq. (7).

B.4 DERIVATION OF CONDITIONAL ELBO IN EQ. (8)

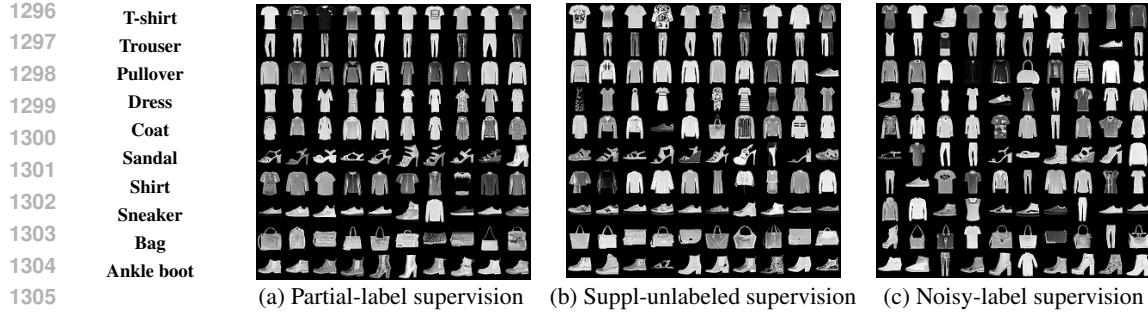
We provide a derivation of conditional ELBO in the following, which is similar to the unconditional ELBO in Ho et al. (2020).

$$\begin{aligned}
 &\log p_{\theta}(\mathbf{x}_0 | z) \\
 &= \log \int \frac{p_{\theta}(\mathbf{x}_{0:T} | z) q(\mathbf{x}_{1:T} | \mathbf{x}_0, z)}{q(\mathbf{x}_{1:T} | \mathbf{x}_0, z)} d\mathbf{x}_{1:T} \\
 &= \log \mathbb{E}_{q(\mathbf{x}_{1:T} | \mathbf{x}_0, z)} \left[\frac{p_{\theta}(\mathbf{x}_T | z) p_{\theta}(\mathbf{x}_{0:T-1} | \mathbf{x}_T, z)}{q(\mathbf{x}_{1:T} | \mathbf{x}_0, z)} \right] \\
 &\geq \mathbb{E}_{q(\mathbf{x}_{1:T} | \mathbf{x}_0, z)} \left[\log \frac{p_{\theta}(\mathbf{x}_T | z) p_{\theta}(\mathbf{x}_{0:T-1} | \mathbf{x}_T, z)}{q(\mathbf{x}_{1:T} | \mathbf{x}_0, z)} \right] \\
 &= \mathbb{E}_{q(\mathbf{x}_{1:T} | \mathbf{x}_0, z)} \left[\log \frac{p_{\theta}(\mathbf{x}_T | z) \prod_{i=0}^{T-1} p_{\theta}(\mathbf{x}_i | \mathbf{x}_{i+1}, z)}{\prod_{i=0}^{T-1} q(\mathbf{x}_{i+1} | \mathbf{x}_i, \mathbf{x}_0, z)} \right] \\
 &= \mathbb{E}_{q(\mathbf{x}_{1:T} | \mathbf{x}_0, z)} \left[\log \frac{p_{\theta}(\mathbf{x}_T | z) \prod_{i=0}^{T-1} p_{\theta}(\mathbf{x}_i | \mathbf{x}_{i+1}, z)}{\prod_{i=0}^{T-1} \frac{q(\mathbf{x}_{i+1} | \mathbf{x}_0, z) q(\mathbf{x}_i | \mathbf{x}_{i+1}, \mathbf{x}_0, z)}{q(\mathbf{x}_i | \mathbf{x}_0, z)}} \right] \\
 &= \mathbb{E}_{q(\mathbf{x}_{1:T} | \mathbf{x}_0, z)} \left[\log \frac{p_{\theta}(\mathbf{x}_T | z) \prod_{i=0}^{T-1} p_{\theta}(\mathbf{x}_i | \mathbf{x}_{i+1}, z)}{\prod_{i=0}^{T-1} q(\mathbf{x}_i | \mathbf{x}_{i+1}, \mathbf{x}_0, z)} - \log q(\mathbf{x}_T | \mathbf{x}_0, z) \right] \\
 &= \mathbb{E}_{q(\mathbf{x}_{1:T} | \mathbf{x}_0, z)} \left[\log \frac{\prod_{i=0}^{T-1} p_{\theta}(\mathbf{x}_i | \mathbf{x}_{i+1}, z)}{\prod_{i=0}^{T-1} q(\mathbf{x}_i | \mathbf{x}_{i+1}, \mathbf{x}_0, z)} - \log \frac{q(\mathbf{x}_T | \mathbf{x}_0, z)}{p_{\theta}(\mathbf{x}_T | z)} \right] \\
 &= \sum_{i=0}^{T-1} \mathbb{E}_{q(\mathbf{x}_i, \mathbf{x}_{i+1} | \mathbf{x}_0, z)} \left[\log \frac{p_{\theta}(\mathbf{x}_i | \mathbf{x}_{i+1}, z)}{q(\mathbf{x}_i | \mathbf{x}_{i+1}, \mathbf{x}_0, z)} \right] - D_{\text{KL}}(q(\mathbf{x}_T | \mathbf{x}_0, z) || p_{\theta}(\mathbf{x}_T | z)) \\
 &= \sum_{i=0}^{T-1} \mathbb{E}_{q(\mathbf{x}_{i+1} | \mathbf{x}_0, z)} \mathbb{E}_{q(\mathbf{x}_i | \mathbf{x}_{i+1}, \mathbf{x}_0, z)} \left[\log \frac{p_{\theta}(\mathbf{x}_i | \mathbf{x}_{i+1}, z)}{q(\mathbf{x}_i | \mathbf{x}_{i+1}, \mathbf{x}_0, z)} \right] - D_{\text{KL}}(q(\mathbf{x}_T | \mathbf{x}_0, z) || p_{\theta}(\mathbf{x}_T | z)) \\
 &= C_3 - \sum_{i=1}^{T-1} \mathbb{E}_{q(\mathbf{x}_{i+1} | \mathbf{x}_0, z)} [D_{\text{KL}}(q(\mathbf{x}_i | \mathbf{x}_{i+1}, \mathbf{x}_0, z) || p_{\theta}(\mathbf{x}_i | \mathbf{x}_{i+1}, z))] \\
 &= -\mathbb{E}_t \left[w_t \|\mathbf{s}_{\theta}(\mathbf{x}_t, z, t) - \nabla \log q(\mathbf{x}_t | \mathbf{x}_0, z)\|_2^2 \right] + C_2.
 \end{aligned}$$

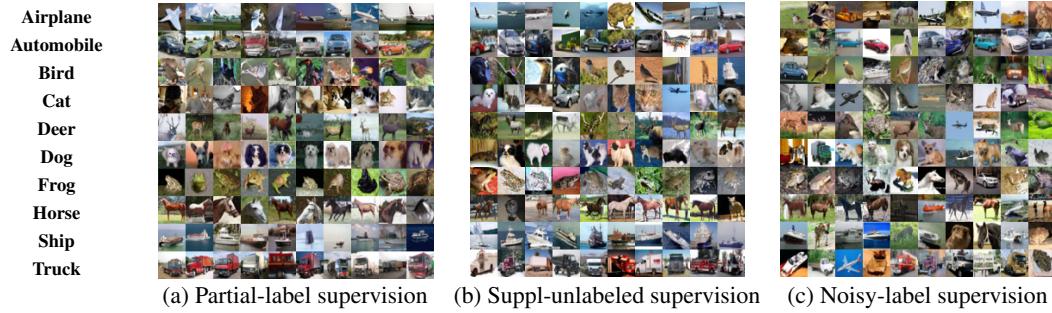
We get the result of Eq. (8).

B.5 DERIVATION OF REMARK 1

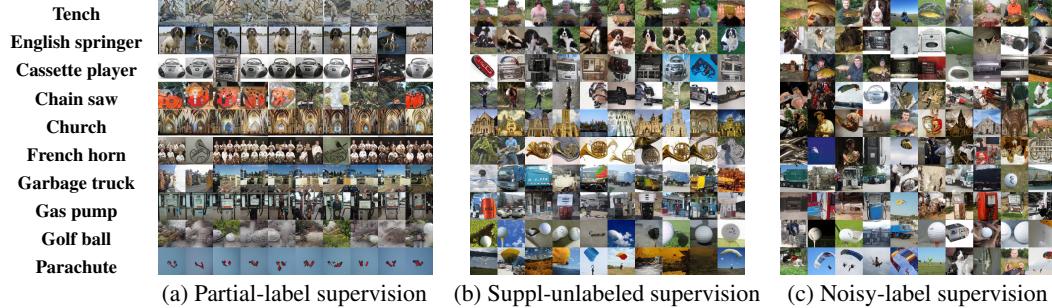
Although this result follows directly from prior studies (Vincent, 2011; Song & Ermon, 2019), we provide a brief derivation here for completeness. Let $\mathcal{L}_{\text{DSM}}(\theta; q(X, Y))$ and $\mathcal{L}_{\text{ESM}}(\theta; q(X, Y))$ denote



1306
1307 Figure 3: Examples of randomly generated Fashion-MNIST images from *Vanilla* models trained
1308 under different types of imprecise supervision.



1320
1321 Figure 4: Examples of randomly generated CIFAR-10 images from *Vanilla* models trained under
1322 different types of imprecise supervision.



1333
1334 Figure 5: Examples of randomly generated ImageNette images from *Vanilla* models trained under
1335 different types of imprecise supervision.

1336 the denoising score matching (DSM) and explicit score matching (ESM) objectives, respectively:

$$\mathcal{L}_{\text{DSM}}(\theta; q(X, Y)) := \mathbb{E}_t \left[\lambda(t) \mathbb{E}_{y \sim q(Y)} \mathbb{E}_{\mathbf{x}_t \sim q_{t|0}(\mathbf{x}_t | \mathbf{x}, y)} \left\| \mathbf{s}_\theta(\mathbf{x}_t, y, t) - \nabla_{\mathbf{x}_t} \log q_{t|0}(\mathbf{x}_t | \mathbf{x}, Y = y) \right\|_2^2 \right],$$

$$\mathcal{L}_{\text{ESM}}(\theta; q(X, Y)) := \mathbb{E}_t \left[\lambda(t) \mathbb{E}_{y \sim q(Y)} \mathbb{E}_{\mathbf{x}_t \sim q_t(\mathbf{x}_t | y)} \left\| \mathbf{s}_\theta(\mathbf{x}_t, y, t) - \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | Y = y) \right\|_2^2 \right].$$

1341 It has been established (Vincent, 2011; Song & Ermon, 2019) that these two formulations differ only
1342 by an additive constant independent of θ :

$$\mathcal{L}_{\text{ESM}}(\theta; q(X, Y)) = \mathcal{L}_{\text{DSM}}(\theta; q(X, Y)) + C_3,$$

1344 where C_3 does not depend on θ . Hence, both objectives admit the same minimizer.

1345 Applying this result to an imprecise-label dataset by identifying $Y = Z$, let $\theta_{\text{ESM}}^* := \arg \min_\theta \mathcal{L}_{\text{ESM}}(\theta; q(X, Z))$. Then the optimal score function satisfies $\mathbf{s}_{\theta_{\text{ESM}}^*}(\mathbf{x}_t, z, t) = \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | Z = z)$. Since the same conclusion holds for \mathcal{L}_{DSM} , we obtain $\mathbf{s}_{\theta_{\text{ESM}}^*} = \mathbf{s}_{\theta_{\text{DSM}}^*} = \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | z)$, which is precisely the statement of Remark 1.

To directly illustrate this bias, we train CDMs under different forms of imprecise supervision by applying Eq. (8) directly, a baseline we refer to as *Vanilla*. We then visualize the images generated by these biased models, as shown in the figures below. The results reveal the following patterns:

- **Partial-label supervision:** The generated images often lack diversity and typically capture only the dominant object. This effect is particularly pronounced on the ImageNette dataset, where samples within the same class appear highly similar. Interestingly, the generated categories generally align with the ground-truth labels, suggesting that diffusion models can still extract correct class information under partial-label supervision. However, the inherent label ambiguity prevents the model from capturing intra-class variation.
- **Noisy-label supervision:** The generated samples tend to contain visual noise. Although the model is able to capture class diversity, corrupted labels cause mismatches between generated samples and their true categories.
- **Supplementary-unlabeled supervision:** The generated images are often both less diverse and noisier. This phenomenon combines the limitations of partial-label supervision with the challenge of abundant unlabeled samples. Because the model has limited access to labeled examples, it relies on averaging confidence across all classes, which reduces its discriminative boundaries and introduces noise.

B.6 PROOF OF THEOREM 1

The derivation here is analogous to that of Theorem 1 in Na et al. (2024), and we provide the full proof below for completeness. First, for all t , the perturbed distribution $q_t(\mathbf{x}_t|z)$ satisfies:

$$q_t(\mathbf{x}_t|z) = \sum_{y=1}^c p(y|z) q_t(\mathbf{x}_t|y) \quad \forall \mathbf{x}_t \in \mathcal{X}, z \in \mathcal{Y}.$$

This implies that the transition from imprecise labels to clean labels is independent of the timesteps. Consequently, Eq. (9) can be derived as follows,

$$\begin{aligned} & \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t|z) \\ &= \frac{\nabla_{\mathbf{x}_t} q_t(\mathbf{x}_t|z)}{q_t(\mathbf{x}_t|z)} \\ &= \frac{\sum_{y=1}^c p(y|z) \nabla_{\mathbf{x}_t} q_t(\mathbf{x}_t|y)}{q_t(\mathbf{x}_t|z)} \\ &= \sum_{y=1}^c \frac{p(y|z) q_t(\mathbf{x}_t|y)}{q_t(\mathbf{x}_t|z)} \cdot \frac{\nabla_{\mathbf{x}_t} q_t(\mathbf{x}_t|y)}{q_t(\mathbf{x}_t|y)} \\ &= \sum_{y=1}^c \frac{p(y|z) q_t(\mathbf{x}_t|y)}{q_t(\mathbf{x}_t|z)} \cdot \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t|y) \\ &= \sum_{y=1}^c p(y|z) \cdot \frac{p(z)}{p(y)} \cdot \frac{p(y|\mathbf{x}_t)}{p(z|\mathbf{x}_t)} \cdot \frac{q_t(\mathbf{x}_t)}{q_t(\mathbf{x}_t)} \cdot \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t|y) \\ &= \sum_{y=1}^c p(z|y) \cdot \frac{p(y|\mathbf{x}_t)}{p(z|\mathbf{x}_t)} \cdot \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t|y) \\ &= \sum_{y=1}^c p(z|y, \mathbf{x}_t) \cdot \frac{p(y|\mathbf{x}_t)}{p(z|\mathbf{x}_t)} \cdot \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t|y) \quad (\text{Conditional indep. of } z \text{ and } \mathbf{x}_t \text{ given } y.) \\ &= \sum_{y=1}^c \frac{p(z|y, \mathbf{x}_t) p(y|\mathbf{x}_t)}{p(z|\mathbf{x}_t)} \cdot \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t|y) \\ &= \sum_{y=1}^c p(y|\mathbf{x}_t, z) \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t|y) \end{aligned}$$

1404 B.7 PROOF OF PROPOSITION 1
14051406 By Remark 1 and Theorem 1, the optimal solution θ_{Gen}^* to Eq. (10) satisfies
1407

1408
$$\sum_{y=1}^c p(y|\mathbf{x}_t, z) \mathbf{s}_{\theta_{\text{Gen}}^*}(\mathbf{x}_t, y, t) = \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | z) = \sum_{y=1}^c p(y|\mathbf{x}_t, z) \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | y),$$

1409

1410 for all $\mathbf{x}_t \in \mathcal{X}$, $z \subseteq \mathcal{Y}$, and $t \in [T]$.
14111412 Next, recall the weighted denoising score matching loss:
1413

1414
$$\mathcal{L}_{\text{Gen}}(\theta) = \mathbb{E}_t \left[w_t \left\| \sum_{y=1}^c p(y|\mathbf{x}_t, z) \mathbf{s}_\theta(\mathbf{x}_t, y, t) - \sum_{y=1}^c p(y|\mathbf{x}_t, z) \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | y) \right\|_2^2 \right].$$

1415

1416 Differentiating with respect to $\mathbf{s}_\theta(\mathbf{x}_t, y, t)$ and setting the derivative to zero yields
1417

1418
$$\frac{\partial}{\partial \mathbf{s}_\theta(\mathbf{x}_t, y, t)} \mathcal{L}_{\text{Gen}}(\theta) = 2w_t p(y|\mathbf{x}_t, z) (\mathbf{s}_{\theta_{\text{Gen}}^*}(\mathbf{x}_t, y, t) - \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | y)) = 0.$$

1419

1420 Since $w_t > 0$, for any y such that $p(y|\mathbf{x}_t, z) > 0$, the optimality condition implies
1421

1422
$$\mathbf{s}_{\theta_{\text{Gen}}^*}(\mathbf{x}_t, y, t) = \nabla_{\mathbf{x}_t} \log q_t(\mathbf{x}_t | y).$$

1423

1424 In particular, under the partial-label learning setting, if $p(y|\mathbf{x}_t, z) = 0$, the loss does not depend on
1425 $\mathbf{s}_\theta(\mathbf{x}_t, y, t)$, and the equality can be established without loss of generality. This completes the proof.
14261427 B.8 PROOF OF THEOREM 2
14281429 We first consider the case where the timestep τ is sampled from a log-normal distribution, as defined
1430 in the EDM framework. Specifically,
1431

1432
$$\ln(\tau) \sim \mathcal{N}(P_{\text{mean}}, P_{\text{std}}^2),$$

1433

1434 where the parameters are set to $P_{\text{mean}} = 1.2$ and $P_{\text{std}} = -1.2$. Accordingly, the probability density
1435 function of τ is given by
1436

1437
$$p(\tau) = \frac{1}{\tau P_{\text{std}} \sqrt{2\pi}} \exp \left(-\frac{(\ln \tau - P_{\text{mean}})^2}{2P_{\text{std}}^2} \right), \quad \tau > 0.$$

1438

1439 The corresponding cumulative distribution function (CDF) is denoted as:
1440

1441
$$F(\tau) = \frac{1}{2} \left[1 + \text{erf} \left(\frac{\ln \tau - P_{\text{mean}}}{P_{\text{std}} \sqrt{2}} \right) \right],$$

1442

1443 where $\text{erf}(x)$ denotes the error function.
14441445 The median of this distribution τ_{mid} is the value at which the CDF equals 0.5, i.e., $F(\tau_{\text{mid}}) = 0.5$.
14461447 To ensure that the selected subinterval allows signal-dominant early timesteps and noise-dominant
1448 later timesteps to complement each other, we require the cumulative probability mass on either side
1449 of the median to be equal. Formally, for subinterval boundaries (l, r) with $r = l + \Delta$, we enforce the
1450 following symmetry condition:
1451

1452
$$F(r) - F(\tau_{\text{mid}}) = F(\tau_{\text{mid}}) - F(l).$$

1453

1454 Rewriting this with $r = l + \Delta$ gives:
1455

1456
$$F(l + \Delta) + F(l) = 2F(\tau_{\text{mid}}) = 1.$$

1457

1458 This implicit equation defines the subinterval $(l, l + \Delta)$ such that the cumulative probability mass is
1459 centered around the median of $p(\tau)$. To compute the left boundary l , we solve:
1460

1461
$$l = \text{Solve}_\tau (F(\tau) + F(\tau + \Delta) - 1 = 0), \quad (21)$$

1462

1463 and set $r = l + \Delta$. The solution can be obtained using any standard root-finding algorithm, such as
1464 the Brent method (Brent, 2013).
1465

1458 We then consider the DDPM setting, where the timestep τ is uniformly sampled from a fixed interval.
 1459 Specifically, we assume $\tau \sim \mathcal{U}(0, 1)$, whose CDF is given by
 1460

$$F(\tau) = \tau$$

1462 Under this distribution, the symmetry condition in Eq. (21) simplifies to
 1463

$$\begin{aligned} l &= \text{Solve}_\tau (F(\tau) + F(\tau + \Delta) - 1 = 0) \\ &= \text{Solve}_\tau (\tau + \tau + \Delta - 1 = 0) \\ &= \frac{1 - \Delta}{2}, \end{aligned}$$

1468 and thus $r = l + \Delta = \frac{1 - \Delta}{2}$. This result implies that the optimal subinterval is symmetric around the
 1469 midpoint of the distribution. In the special case where only a single timestep is used (i.e., $\Delta \rightarrow 0$),
 1470 the best estimate of the conditional ELBO occurs exactly at the median. As the sampled timestep
 1471 deviates further from the midpoint, classification accuracy tends to degrade. This observation aligns
 1472 with the empirical findings of Li et al. (2023), who reported that classification accuracy is maximized
 1473 near the median and declines towards the edges. Their use of evenly spaced timesteps centered
 1474 around the median further supports our strategy.

1475 B.9 PROOF OF THEOREM 3

1476 For clarity, we abbreviate $\hbar(\tau, y)$ as $\hbar(\tau)$, since the proof does not depend explicitly on y . Define the
 1477 weighted integral of \hbar and the normalization factor over an interval $[l, r]$ as
 1479

$$A(l, r) := \int_l^r \hbar(\tau) p(\tau) d\tau, \quad Z(l, r) := \int_l^r p(\tau) d\tau,$$

1480 so that the local expectation can be written as $\mu' = A(l, r)/Z(l, r)$. Let $\mu'' = \mathbb{E}_{\tau \sim p(\tau)}[\hbar(\tau)]$ denote
 1481 the global expectation. The squared error objective in Eq. (16) then becomes
 1484

$$g(l, r) := (\mu' - \mu'')^2,$$

1485 subject to the probability-mass constraint $Z(l, r) = \alpha$.
 1487

1488 We apply the method of Lagrange multipliers with
 1489

$$L(l, r, \lambda) := (\mu' - \mu'')^2 + \lambda \left(\int_l^r p(\tau) d\tau - \alpha \right).$$

1492 By the Leibniz rule, the derivatives of $A(l, r)$ and $Z(l, r)$ with respect to the interval boundaries are
 1493

$$\frac{\partial A(l, r)}{\partial l} = -\hbar(l) p(l), \quad \frac{\partial A(l, r)}{\partial r} = \hbar(r) p(r), \quad \frac{\partial Z(l, r)}{\partial l} = -p(l), \quad \frac{\partial Z(l, r)}{\partial r} = p(r).$$

1496 Hence, the derivatives of $\mu' = A/Z$ are
 1497

$$\frac{\partial \mu'}{\partial l} = \frac{p(l)}{Z(l, r)} (\mu' - \hbar(l)), \quad \frac{\partial \mu'}{\partial r} = \frac{p(r)}{Z(l, r)} (\hbar(r) - \mu').$$

1500 Differentiating L w.r.t. l and r gives
 1501

$$\begin{aligned} \frac{\partial L}{\partial l} &= 2(\mu' - \mu'') \cdot \frac{p(l)}{Z(l, r)} (\mu' - \hbar(l)) - \lambda p(l), \\ \frac{\partial L}{\partial r} &= 2(\mu' - \mu'') \cdot \frac{p(r)}{Z(l, r)} (\hbar(r) - \mu') + \lambda p(r). \end{aligned}$$

1507 Setting both derivatives to zero yields the necessary conditions
 1508

$$2(\mu' - \mu'')(\mu' - \hbar(l)) = \lambda Z(l, r), \quad 2(\mu' - \mu'')(\hbar(r) - \mu') = \lambda Z(l, r).$$

1510 Equating the two expressions gives
 1511

$$\mu' - \hbar(l) = \hbar(r) - \mu' \implies \mu' = \frac{1}{2}(\hbar(l) + \hbar(r)).$$

1512 Substituting back, we obtain
 1513
 1514
$$\int_l^r p(\tau) \hbar(\tau, y) d\tau = \frac{Z(l, r)}{2} (\hbar(l, y) + \hbar(r, y)), \quad Z(l, r) := \int_l^r p(\tau) d\tau.$$

 1515

1516 Equivalently, the necessary optimality condition is $\text{ERR}(l^*, r^*, y) = 0$. Since $Z(l^*, r^*) > 0$, this is
 1517 also equivalent to
 1518

$$\text{ERR}(l^*, r^*, y) := \mathbb{E}_{\tau \sim p(\tau | \tau \in [l^*, r^*])} [\hbar(\tau, y)] - \frac{1}{2} (\hbar(l^*, y) + \hbar(r^*, y)) = 0.$$

1521 This establishes the necessary condition for an optimal subinterval.
 1522

1523 C DISCCUSION

1525 C.1 ANALYSIS OF EARLY-LEARNING REGULARIZATION IN EQ. (15)

1527 The effectiveness of Eq. (15) can be better understood by examining the form of its gradient. For
 1528 clarity, we restate the loss with the following notation: given a noisy-labeled input (\mathbf{x}, \hat{y}) , we denote
 1529 the model’s output probabilities as $f_\theta(\mathbf{x})$ and the corresponding EMA target as $f_\phi(\mathbf{x})$.

1530 Let $\hat{\mathbf{y}} \in \mathbb{R}^c$ be the one-hot vector corresponding to the noisy label \hat{y} . Then the loss over the whole
 1531 dataset $\mathcal{D} = \{(\mathbf{x}^{[i]}, \hat{\mathbf{y}}^{[i]})\}_{i=1}^n$ can be computed according to Eq. (15) as
 1532

$$\mathcal{L}_{\text{Cls}}^{\text{NL}}(\mathcal{D}) = -\frac{1}{n} \sum_{i=1}^n \langle \text{sg}(\mathbf{r}^{[i]}), \log f_\theta(\mathbf{x}^{[i]}) \rangle, \quad \mathbf{r}^{[i]} = \hat{\mathbf{y}}^{[i]} - \lambda \frac{f_\theta(\mathbf{x}^{[i]}) \odot (\delta^{[i]} \mathbf{1} - f_\phi(\mathbf{x}^{[i]}))}{1 - \delta^{[i]}}, \quad (22)$$

1535 where $\delta^{[i]} = \langle f_\theta(\mathbf{x}^{[i]}), f_\phi(\mathbf{x}^{[i]}) \rangle$, $\text{sg}(\cdot)$ denotes the stop-gradient operator, and \odot is the Hadamard
 1536 product. By construction $\mathbf{r}^{[i]}$ is treated as a *constant* w.r.t. θ due to the stop-gradient.

1537 **Lemma 1.** *Let $\psi_\theta(\mathbf{x})$ denote the pre-softmax logits such that $f_\theta(\mathbf{x}) = \text{softmax}(\psi_\theta(\mathbf{x}))$. For the loss
 1538 in Eq. (15), the gradients are*

$$\frac{\partial \mathcal{L}_{\text{Cls}}^{\text{NL}}(\mathbf{x}^{[i]})}{\partial \psi_\theta(\mathbf{x}^{[i]})} = f_\theta(\mathbf{x}^{[i]}) - \text{sg}(\mathbf{r}^{[i]}), \quad \text{for each } i = 1, \dots, n, \quad (23)$$

1540 and, by the chain rule,

$$\nabla_\theta \mathcal{L}_{\text{Cls}}^{\text{NL}}(\mathcal{D}) = \frac{1}{n} \sum_{i=1}^n J_{\mathbf{z}_\theta}(\mathbf{x}^{[i]})^\top \left[f_\theta(\mathbf{x}^{[i]}) - \text{sg}(\mathbf{r}^{[i]}) \right], \quad (24)$$

1543 where $J_{\mathbf{z}_\theta}(\mathbf{x}) = \partial \mathbf{z}_\theta(\mathbf{x}) / \partial \theta$ is the Jacobian of the logits w.r.t. the parameters. Equivalently, expanding
 1544 $\mathbf{r}^{[i]}$ gives

$$\nabla_\theta \mathcal{L}_{\text{Cls}}^{\text{NL}}(\mathcal{D}) = \frac{1}{n} \sum_{i=1}^n J_{\mathbf{z}_\theta}(\mathbf{x}^{[i]})^\top \left[f_\theta(\mathbf{x}^{[i]}) - \hat{\mathbf{y}}^{[i]} + \lambda \text{sg} \left(\frac{f_\theta(\mathbf{x}^{[i]}) \odot (\delta^{[i]} \mathbf{1} - f_\phi(\mathbf{x}^{[i]}))}{1 - \delta^{[i]}} \right) \right]. \quad (25)$$

1551 *Proof.* For any $i \in \{1, \dots, n\}$, let us first verify that $\mathbf{r}^{[i]}$ sums to 1. With

$$\mathbf{r}^{[i]} = \hat{\mathbf{y}}^{[i]} - \lambda \frac{f_\theta(\mathbf{x}^{[i]}) \odot (\delta^{[i]} \mathbf{1} - f_\phi(\mathbf{x}^{[i]}))}{1 - \delta^{[i]}},$$

1555 we sum over classes and using $\langle f_\theta(\mathbf{x}^{[i]}), \mathbf{1} \rangle = 1$ yields

$$\mathbf{1}^\top \mathbf{r}^{[i]} = 1 - \lambda \frac{\delta^{[i]} - \langle f_\theta(\mathbf{x}^{[i]}), f_\phi(\mathbf{x}^{[i]}) \rangle}{1 - \delta^{[i]}} = 1,$$

1559 so $\mathbf{r}^{[i]}$ lies on the simplex (hence Eq. (22) is an ordinary cross-entropy with a fixed target). Let
 1560 $\psi^{[i]} = \psi_\theta(\mathbf{x}^{[i]})$ be the logits and recall $\frac{\partial \text{logsoftmax}(\psi)}{\partial \psi} = I - \text{softmax}(\psi) \mathbf{1}^\top$. For the per-sample
 1561 loss $\ell^{[i]} = -\langle \text{sg}(\mathbf{r}^{[i]}), \log \text{softmax}(\psi^{[i]}) \rangle$, the derivative w.r.t. logits is

$$\frac{\partial \ell^{[i]}}{\partial \psi^{[i]}} = \text{softmax}(\psi^{[i]}) - \text{sg}(\mathbf{r}^{[i]}) = f_\theta(\mathbf{x}^{[i]}) - \text{sg}(\mathbf{r}^{[i]}),$$

1563 which is Eq. (23). Applying the chain rule and averaging over i gives Eq. (24). Replacing $\text{sg}(\mathbf{r}^{[i]})$ by
 1564 its explicit form produces Eq. (25). \square

1566
 1567 **Remark.** Eq. (25) shows that $\mathcal{L}_{\text{Cls}}^{\text{NL}}$ behaves like the standard cross-entropy gradient plus an ELR-like
 1568 corrective term. This term amplifies gradients on clean samples and counteracts gradients on noisy
 1569 samples. Specifically, we expand this ELR-like corrective term into:
 1570

$$\mathbf{g}_y^{[i]} := \frac{f_\theta(\mathbf{x}^{[i]})}{1 - \langle f_\theta(\mathbf{x}^{[i]}), f_\phi(\mathbf{x}^{[i]}) \rangle} \sum_{k=1}^c (f_\phi(\mathbf{x}^{[i]})_k - f_\phi(\mathbf{x}^{[i]})_y) f_\theta(\mathbf{x}^{[i]})_k. \quad (26)$$

1571
 1572 If y^* is the true class, then the y^* th entry of $f_\phi(\mathbf{x}^{[i]})$ tends to be dominant during early-learning. In
 1573 that case, the y^* th entry of $\mathbf{g}^{[i]}$ is negative. This is useful both for examples with clean labels and for
 1574 examples with noisy labels. For examples with clean labels, the cross-entropy term $f_\theta(\mathbf{x}^{[i]}) - \hat{\mathbf{y}}^{[i]}$
 1575 tends to vanish after the early-learning stage because $f_\theta(\mathbf{x}^{[i]})$ is very close to $\hat{\mathbf{y}}^{[i]}$, allowing examples
 1576 with wrong labels to dominate the gradient. Adding $\mathbf{g}^{[i]}$ counteracts this effect by ensuring that
 1577 the magnitudes of the coefficients on examples with clean labels remain large. Thus, $\mathbf{g}^{[i]}$ fulfills the
 1578 two desired properties that boosting the gradient of examples with clean labels, and neutralizing the
 1579 gradient of the examples with false labels.
 1580

1581 C.2 CLASS-PRIOR ESTIMATION IN IMPRECISE-LABEL DATASETS

1582 When the class priors $p(y)$ (here we slightly abuse notation and denote them as π_y) are not directly
 1583 accessible to the learning algorithm, they can be estimated using off-the-shelf estimation methods (Luo
 1584 et al., 2024; Wang et al., 2022a). In this section, we present the problem formulation and outline how
 1585 class priors can be estimated in practice.
 1586

1587 C.2.1 CLASS-PRIOR ESTIMATION IN PARTIAL-LABEL DATASETS

1588 In partial-label learning, each instance is associated with a candidate label set rather than a single
 1589 ground-truth label. This label ambiguity makes it difficult to estimate the class prior distribution,
 1590 since simply counting training samples per class is no longer feasible. To address this issue, we adopt
 1591 an iterative estimation strategy that updates the class prior in a moving-average manner.
 1592

1593 We use the model’s predicted labels as a proxy for class prior estimation. Since predictions in the
 1594 early stage of training are often unreliable, we design a moving-average update rule that gradually
 1595 stabilizes the estimated distribution. The update starts from a uniform prior $\mathbf{r} = [1/c, \dots, 1/c]$, and
 1596 is refined at each training epoch as
 1597

$$\mathbf{r} \leftarrow \mu \mathbf{r} + (1 - \mu) \mathbf{z}, \quad \mathbf{z}_j = \frac{1}{n} \sum_{i=1}^n \mathbb{I}\left(j = \arg \max_{y \in S_i} f_j(x_i)\right), \quad (27)$$

1598 where $\mu \in [0, 1]$ is a momentum parameter, S_i is the candidate label set for sample x_i , and $f_j(x_i)$
 1599 denotes the model prediction for class j . This rule progressively refines \mathbf{r} as the model improves,
 1600 leading to more accurate and stable class-prior estimates.
 1601

1602 C.2.2 CLASS-PRIOR ESTIMATION IN SUPPLEMENTARY-UNLABELED DATASETS

1603 In the case of supplementary-unlabeled datasets, which also is called semi-supervised datasets,
 1604 the estimation of class-prior is relatively straightforward. We assume that the distribution of the
 1605 labeled dataset is consistent with that of the unlabeled dataset. Therefore, the class-prior can be
 1606 directly obtained by counting the class distribution over the labeled dataset, which serves as a reliable
 1607 approximation of the overall data distribution.
 1608

1609 C.2.3 CLASS-PRIOR ESTIMATION IN NOISY-LABEL DATASETS

1610 We consider the widely adopted class-dependent label noise setting (Yao et al., 2020), where the
 1611 observed noisy label of each $\mathbf{x} \in \mathcal{X}$ depends only on its underlying clean label. Formally, the
 1612 transition probability from class i to class j is defined as
 1613

$$P(\tilde{Y} = e_j | Y = e_i, X = \mathbf{x}) = P(\tilde{Y} = e_j | Y = e_i) = T_{ij}, \quad \forall i, j \in [[c]],$$

1614 where $\mathbf{T} = [T_{ij}] \in [0, 1]^{c \times c}$ is the noise transition matrix. To make the estimation of \mathbf{T} feasible, we
 1615 follow prior work and impose the following assumptions.
 1616

1620 **Assumption 1** (Sufficiently Scattered Assumption (Li et al., 2021b)). The clean class posterior $P(Y|X) = [P(Y = e_1|X), \dots, P(Y = e_c|X)]^\top \in [0, 1]^c$ is said to be sufficiently scattered if there
 1621 exists a set $\mathcal{H} = \{\mathbf{x}_1, \dots, \mathbf{x}_m\}$ such that the matrix $\mathbf{H} = [P(Y|X = \mathbf{x}_1), \dots, P(Y|X = \mathbf{x}_m)]$
 1622 satisfies: (i) $\mathcal{Q} \subseteq \text{cone}\{\mathbf{H}\}$, where $\mathcal{Q} = \{\mathbf{v} \in \mathbb{R}^c \mid \mathbf{v}^\top \mathbf{1} \geq \sqrt{c-1} \|\mathbf{v}\|_2\}$, and $\text{cone}\{\mathbf{H}\}$ denotes
 1623 the convex cone generated by the columns of \mathbf{H} ; (ii) $\text{cone}\{\mathbf{H}\} \not\subseteq \text{cone}\{\mathbf{U}\}$ for any unitary matrix
 1624 $\mathbf{U} \in \mathbb{R}^{c \times c}$ that is not a permutation matrix.
 1625

1626 **Assumption 2** (Nonsingular \mathbf{T}). The transition matrix \mathbf{T} is nonsingular, i.e., $\text{Rank}(\mathbf{T}) = c$.
 1627

1628 Assumption 1 ensures that the clean posteriors are sufficiently scattered so that the ground-truth \mathbf{T}
 1629 can be identified, while Assumption 2 guarantees the invertibility of \mathbf{T} .
 1630

1631 Let ϵ denote the noise rate. For symmetric label noise, we have $T_{ii} = 1 - \epsilon$ and $T_{ij} = \frac{\epsilon}{c-1}$ with
 1632 $j \neq i$. In practice, the transition matrix can be estimated by solving the following optimization
 1633 problem (Li et al., 2021b):
 1634

$$\min_{\theta, \widehat{\mathbf{T}}} L(\theta, \widehat{\mathbf{T}}) = \frac{1}{n} \sum_{i=1}^n \ell(\widehat{\mathbf{T}}^\top h_\theta(\mathbf{x}_i), \widetilde{y}_i) + \lambda \cdot \log \det(\widehat{\mathbf{T}}), \quad (28)$$

1637 where ℓ is a loss function (typically cross-entropy), $h_\theta(\cdot)$ is the output of a neural network parameter-
 1638 ized by θ , and the regularization term $\log \det(\widehat{\mathbf{T}})$ encourages the estimated transition matrix to have
 1639 minimal simplex volume. Here $\lambda > 0$ is a trade-off hyperparameter. By Assumption 1, the solution
 1640 $\widehat{\mathbf{T}}$ converges to the true \mathbf{T} given sufficient noisy data (Theorem 1 in (Li et al., 2021b)).
 1641

1642 Once the transition matrix \mathbf{T} is estimated, the clean class prior $\pi = [\pi_1, \dots, \pi_c]^\top$ can be obtained by
 1643 solving the following system of linear equations:
 1644

$$\begin{cases} \widetilde{\pi}_1 = T_{11}\pi_1 + T_{21}\pi_2 + \dots + T_{c1}\pi_c \\ \widetilde{\pi}_2 = T_{12}\pi_1 + T_{22}\pi_2 + \dots + T_{c2}\pi_c \\ \vdots \\ \widetilde{\pi}_c = T_{1c}\pi_1 + T_{2c}\pi_2 + \dots + T_{cc}\pi_c \end{cases}, \quad (29)$$

1649 where $\widetilde{\pi}_i = P(\widetilde{Y} = e_i)$ is the noisy class prior of the i -th class. The empirical estimate of $\widetilde{\pi}_i$ can be
 1650 computed as
 1651

$$\widehat{\pi}_i = \frac{1}{n} \sum_{j=1}^n \mathbf{1}\{\widetilde{y}_j = e_i\}, \quad \forall i \in [[c]]. \quad (30)$$

1654 Solving this system yields the clean class prior π , which is then used in subsequent modeling.
 1655

D IMPLEMENTATION DETAILS

1656 Our implementation is based on PyTorch 1.12 (Paszke et al., 2019), and all experiments were
 1657 conducted on NVIDIA Tesla A100 GPUs with CUDA 12.4.
 1658

1659 **Imprecise-label construction.** For all class-dependent partial-label datasets, we construct a
 1660

1661 10×10 circulant transition matrix $\begin{bmatrix} 1 & q+0.2 & q & q-0.2 & \dots & q+0.2 & q & q-0.2 \\ q-0.2 & 1 & q+0.2 & q & \dots & q & q-0.2 & q+0.2 \\ q & q-0.2 & 1 & q+0.2 & \dots & q-0.2 & q+0.2 & q \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ q+0.2 & q & q-0.2 & 1 & \dots & q & q-0.2 & 1 \end{bmatrix}$, where
 1662

1663 each row maps a true label to a candidate set of labels with varying probabilities, and q is set
 1664 to 0.5. For noisy-label datasets with asymmetric noise (40% flip probability), we adopt the
 1665 following mappings: *Fashion-MNIST*: ‘Pullover’ → ‘Sneaker’, ‘Dress’ → ‘Bag’, ‘Sandal’ → ‘Shirt’,
 1666 ‘Shirt’ → ‘Sandal’. *CIFAR-10*: ‘Truck’ → ‘Automobile’, ‘Bird’ → ‘Airplane’, ‘Deer’ → ‘Horse’,
 1667 ‘Cat’ → ‘Dog’, ‘Dog’ → ‘Cat’. *ImageNette*: ‘Tench’ → ‘English springer’, ‘Cassette player’ → ‘Garbage
 1668 truck’, ‘Chain saw’ → ‘Church’, ‘Golf ball’ → ‘Parachute’, ‘Parachute’ → ‘Golf ball’.
 1669

1670 **Model setup.** The overall diffusion framework follows EDM (Karras et al., 2022), and the training
 1671 hyperparameters are kept consistent with those reported therein. For all experiments, we adopt the
 1672 DDPM++ network architecture with a U-Net backbone. Specifically, we employ the Adam optimizer
 1673

1674 with a learning rate of $1e-3$, parameters $(\beta_1, \beta_2) = (0.9, 0.999)$, and $\epsilon = 1e-8$. The EMA decay is
 1675 set to 0.5. We use a batch size of 128 for Fashion-MNIST, 64 for CIFAR-10, and 16 for ImageNette.
 1676 For the diffusion classifier, we set the timestep interval length Δ to 6.4. All models are trained from
 1677 scratch for 200k iterations.
 1678

1679 E EXPERIMENTS

1681 E.1 EVALUATION METRICS

1683 We evaluate the trained CDMs using four unconditional metrics, including Fréchet Inception Dis-
 1684 tance (FID) (Heusel et al., 2017), Inception Score (IS) (Salimans et al., 2016), Density, and Cov-
 1685 erage (Naeem et al., 2020), and three conditional metrics, namely CW-FID, CW-Density, CW-
 1686 Coverage (Chao et al., 2022). All metrics are computed using the official implementation of
 1687 DLSM (Chao et al., 2022). Although these metrics have been introduced in related work (Na
 1688 et al., 2024), we briefly recap them here for completeness and clarity.

1689 **Unconditional metrics.** Unconditional metrics evaluate generated samples without reference to class
 1690 labels. In our experiments, images are first generated conditionally per class but then pooled without
 1691 labels when computing the metrics. This evaluation protocol is consistent with prior studies (Kaneko
 1692 et al., 2019; Chao et al., 2022).

- 1693 • FID measures the distance between real and generated image distributions in the pre-trained
 1694 feature space (Szegedy et al., 2016), indicating the fidelity and diversity of generated images.
 1695
- 1696 • IS evaluates whether generated images belong to distinct classes and whether each image is
 1697 class-consistent, reflecting the realism and class separability of generated images.
 1698
- 1699 • Density and Coverage are reliable versions of Precision and Recall (Naeem et al., 2020),
 1700 respectively. Density measures how well generated samples cover real data distribution,
 1701 while Coverage assesses how well real samples are represented by generated ones.

1702 **Conditional metrics.** To measure conditional consistency, we adopt class-wise (CW) variants
 1703 of the above metrics, which compute each metric separately within each class and then average
 1704 across classes. Notably, CW-FID (also called intra-FID) is widely used in conditional generative
 1705 modeling (Miyato & Koyama, 2018; Kaneko et al., 2019), and has been highlighted as a key measure
 1706 of conditional distribution quality.

1707 **Remark:** It should be noted that the Fashion-MNIST dataset is not suitable for evaluation using these
 1708 metrics, so we do not perform evaluation on the Fahsion-MNIST dataset.

1709 E.2 FULL RESULTS IN WEAKLY SUPERVISED LEARNING

1711 Building on the experiments presented in the main text, we further provide an extended comparison
 1712 with a broader set of methods to ensure a comprehensive evaluation. The details are summarized as
 1713

1714 **Partial-label learning.** We compare against ten representative baselines: *PRODEN* (Lv et al., 2020),
 1715 *CAVL* (Zhang et al., 2021b), *POP* (Xu et al., 2023), *CC* (Feng et al., 2020), *LWS* (Wen et al., 2021),
 1716 *IDGP* (Qiao et al., 2023), *PiCO* (Wang et al., 2023), *ABLE* (Xia et al., 2022), *CRDPLL* (Wu et al.,
 1717 2022), and *DIRK* (Wu et al., 2024). For a fair comparison, we follow the hyperparameter settings
 1718 used in *PLENCH* (Wang et al., 2025b). The complete results are reported in Table 5.

1719 **Semi-supervised learning.** We follow the training and evaluation protocols of *USB* (Wang et al.,
 1720 2022c), a widely adopted benchmark for fair and unified SSL comparisons. Our baselines cover a
 1721 broad spectrum of recent approaches. First, we include confidence-thresholding methods such as
 1722 *FixMatch* (Sohn et al., 2020), *FlexMatch* (Zhang et al., 2021a), *FreeMatch* (Wang et al., 2022d),
 1723 *ReMixMatch* (Berthelot et al., 2019a), *Dash* (Xu et al., 2021) and *UDA* (Xie et al., 2020). Second,
 1724 we consider contrastive-learning based and pseudo-label based methods, including *CoMatch* (Li
 1725 et al., 2021a), *SoftMatch* (Chen et al., 2023) and *SimMatch* (Zheng et al., 2022). Finally, we add
 1726 several classical and widely studied SSL approaches, including *Pseudo-Labeling* (Lee et al., 2013),
 1727 *VAT* (Miyato et al., 2018) and *Mean Teacher* (Tarvainen & Valpola, 2017). This diverse collection of
 1728 baselines allows us to rigorously examine whether our framework remains competitive against both
 1729 state-of-the-art and classical SSL methods under consistent experimental setups.

1728
 1729 **Table 5:** Classification results on Fashion-MNIST, CIFAR-10, and ImageNette datasets under various
 1730 types of partial-label supervision. **Bold** numbers indicate the best performance.

Method	Fashion-MNIST		CIFAR-10		ImageNette	
	Random	Class-50%	Random	Class-50%	Random	Class-50%
<i>PRODEN</i>	93.31 \pm 0.07	93.44 \pm 0.21	90.02 \pm 0.22	90.44 \pm 0.44	84.75 \pm 0.13	83.50 \pm 0.60
<i>CAVL</i>	93.09 \pm 0.17	92.67 \pm 0.25	87.28 \pm 0.64	87.16 \pm 0.58	41.69 \pm 4.12	46.46 \pm 7.15
<i>POP</i>	93.59 \pm 0.17	93.57 \pm 0.19	89.13 \pm 0.22	90.19 \pm 0.10	84.65 \pm 0.55	84.29 \pm 0.17
<i>CC</i>	93.17 \pm 0.32	92.65 \pm 0.29	88.40 \pm 0.24	89.12 \pm 0.23	81.11 \pm 0.50	80.74 \pm 0.68
<i>IDGP</i>	92.26 \pm 1.25	93.07 \pm 0.16	89.65 \pm 0.53	90.83 \pm 0.34	84.07 \pm 0.26	82.18 \pm 0.13
<i>PiCO</i>	93.32 \pm 0.12	93.32 \pm 0.33	86.40 \pm 0.89	87.51 \pm 0.66	82.15 \pm 0.23	84.41 \pm 0.93
<i>ABLE</i>	93.02 \pm 0.26	93.20 \pm 0.16	90.77 \pm 0.33	90.74 \pm 0.48	71.81 \pm 2.46	76.53 \pm 1.28
<i>CRDPLL</i>	94.03 \pm 0.14	93.80 \pm 0.23	92.74 \pm 0.26	92.89 \pm 0.27	84.31 \pm 0.25	88.08 \pm 0.34
<i>DIRK</i>	94.11 \pm 0.22	93.99 \pm 0.24	93.48 \pm 0.14	93.22 \pm 0.37	87.90 \pm 0.11	87.47 \pm 0.17
<i>Vanilla</i>	80.20 \pm 1.29	66.03 \pm 1.43	60.25 \pm 0.17	56.34 \pm 0.50	56.04 \pm 0.61	59.47 \pm 0.51
<i>DMIS^{CE}</i>	84.24 \pm 0.37	78.45 \pm 0.46	91.47 \pm 0.15	90.52 \pm 0.35	84.49 \pm 0.05	82.34 \pm 0.27
<i>DMIS</i>	94.27\pm0.55	94.20\pm0.15	94.70\pm0.49	93.53\pm0.12	89.31\pm0.21	88.42\pm0.43

1741
 1742 **Table 6:** Classification results on Fashion-MNIST, CIFAR-10, and ImageNette datasets under various
 1743 types of supplementary-unlabeled supervision. **Bold** numbers indicate the best performance.

Method	Fashion-MNIST		CIFAR-10		ImageNette	
	Random-1%	Random-10%	Random-1%	Random-10%	Random-1%	Random-10%
<i>Pseudo-Labeling</i>	83.53 \pm 0.46	89.59 \pm 0.23	50.10 \pm 0.95	72.92 \pm 0.17	43.00 \pm 0.82	68.03 \pm 0.32
<i>Mean Teacher</i>	82.34 \pm 0.09	89.91 \pm 0.15	47.69 \pm 0.27	73.01 \pm 0.78	40.53 \pm 1.56	65.72 \pm 0.55
<i>VAT</i>	83.31 \pm 0.61	89.35 \pm 0.12	49.64 \pm 0.90	71.07 \pm 1.27	38.63 \pm 8.39	63.93 \pm 5.18
<i>UDA</i>	84.28 \pm 0.41	90.83 \pm 0.34	69.20 \pm 1.41	80.50 \pm 0.55	50.52 \pm 3.79	72.53 \pm 1.17
<i>FixMatch</i>	84.32 \pm 0.33	90.76 \pm 0.38	67.48 \pm 1.42	80.00 \pm 0.63	50.41 \pm 4.43	71.32 \pm 1.93
<i>Dash</i>	84.73 \pm 0.09	91.16 \pm 0.20	70.14 \pm 0.69	81.50 \pm 0.68	57.68 \pm 2.19	74.66 \pm 0.81
<i>CoMatch</i>	85.31 \pm 0.29	90.52 \pm 0.12	61.45 \pm 1.46	77.79 \pm 0.53	63.88 \pm 0.78	73.20 \pm 0.46
<i>FlexMatch</i>	84.43 \pm 0.30	90.69 \pm 0.03	70.72 \pm 0.93	81.35 \pm 0.48	61.39 \pm 0.70	73.08 \pm 0.13
<i>FreeMatch</i>	84.30 \pm 0.37	90.92 \pm 0.24	70.15 \pm 0.44	80.99 \pm 0.56	60.37 \pm 1.11	73.14 \pm 1.03
<i>SimMatch</i>	84.69 \pm 0.17	91.18 \pm 0.13	73.33 \pm 1.02	82.90 \pm 0.43	58.12 \pm 2.66	76.12 \pm 0.45
<i>SoftMatch</i>	84.72 \pm 0.23	91.22 \pm 0.11	73.24 \pm 0.82	88.66 \pm 0.60	58.50 \pm 2.31	75.75 \pm 0.25
<i>Vanilla</i>	78.37 \pm 3.72	90.50 \pm 1.00	53.49 \pm 0.15	85.13 \pm 0.12	49.55 \pm 0.99	74.70 \pm 0.53
<i>DMIS^{CE}</i>	82.92 \pm 0.17	91.07 \pm 0.18	75.40 \pm 0.54	89.85 \pm 0.08	62.64 \pm 0.24	71.39 \pm 0.45
<i>DMIS</i>	85.92\pm0.13	92.97\pm0.21	76.30\pm0.17	92.47\pm0.39	68.23\pm0.19	77.30\pm0.15

1757
 1758 **Table 7:** Classification results on Fashion-MNIST, CIFAR-10, and ImageNette datasets under various
 1759 types of noisy-label supervision. **Bold** numbers indicate the best performance.

Method	Fashion-MNIST		CIFAR-10		ImageNette	
	Sym-40%	Asym-40%	Sym-40%	Asym-40%	Sym-40%	Asym-40%
<i>CE</i>	76.18 \pm 0.26	82.01 \pm 0.06	67.22 \pm 0.26	76.98 \pm 0.42	58.43 \pm 0.77	71.81 \pm 0.38
<i>Co-learning</i>	90.85 \pm 0.63	84.10 \pm 2.01	84.97 \pm 0.53	80.36 \pm 1.09	76.16 \pm 0.96	75.37 \pm 0.49
<i>Co-teaching</i>	92.17 \pm 0.34	92.78 \pm 0.25	86.54 \pm 0.57	79.38 \pm 0.39	66.55 \pm 1.00	75.12 \pm 0.50
<i>Co-teaching+</i>	91.05 \pm 0.06	91.62 \pm 0.20	67.28 \pm 1.85	79.43 \pm 0.47	75.79 \pm 0.79	75.17 \pm 0.40
<i>SCE</i>	93.62 \pm 0.22	88.60 \pm 0.20	82.82 \pm 0.40	81.54 \pm 0.64	77.99 \pm 0.39	74.81 \pm 1.04
<i>GCE</i>	93.64 \pm 0.03	87.48 \pm 0.09	85.00 \pm 0.27	77.97 \pm 3.69	81.18 \pm 0.35	72.61 \pm 1.14
<i>Decoupling</i>	92.24 \pm 0.23	92.10 \pm 0.44	82.24 \pm 0.28	79.89 \pm 0.58	75.53 \pm 0.69	78.24 \pm 0.21
<i>ELR</i>	93.13 \pm 0.13	92.82 \pm 0.09	85.68 \pm 0.13	81.32 \pm 0.31	84.03 \pm 2.86	73.51 \pm 0.31
<i>JoCoR</i>	84.05 \pm 1.11	89.45 \pm 4.43	77.92 \pm 3.92	78.68 \pm 0.07	67.82 \pm 1.97	74.67 \pm 0.43
<i>Mixup</i>	92.21 \pm 0.03	92.01 \pm 1.02	84.26 \pm 0.64	83.21 \pm 0.85	76.65 \pm 1.62	77.16 \pm 0.71
<i>PENCIL</i>	90.85 \pm 0.58	91.77 \pm 0.69	85.91 \pm 0.26	84.89 \pm 1.49	81.94 \pm 1.26	77.20 \pm 1.15
<i>Vanilla</i>	90.11 \pm 1.24	85.41 \pm 0.96	80.22 \pm 0.10	86.31 \pm 0.10	55.86 \pm 1.95	53.91 \pm 1.07
<i>DMIS^{CE}</i>	82.76 \pm 0.57	83.39 \pm 0.24	84.75 \pm 0.36	84.21 \pm 0.18	80.47 \pm 0.56	77.21 \pm 0.19
<i>DMIS</i>	93.40\pm0.40	93.20\pm0.30	88.63\pm0.12	88.83\pm0.33	84.12\pm0.18	79.30\pm0.27

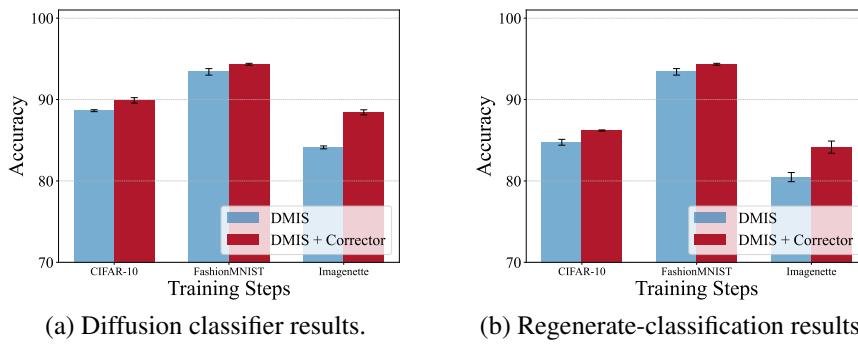
1773
 1774 **Noisy-label learning.** We further benchmark our method against nine widely used approaches:
 1775 *Coteaching* (Han et al., 2018), *Coteaching+* (Yu et al., 2019), *SCE* (Wang et al., 2019), *GCE* (Zhang
 1776 & Sabuncu, 2018), *Decoupling* (Malah & Shalev-Shwartz, 2017), *ELR* (Liu et al., 2020), and
 1777 *JoCoR* (Wei et al., 2020). These methods cover a range of strategies, from sample selection and
 1778 reweighting to robust loss design, thus providing a diverse and rigorous benchmark.

1779 Across all three weakly supervised scenarios, our method consistently achieves the best performance
 1780 compared to existing baselines, reinforcing both its robustness and versatility under different forms
 1781 of imprecise supervision.

1782 E.3 INTEGRATION WITH EXISTING IMPRECISE-LABEL CORRECTORS
1783

1784 Existing weakly supervised learning methods often rely on pseudo-labeling strategies that aim to
1785 correct imprecise labels by assigning refined labels to training samples. From this perspective, our
1786 approach is orthogonal to such methods: while pseudo-labeling seeks to approximate the true labels
1787 as closely as possible, our framework focuses on robustly learning from the remaining label noise. In
1788 practice, pseudo-labeling methods inevitably produce imperfect corrections. While most samples
1789 may be relabeled correctly, a non-negligible portion of instances still receive erroneous pseudo-labels
1790 because no classifier is perfect. As a result, imprecise supervision is effectively transformed into a
1791 noisy-label supervision.

1792 This naturally complements our framework: by combining a pseudo-label corrector with *DMIS*, one
1793 can first reduce label uncertainty through correction and then leverage the robustness of diffusion
1794 models to learn from the residual noise. To validate this premise, we conduct a case study where a
1795 noisy-label learning method trained on CIFAR-10 with 40% symmetric noise achieves a pseudo-label
1796 accuracy of 80% on the training set. Using this pseudo-labeled dataset as input, our *DMIS* framework
1797 further improves the classification performance. As illustrated in Figure 6, integrating pseudo-label
1798 correction with *DMIS* consistently improves the performance across all datasets. Thus, we believe
1799 that our framework addresses the challenge of imprecise labels through the lens of diffusion model,
1800 offering a complementary perspective to conventional noisy-label approaches.

1811 Figure 6: Test accuracy before and after applying pseudo-label correction with *DMIS*.
18121813 E.4 COMPARISON OF ACCURACY CURVES BETWEEN *DMIS* AND *Vanilla*
1814

1815 To better illustrate the difference between the *Vanilla* method and our proposed *DMIS*, we plot the test
1816 accuracy curves during training, as shown in Figure 7. Across all settings, the *Vanilla* model exhibits
1817 an initial rise in accuracy followed by a gradual decline as training progresses, suggesting that it
1818 struggles to maintain stable performance under prolonged training. In contrast, *DMIS* consistently
1819 sustains high accuracy throughout training, showing its robustness across diverse supervision types.

1820 Specifically, in the noisy-label setting, the decline of *Vanilla* is especially pronounced, reflecting
1821 its sensitivity to label corruption. In partial-label learning, *Vanilla* also exhibits instability, whereas
1822 *DMIS* maintains reliable performance. Even in semi-supervised learning, where labels are clean but
1823 scarce, *DMIS* achieves higher and more stable accuracy compared to *Vanilla*, demonstrating that our
1824 framework is not only noise-robust but also effective in leveraging limited supervision.

1825 E.5 VISUALIZATION OF NOISY CONDENSED DATASETS
1826

1827 We visualize the condensed images on CIFAR-10 and Fashion-MNIST in Figure 8 and Figure 9,
1828 respectively. It is evident that datasets generated by our method exhibit both higher diversity and
1829 stronger realism compared to other approaches. In particular, for the condensed Fashion-MNIST
1830 images, methods such as *DC* and *DM* often produce samples that do not faithfully correspond to
1831 their assigned class, resulting in condensed datasets that still contain noisy labels and thus degrade
1832 performance. By contrast, our proposed *DMIS* generates class-consistent and visually recognizable
1833 samples across categories, yielding condensed datasets that better preserve label fidelity and semantic
1834

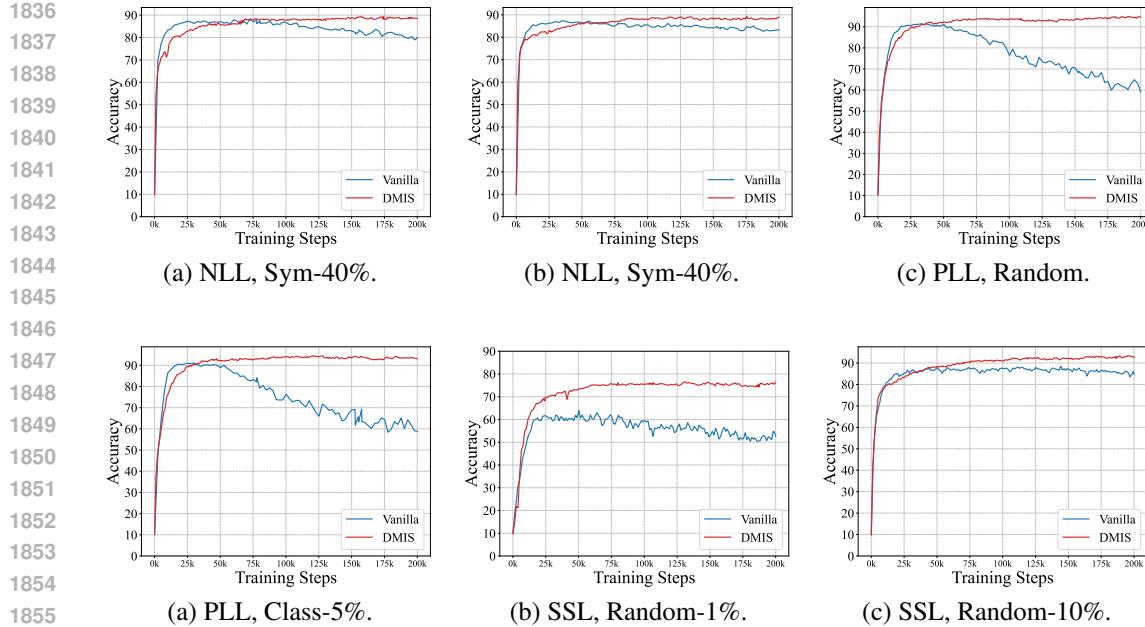


Figure 7: Test accuracy curves of the *Vanilla* and *DMIS* models on CIFAR-10 under different forms of imprecise supervision, including noisy-label learning (NLL), partial-label learning (PLL), and semi-supervised learning (SSL).

alignment. These visualizations further support the quantitative results, highlighting the advantage of generative condensation under noisy supervision.

E.6 ADDITIONAL RESULTS ON DATASET CONDENSATION UNDER DIFFERENT FORMS OF IMPRECISE SUPERVISION

To illustrate the extreme case of noisy dataset condensation, we report the results when the IPC is set to 1. As shown in Table 8, *DMIS* consistently achieves the best performance across all datasets and noise types, even under the extreme case of $IPC = 1$. Notably, while most existing condensation methods collapse under severe supervision noise, our method maintains a clear advantage, outperforming the strongest baselines by a large margin. These results further demonstrate the robustness of *DMIS* in distilling informative representations despite highly limited and imprecisely labeled data.

Table 8: Classification results (test accuracy, %) on noisy-label Fashion-MNIST, CIFAR-10, and ImageNette datasets. ‘IPC’ indicates the number of images per class in the condensed dataset. **Bold** numbers indicate the best performance.

Dataset	Type	IPC	DC	DSA	DM	MTT	RDED	SRE2L	DMIS
F-MNIST	Sym-40%	1	15.21 ± 0.75	19.55 ± 0.58	15.56 ± 0.20	10.86 ± 1.90	18.07 ± 3.33	14.33 ± 1.20	33.18 ± 2.15
	Asym-40%	1	20.17 ± 0.29	17.61 ± 0.89	23.91 ± 0.36	7.39 ± 0.84	13.20 ± 0.83	13.13 ± 0.21	25.78 ± 0.70
CIFAR-10	Sym-40%	1	8.99 ± 1.59	10.00 ± 0.00	14.41 ± 1.03	9.99 ± 0.00	11.20 ± 0.41	11.06 ± 0.83	11.81 ± 1.04
	Asym-40%	1	11.88 ± 1.55	10.00 ± 0.00	10.00 ± 0.00	9.96 ± 0.05	13.96 ± 1.38	15.49 ± 0.34	15.88 ± 0.57
ImageNette	Sym-40%	1	9.87 ± 0.00	9.87 ± 0.00	9.87 ± 0.00	19.17 ± 2.35	12.98 ± 1.16	11.90 ± 0.78	19.32 ± 0.84
	Asym-40%	1	9.87 ± 0.00	9.87 ± 0.00	9.87 ± 0.00	17.36 ± 0.10	12.98 ± 0.42	18.55 ± 2.19	21.13 ± 0.95

Furthermore, even when samples are imprecisely annotated with candidate labels, our method is still able to perform effective condensation on partial-label datasets. In contrast, most existing dataset condensation methods rely on the assumption of having single labels for each instance and therefore fail under this type of supervision. The only exception lies in decoupled condensation approaches such as *RDED* and *SRE2L*, where a teacher model can still be trained on partial-label data. We present the corresponding results under partial-label supervision below in Table 9.

As shown in the table, our method consistently achieves substantial improvements across both Random and Class-50% candidate set generation strategies, and under all IPC settings. In particular, on Fashion-MNIST, our approach yields dramatic performance gains, reaching above **87%** accuracy even with partial label supervision, whereas both *RDED* and *SRe2L* fail to exceed 16% under the same setting. On the more challenging CIFAR-10 benchmark, our method also demonstrates strong robustness, especially under larger IPCs where the gap over baseline methods becomes increasingly pronounced (e.g., over **20%** absolute improvement at $IPC = 100$). These results highlight that our condensation strategy can effectively leverage weak supervision and generate compact yet highly informative synthetic datasets, even when label noise is introduced by the partial-label scenario.

Table 9: Classification results (test accuracy, %) on partial-label Fashion-MNIST and CIFAR-10 datasets under different IPCs. ‘Random’ and ‘Class-50%’ denote two candidate set generation strategies. **Bold** numbers indicate the best performance.

Dataset	Random				Class-50%			
	IPC	<i>RDED</i>	<i>SRe2L</i>	Ours	IPC	<i>RDED</i>	<i>SRe2L</i>	Ours
F-MNIST	1	10.48 ± 0.82	9.72 ± 1.02	44.06 ± 1.64	1	10.73 ± 0.78	10.93 ± 0.76	33.99 ± 2.48
	10	13.17 ± 4.66	8.80 ± 0.70	72.02 ± 0.77	10	14.58 ± 1.19	10.83 ± 0.35	70.46 ± 0.26
	50	13.06 ± 2.36	10.30 ± 0.40	83.98 ± 0.12	50	15.55 ± 0.55	11.33 ± 0.76	79.67 ± 0.13
	100	13.39 ± 2.40	9.34 ± 1.77	87.30 ± 0.31	100	11.90 ± 2.67	11.05 ± 0.77	81.42 ± 0.25
CIFAR-10	1	15.32 ± 1.72	20.69 ± 0.88	16.31 ± 1.54	1	14.52 ± 1.06	15.55 ± 0.91	16.61 ± 1.14
	10	26.28 ± 0.29	19.45 ± 1.12	30.50 ± 0.27	10	10.00 ± 0.00	18.56 ± 0.05	25.00 ± 0.70
	50	34.96 ± 0.92	20.56 ± 0.71	44.39 ± 0.65	50	25.59 ± 1.32	19.39 ± 1.09	45.94 ± 1.27
	100	28.88 ± 2.56	19.65 ± 1.09	58.46 ± 0.64	100	25.81 ± 1.64	18.65 ± 1.51	58.06 ± 0.32

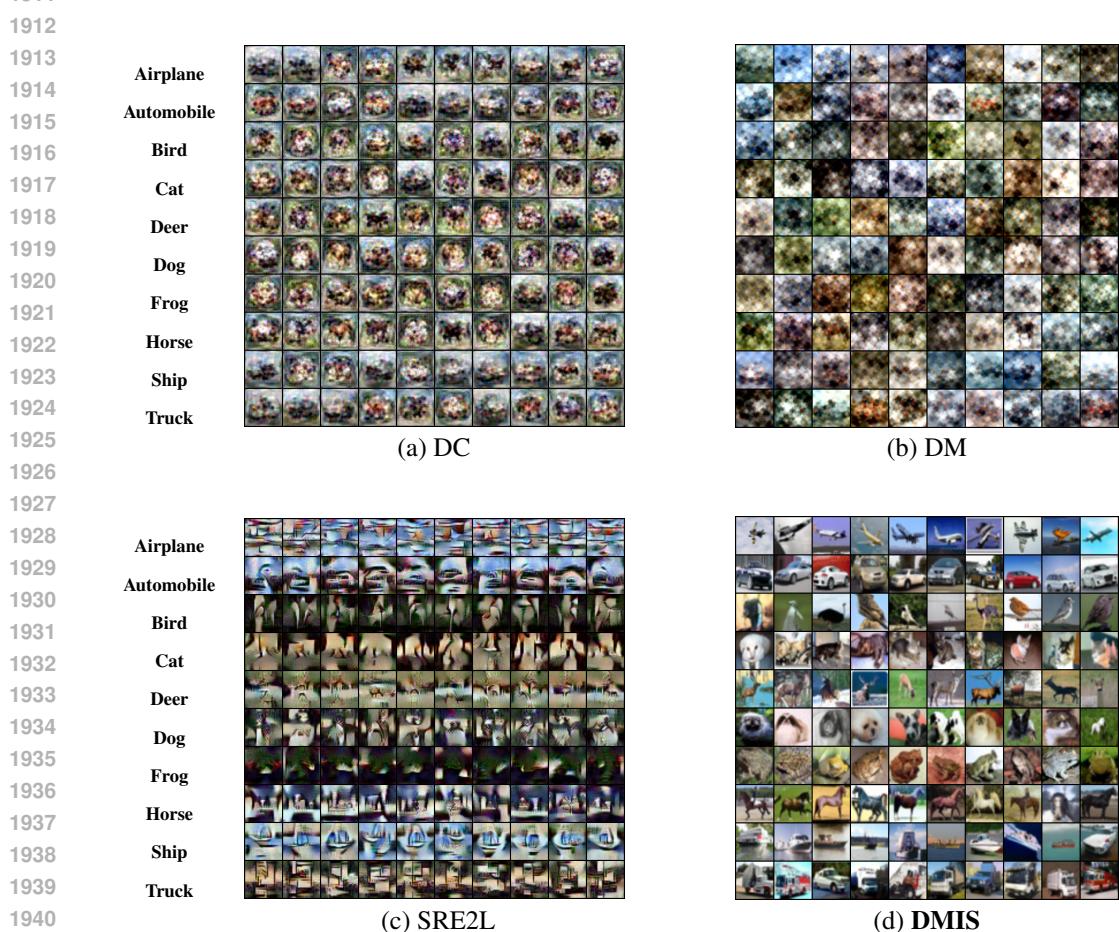


Figure 8: Visualization of condensed CIFAR-10 images generated by *DC*, *DM*, *SRE2L*, and our method ***DMIS***.

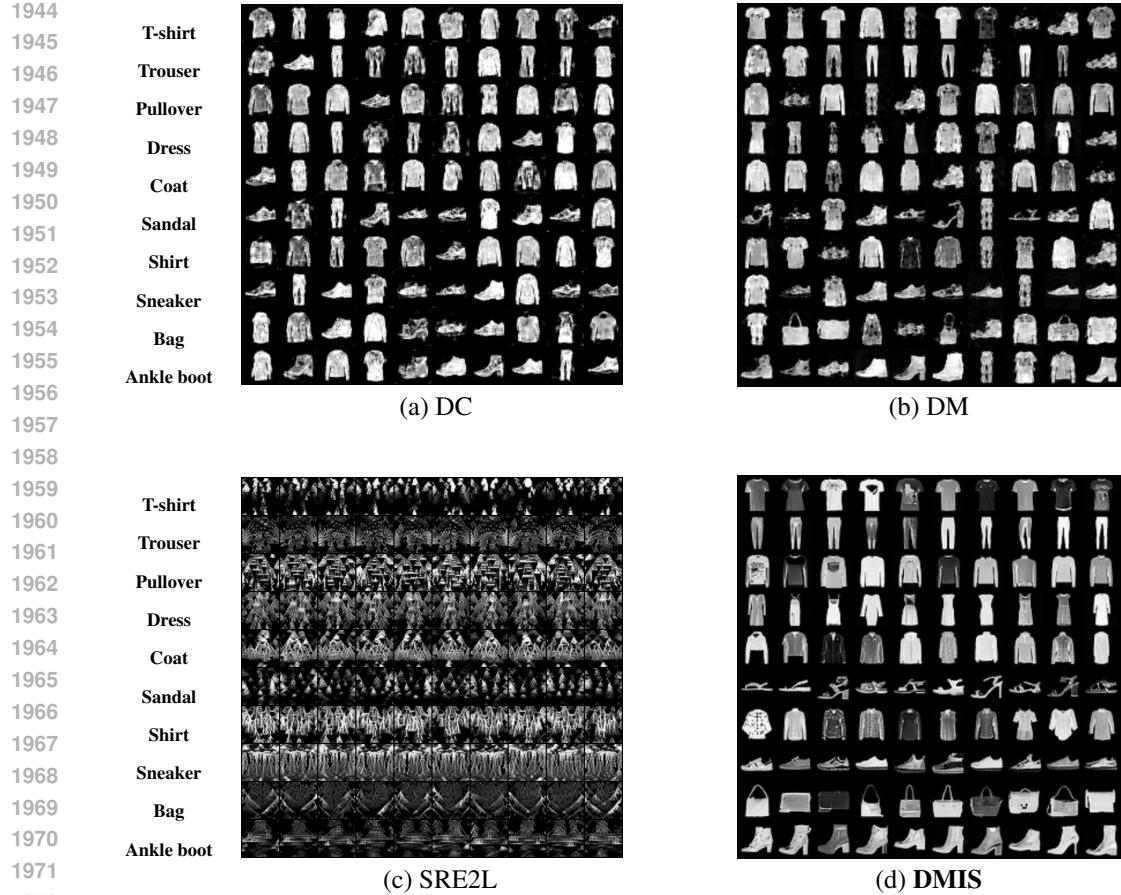


Figure 9: Visualization of condensed Fashion-MNIST images generated by *DC*, *DM*, *SRE2L*, and our method ***DMIS***.

F ADDITIONAL EXPERIMENTS RESULTS

F.1 THE FULL RESULTS OF TABLE 1.

Table 10: Generative results on CIFAR-10 and ImageNette under various settings. ‘uncond’ and ‘cond’ indicate unconditional and conditional metrics. **Bold** numbers indicate better performance.

Metric		Noisy-label supervision				Partial-label supervision				Suppl-unlabeled supervision				Clean	
		Sym-40%		Asym-40%		Random		Class-50%		Random-1%		Random-10%			
		Vanilla	DMIS	Vanilla	DMIS	Vanilla	DMIS	Vanilla	DMIS	Vanilla	DMIS	Vanilla	DMIS	Vanilla	DMIS
CIFAR-10	cond	FID (\downarrow)	3.33 \pm 0.06	3.47 \pm 0.11	3.23 \pm 0.07	3.10 \pm 0.11	7.76 \pm 0.25	2.26 \pm 0.08	11.75 \pm 0.42	2.77 \pm 0.09	3.16 \pm 0.07	3.12 \pm 0.05	2.93 \pm 0.11	2.89 \pm 0.08	2.05 \pm 0.05
	cond	IS (\downarrow)	9.56 \pm 0.08	9.68 \pm 0.05	9.02 \pm 0.06	9.73 \pm 0.06	9.09 \pm 0.15	9.80 \pm 0.04	9.62 \pm 0.11	9.68 \pm 0.07	10.03 \pm 0.09	10.57 \pm 0.06	9.80 \pm 0.08	9.83 \pm 0.05	10.61 \pm 0.04
	cond	Density (\uparrow)	101.39 \pm 0.85	109.75 \pm 0.62	100.06 \pm 0.91	109.69 \pm 0.55	103.21 \pm 1.12	106.49 \pm 0.48	108.76 \pm 0.75	109.06 \pm 0.52	97.19 \pm 1.05	108.18 \pm 0.66	99.96 \pm 0.88	108.87 \pm 0.59	112.59 \pm 0.45
	cond	Coverage (\uparrow)	81.12 \pm 0.35	81.21 \pm 0.28	80.71 \pm 0.41	81.30 \pm 0.32	68.45 \pm 0.65	82.69 \pm 0.25	64.90 \pm 0.72	81.56 \pm 0.28	78.25 \pm 0.32	81.85 \pm 0.38	80.85 \pm 0.38	82.24 \pm 0.27	83.27 \pm 0.22
	cond	CW-FID (\downarrow)	29.12 \pm 0.35	29.12 \pm 0.45	29.12 \pm 0.35	29.12 \pm 0.35	27.04 \pm 0.37	30.32 \pm 0.37	32.56 \pm 0.41	16.25 \pm 0.65	16.12 \pm 0.55	11.19 \pm 0.51	11.77 \pm 0.42	9.93 \pm 0.55	11.70 \pm 0.52
	cond	CW-Density (\downarrow)	99.98 \pm 0.82	107.23 \pm 0.65	90.85 \pm 0.75	107.07 \pm 0.58	70.04 \pm 0.92	106.75 \pm 0.61	102.43 \pm 0.88	108.66 \pm 0.55	100.73 \pm 0.68	96.29 \pm 0.72	107.94 \pm 0.62	111.70 \pm 0.52	111.70 \pm 0.52
ImageNette	cond	CW-Coverage (\uparrow)	75.39 \pm 0.45	75.39 \pm 0.35	79.63 \pm 0.42	79.65 \pm 0.38	65.45 \pm 0.68	82.09 \pm 0.32	61.45 \pm 0.36	75.03 \pm 0.52	71.24 \pm 0.36	75.84 \pm 0.41	80.80 \pm 0.45	81.12 \pm 0.39	83.91 \pm 0.30
	cond	FID (\downarrow)	14.11 \pm 0.55	13.44 \pm 0.48	13.93 \pm 0.52	13.91 \pm 0.45	79.13 \pm 2.15	72.62 \pm 1.85	91.28 \pm 2.45	79.12 \pm 1.95	23.88 \pm 0.85	19.26 \pm 0.72	14.32 \pm 0.58	12.84 \pm 0.46	11.52 \pm 0.42
	cond	IS (\downarrow)	12.69 \pm 0.13	13.21 \pm 0.12	12.51 \pm 0.14	13.73 \pm 0.11	9.19 \pm 0.25	9.40 \pm 0.18	9.27 \pm 0.22	9.11 \pm 0.24	13.23 \pm 0.16	13.72 \pm 0.13	12.80 \pm 0.15	13.16 \pm 0.12	13.81 \pm 0.10
	cond	Density (\uparrow)	101.31 \pm 0.98	112.52 \pm 0.82	111.66 \pm 0.88	106.78 \pm 0.92	95.33 \pm 1.25	99.83 \pm 0.75	94.29 \pm 1.32	102.58 \pm 0.85	115.94 \pm 1.15	125.68 \pm 0.95	105.27 \pm 0.85	109.23 \pm 0.78	117.23 \pm 0.65
	cond	Coverage (\uparrow)	76.62 \pm 0.42	76.81 \pm 0.38	78.32 \pm 0.45	79.81 \pm 0.35	21.44 \pm 0.85	32.48 \pm 0.55	16.69 \pm 0.92	22.30 \pm 0.65	53.53 \pm 0.72	55.39 \pm 0.58	73.79 \pm 0.48	75.55 \pm 0.42	80.12 \pm 0.35
	cond	CW-FID (\downarrow)	80.31 \pm 0.28	60.12 \pm 0.85	62.26 \pm 1.95	58.63 \pm 1.65	157.76 \pm 0.95	63.58 \pm 2.15	163.45 \pm 1.5	67.84 \pm 0.25	71.24 \pm 2.15	76.12 \pm 0.25	99.47 \pm 0.34	94.37 \pm 0.25	103.35 \pm 0.35
CIFAR-10	cond	CW-Density (\downarrow)	99.98 \pm 0.85	99.98 \pm 0.85	99.98 \pm 0.85	99.98 \pm 0.85	99.98 \pm 0.85	99.98 \pm 0.85	99.98 \pm 0.85	99.98 \pm 0.85	99.98 \pm 0.85	99.98 \pm 0.85	99.98 \pm 0.85	99.98 \pm 0.85	99.98 \pm 0.85
	cond	CW-Coverage (\uparrow)	67.89 \pm 0.55	71.94 \pm 0.45	74.18 \pm 0.52	75.82 \pm 0.48	19.76 \pm 0.85	24.35 \pm 0.55	15.88 \pm 0.92	18.93 \pm 0.62	51.73 \pm 0.75	52.15 \pm 0.65	72.61 \pm 0.58	74.85 \pm 0.52	78.48 \pm 0.45

F.2 COMPARISON AGAINST OTHER NOISE-ROBUST DIFFUSION METHODS.

We have also shown comparisons against noise-robust diffusion methods (Na et al., 2024; Dufour et al., 2024) that are designed to handle noisy-label data. For image generation task, we compare them and our DMIS model on CIFAR-10 with 40% symmetric and asymmetric label noise, reporting FID, IS under the same architecture and training budget. For noisy-label learning task, we compare them and our DMIS as data generators for downstream classification. We use each model to synthesize the same number of labeled samples and then train a Wide-ResNet-40-10 classifier on top of these

	Metric	Sym-40%			Asym-40%		
		FID	IS	Accuracy	FID	IS	Accuracy
2001	DMIS (Ours)	3.47	9.80	88.63	3.10	9.73	88.83
2002	CAD (Dufour et al., 2024)	4.10	9.68	81.75	3.87	9.16	82.33
2003	TDSM (Na et al., 2024)	3.85	9.40	66.40	3.96	10.12	72.32

Table 11: Results under 40% symmetric and asymmetric noise.

synthetic datasets. Importantly, both two compared methods assume access to additional prior information, which can give them an advantage in this setting. Despite this, our method still achieves the best overall performance under the same backbone and training budget. This suggests that our approach is competitive while relying on strictly weaker assumptions about the available supervision.

F.3 TOP- k TRUNCATION FOR LARGE LABEL SPACES

When extending to datasets with many classes, a straightforward implementation becomes expensive because it requires estimating and backpropagating a per-class objective at every step. To reduce this cost in practice, we restrict gradients to classes that carry non-negligible probability mass.

Concretely, we apply a top- k strategy to both the diffusion posterior and the pseudo-label distribution: only the k largest entries are retained, while all remaining entries are zero-masked and do not contribute to the gradient. In this way, the effective complexity scales with the number of active classes k per example, rather than with the total number of classes.

To assess the impact of this approximation, we conduct an experiment on the Caltech-15 dataset (Pan et al., 2023) with 40% symmetric label noise and set $k = 10$. As shown below, the top- k variant achieves performance comparable to the full model while reducing computational cost, indicating that this strategy is a practical mechanism for scaling our method to larger label spaces.

	Generation Metric				Classification Metric	
	FID	IS	Density	Coverage		
2029	DMIS (Ours)	4.25	12.39	103.83	96.20	78.92

Table 12: Performance of DMIS under Caltech-15 dataset with 40% symmetric noise.

F.4 EXPERIMENTS BEYOND SYNTHETIC CLASS-CONDITIONAL NOISE

We primarily use synthetic noisy labels to obtain a controlled setting that supports our theory, where both the noise rate and the noise type (e.g., symmetric, class-dependent) can be precisely specified.

Starting from such controlled synthetic-noise regimes is a necessary first step to validate both the theoretical predictions and the basic empirical behavior of our method. To further demonstrate its practicality under more realistic supervision, we also evaluate DMIS on real noisy-label and partial-label benchmarks. Specifically, we report results on the real noisy-label dataset CIFAR-10N (Wei et al., 2022) and the real partial-label dataset PLCIFAR-10 (Wang et al., 2025b), whose labels are provided by human annotators.

In addition, we consider instance-dependent label noise on CIFAR-10, following standard instance-dependent noise protocols in the noisy-label literature. The results below show that DMIS maintains strong generative quality and competitive classification accuracy under these more complex and realistic noise conditions.

2052
 2053
 2054
 2055
 2056
 2057
 2058
 2059
 2060
 2061
 2062
 2063
 2064
 2065
 2066
 2067
 2068
 2069
 2070
 2071
 2072
 2073
 2074
 2075

Dataset	FID	IS	Test accuracy (%)
CIFAR10-N	3.22	9.66	93.21
Instance-dependent CIFAR10	4.85	9.21	81.32
PLCIFAR10	2.95	9.82	93.65

2081 Table 13: Performance of DMIS under more complex imprecise supervision datasets.
 2082

2083
 2084
 2085
 2086
 2087
 2088
 2089
 2090
 2091
 2092
 2093
 2094
 2095
 2096
 2097
 2098
 2099
 2100
 2101
 2102
 2103
 2104
 2105